



ExtremeXOS Release Notes

Software Version ExtremeXOS 15.5.5-Patch1-1

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 15.5.2	7
Re-authentication Using Simple Network Management Protocol (SNMP) for Network Management System (NMS) Operation	8
Supported Platforms	8
Multi-chassis Link Aggregation Group—Link Aggregation Control Protocol (MLAG-LACP) Enhancements	8
Supported Platforms	8
New and Corrected Features in ExtremeXOS 15.5.1	9
Support Linux File System	10
Supported Platforms	10
New CLI Commands	11
Changed CLI Commands	11
Non-Extreme Networks Optics Licensing	12
Changed CLI Commands	13
Ethernet Ring Protection Switching (ERPS) Using Y.1731 Continuity Check Messages (CCMs)	14
Supported Platforms	15
Limitations	15
Return-to-Normal Simple Network Management Protocol (SNMP) Notifications	15
Supported Platforms	15
Simple Network Management Protocol (SNMP) Notification Logs	16
Supported RFCs	16
Supported Platforms	16
Limitations	16
New CLI Commands	17
Bidirectional Forwarding Detection (BFD) Up/Down Traps	18
Supported Drafts	18
Supported Platforms	18
New CLI Commands	18
Changed CLI Commands	19
Internet Protocol (IP) v6 Multi-cast Listener Discovery Protocol (MLD) Source-Specific Multi-cast (SSM) Map	19
Supported Platforms	19
Limitations	19
New CLI Commands	20
Changed CLI Commands	20
Access Control List (ACL) Enhancements	21
Enhance Clear Counters Per Port Command	21
Supported Platforms	21
Changed CLI Commands	21

Multi-switch Link Aggregation Group (MLAG) Out-of-Band Keep-Alive Protocol	22
Supported Platforms	22
Limitations	22
New CLI Commands	22
Changed CLI Commands	23
Multi-switch Link Aggregation Group (MLAG) Support for More than One Peer	23
Supported Platforms	23
Limitations	23
Multi-switch Link Aggregation Group (MLAG) MD5 Hash for TCP Checkpointing Connecting	24
Supported Platforms	24
Limitations	24
New CLI Commands	24
Changed CLI Commands	24
Layer 2 Protocol Tunneling (L2PT) and Filtering	25
Supported Platforms	25
Limitations	25
New CLI Commands	26
Changed CLI Commands	26
Y.1731 Compliant Performance Monitoring SNMP MIBs	27
Supported Platforms	27
Limitations	27
Dynamic Host Configuration Protocol (DHCPv6) RFC4649 Relay Agent Remote-ID Option	28
Supported Platforms	28
Multiprotocol Label Switching (MPLS) Pseudowire—Label-Switched Path (PW-LSP) MIB Counters	29
Joint Interoperability Test Command (JITC) Compliance	29
New Hardware Supported in ExtremeXOS 15.5.2	30
New Hardware Supported in ExtremeXOS 15.5.1	30
ExtremeXOS Hardware and Software Compatibility Matrix	30
Upgrading to ExtremeXOS	31
Downloading Supported MIBs	31
Tested Third-Party Products	31
Tested RADIUS Servers	31
Tested Third-Party Clients	32
PoE Capable VoIP Phones	32
Extreme Switch Security Assessment	33
DoS Attack Assessment	33
ICMP Attack Assessment	33
Port Scan Assessment	33
Service Notifications	33

Limits.....35

Supported Limits	35
------------------------	----

Open Issues, Known Behaviors, and Resolved Issues	83
Open Issues	86
Known Behaviors	91
Resolved Issues in ExtremeXOS 15.5.5-Patch1-1	92
Resolved Issues in ExtremeXOS 15.5.5	95
Resolved Issues in ExtremeXOS 15.5.4-Patch1-5	98
Resolved Issues in ExtremeXOS 15.5.4-Patch1-4	99
Resolved Issues in ExtremeXOS 15.5.4	103
Resolved Issues in ExtremeXOS 15.5.3-Patch1-6	106
Resolved Issues in ExtremeXOS 15.5.3-Patch1-5	107
Resolved Issues in ExtremeXOS 15.5.3-Patch1-2	109
Resolved Issues in ExtremeXOS 15.5.3	111
Resolved Issues in ExtremeXOS 15.5.2-Patch1-5	113
Resolved Issues in ExtremeXOS 15.5.2-Patch1-1	115
Resolved Issues in ExtremeXOS 15.5.2	117
Resolved Issues in ExtremeXOS 15.5.1	122
ExtremeXOS Documentation Corrections.....	135
ACL Egress Counters Limitation	136
ACLs General	136
ACL Policy Redirect	137
Configure Access-List VLAN-ACL-Precedence Command Usage Guidelines	137
Configure IP-MTU VLAN Command Syntax Description	138
Command Reference	138
User Guide	138
Debounce Commands	139
Configure stack-ports debounce time	139
Description	139
Syntax Description	139
Default	139
Usage Guidelines	139
Example	139
History:	139
Platforms Availability	139
Show stack-ports debounce	140
Description	140
Syntax Description	140
Default	140
Usage Guidelines	140
Example	140
History	140
Platform Availability	140
ELRP	141
ELRP and QoS	142
L2VPN Sharing Commands	142
Link Aggregation (LAG) Limit for Multiprotocol Label Switching (MPLS) Terminated Packets	143

Mirroring	143
MLAG	144
MLAG PIM-SM *,G Forwarding Limitation	144
NetLogin	144
PoE Power Delivery	145
Policies and Security	145
Rate Limiting/Meters	146
Routing Policies	147
Synchronize Command	148
TACACS Server	149
Unconfigure Switch Erase Command	151
VRRP	152
VRRP Guidelines	153

1 Overview

These release notes document ExtremeXOS® 15.5.5-Patch1-1 which resolves software deficiencies.

This chapter contains the following sections:

- [New and Corrected Features in ExtremeXOS 15.5.2 on page 7](#)
- [New and Corrected Features in ExtremeXOS 15.5.1 on page 9](#)
- [Joint Interoperability Test Command \(JITC\) Compliance on page 29](#)
- [New Hardware Supported in ExtremeXOS 15.5.2 on page 30](#)
- [New Hardware Supported in ExtremeXOS 15.5.1 on page 30](#)
- [ExtremeXOS Hardware and Software Compatibility Matrix on page 30](#)
- [Upgrading to ExtremeXOS on page 31](#)
- [Downloading Supported MIBs on page 31](#)
- [Tested Third-Party Products on page 31](#)
- [Extreme Switch Security Assessment on page 33](#)
- [Service Notifications on page 33](#)


New and Corrected Features in ExtremeXOS 15.5.2

This section lists the new and corrected features supported in the ExtremeXOS 15.5.2 software:



NOTE

For ExtremeXOS 15.5.2, Identity Management, XNV, and NTP are not supported on the Summit X430 series switches, even with an L2 Edge license.

-
- [Re-authentication Using Simple Network Management Protocol \(SNMP\) for Network Management System \(NMS\) Operation on page 8](#)
 - [Multi-chassis Link Aggregation Group—Link Aggregation Control Protocol \(MLAG-LACP\) Enhancements on page 8](#)
- 

Re-authentication Using Simple Network Management Protocol (SNMP) for Network Management System (NMS) Operation

This feature adds a proprietary MIB (EXTREME-MAC-AUTH-MIB) to ExtremeXOS that allows an NMS to force re-authentication of clients authenticated using MAC-based or dot1x authentication.

Supported Platforms

All platforms

Multi-chassis Link Aggregation Group—Link Aggregation Control Protocol (MLAG-LACP) Enhancements

When MLAG LACP is configured, if both MLAG peers go down and then one of them never boots up, the connection between the remote node and the single MLAG peer stays logically down. This feature requires the MLAG peer in active state to send LACP PDUs with its MAC address or configured LACP addresses even when MLAG peering is not established, thus maintaining the connection between remote node and single MLAG peer.

It is recommended that you configure LACP MAC on both MLAG peers to minimize the traffic disruption when the second MLAG peer comes up.

Supported Platforms

All platforms



New and Corrected Features in ExtremeXOS 15.5.1

This section lists new and corrected features in the ExtremeXOS 15.5.1 software:

- Support Linux File System on page 10
- Non-Extreme Networks Optics Licensing on page 12
- Ethernet Ring Protection Switching (ERPS) Using Y.1731 Continuity Check Messages (CCMs) on page 14
- Return-to-Normal Simple Network Management Protocol (SNMP) Notifications on page 15
- Simple Network Management Protocol (SNMP) Notification Logs on page 16
- Bidirectional Forwarding Detection (BFD) Up/Down Traps on page 18
- Internet Protocol (IP) v6 Multi-cast Listener Discovery Protocol (MLD) Source-Specific Multi-cast (SSM) Map on page 19
- Access Control List (ACL) Enhancements on page 21
- Enhance Clear Counters Per Port Command on page 21
- Multi-switch Link Aggregation Group (MLAG) Out-of-Band Keep-Alive Protocol on page 22
- Multi-switch Link Aggregation Group (MLAG) Support for More than One Peer on page 23
- Multi-switch Link Aggregation Group (MLAG) MD5 Hash for TCP Checkpointing Connecting on page 24
- Layer 2 Protocol Tunneling (L2PT) and Filtering on page 25
- Y.1731 Compliant Performance Monitoring SNMP MIBs on page 27
- Dynamic Host Configuration Protocol (DHCPv6) RFC4649 Relay Agent Remote-ID Option on page 28
- Multiprotocol Label Switching (MPLS) Pseudowire—Label-Switched Path (PW-LSP) MIB Counters on page 29



Support Linux File System

This feature introduces the concept of the “current working directory” relative to an ExtremeXOS login session, which can be modified and viewed with the `cd` (change directory) and `pwd` commands (print working directory).

The user “root” directory is `/usr/local` and you cannot navigate above this directory. The user “root” file system’s file system’s contents can vary depending on the platform and availability of removable or non-removable storage devices. Only the exposed file systems appear in this root directory, and the ExtremeXOS-internal file systems are not visible or accessible.

Under the `local` directory there are three standard sub-directories:

Folder Name	Contents	Corresponds to Deprecated Directory
<code>cfg</code>	Configuration files	NOTE: <ul style="list-style-type: none"> In earlier ExtremeXOS versions, equivalent to the current directory. Currently, default directory when you log on.
<code>ext</code>	External memory (for example, USB memory stick)	<code>memorycard</code>
<code>tmp</code>	Stores files temporarily,—typically core files (process crash dumps, etc.) and trace files	<code>internal-memory</code>

Sample root directory for existing and new file systems:

```
x480-48t(SS).22 # ls /usr/local
drwxrwxr--    5 root  rw                0 Apr 23 15:00 cfg
drwxr-xr-x    9 root  root            8192 Jan  1 1970 ext
drwxr-xr-x    2 root  root            2048 Apr 23 15:00 tmp
```

Supported Platforms

All platforms.

New CLI Commands

- `pwd` (print working directory)
- `cd` (change directory)
- `mkdir` (make directory)
- `rmdir` (remove directory)

Changed CLI Commands

With the deprecation of the keywords (shown in **bold**), these commands are analogous to their same-named Linux/UNIX shell commands:

- `ls { [internal-memory | memorycard] } { <file-name> }`
- `rm { [internal-memory | memorycard] } <file-name>`
- `cp [<old-name> <new-name>`
 - | **internal-memory** <old-name-internal> **memorycard** <new-name-memorycard>
 - | **internal-memory** <old-name-internal> **internal-memory** <new-name-internal>
 - | **memorycard** <old-name-memorycard> **memorycard** <new-name-memorycard>
 - | **memorycard** <old-name-memorycard> <new-name>
 - | <old-name> **memorycard** {<new-name-memorycard>}]
- `mv [<old-name> <new-name>`
 - | **internal-memory** <old-name-internal> **memorycard** <new-name-memorycard>
 - | **internal-memory** <old-name-internal> **internal-memory** <new-name-internal>
 - | **memorycard** <old-name-memorycard> **memorycard** <new-name-memorycard>
 - | **memorycard** <old-name-memorycard> <new-name>
 - | <old-name> **memorycard** {<new-name-memorycard>}]



```

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

```

Non-Extreme Networks Optics Licensing

Starting with ExtremeXOS 15.4, Extreme Networks began issuing a purchasable software license, “non-Extreme Optics” to use optical devices from third-party vendors on 40G and 100G ports. Previously, this feature only issued warning messages, but starting with ExtremeXOS 15.5, this feature now actually restricts a port’s egress rate to 25% of capacity when attached to an unlicensed device.



NOTE

This feature restriction does not apply to stacking ports.

Like other feature licenses, once the license is applied to a specific switch, that switch is permanently enabled to allow the unrestricted use of non-Extreme Networks optical devices. Without the license, ports which are attached to such devices are continuously restricted.

The devices subject to restriction are:

- QSFP+ SR4
- QSFP+ LR4
- CFP2 100G LR4
- CFP2 100G SR10
- QSFP+ passive copper 0.5m cable
- QSFP+ passive copper 1m cable
- QSFP+ passive copper 2m cable
- QSFP+ passive copper 3m cable
- QSFP+ passive copper 5m cable
- QSFP+ active optical 5m cable
- QSFP+ active optical 10m cable
- QSFP+ active optical 20m cable
- QSFP+ active optical 50m cable
- QSFP+ active optical 100m cable
- QSFP+ to SFP+ fan out passive copper 1m cable—QSFP+ end
- QSFP+ to SFP+ fan out passive copper 2m cable—QSFP+ end
- QSFP+ to SFP+ fan out passive copper 3m cable—QSFP+ end
- QSFP+ to SFP+ fan out passive copper 5m cable—QSFP+ end

Changed CLI Commands

The `show ports configuration` and `show ports information detail` commands include flags showing the summary status of non-Extreme Networks optical devices:

- Licensed—the optics license is installed (the 40G+ device is unrestricted, but is still unsupported by Extreme Networks TAC).
- Restricted—the non-Extreme Networks 40G+ device is restricted (rate limited).
- Unsupported—the non-Extreme Networks device is unsupported by Extreme Networks, but is unrestricted because it does not require a license (1G/10G).



Ethernet Ring Protection Switching (ERPS) Using Y.1731 Continuity Check Messages (CCMs)

This feature supports ERPS/G.8032 using Connectivity Fault Management (CFM) that sends Y.1731 Continuity Check Messages (CCMs) to detect connectivity failures, and thus link status. This is in addition to the current support for IEEE 802.1ag CCMs.

CFM functionality is the same in both specifications (Y.1731 and IEEE 802.1ag). Protocol Entities are also the same, but referred to with different terminology, except that Y.1731 does not use Maintenance Domain Name. [Table 1](#) shows the terminology differences.

Table 1: Y.1731 vs. IEEE 802.1ag Terminology

ITU-T G.8013/ Y.1731 Term	IEEE 802.1ag Term	Comments
MEG	MA	
MEG ID	MAID (Domain Name + Short MA Name)	Unlike in IEEE 802.1ag, the MEG ID does not imply a split between domain name and a short MEG name in ITU-T Y.1731.
MEG level	MA level	

MEG ID Formats:

- 32:ICC-based
- 33–63:Reserved for future use by ITU-T

Table 2: CCM PDU Format in IEEE 802.1ag vs. ITU-T Y.1731

Bytes	IEEE 802.1ag CCM PDU Format	ITU-T Y.1731 PDU Format
Byte 1	Domain level (3 bits) + CCM version (5 bits)	MEG level (3 bits) + CCM version (5 bits, 0 always)
Byte 2	OpCode (1 for CCM)	OpCode (1 for CCM)
Byte 3	Flags [RDI (1 bit) + reserved (5 bits) + CCM Interval (3 bits)]	Flags [RDI (1 bit) + reserved (5 bits) + CCM period (3 bits)]
Byte 4	First TLV offset	TLV offset (set to 70)
Byte 5–8	Sequence number	Sequence number (all zeros)
Byte 9–10	MEP ID	MEP ID
Byte 11–58	MAID MA details (format + length + value) with/without MD details (format + length + value)	MEG ID [reserved/no domain present (1 byte - 01) + MEG format (1 byte) + MEG length (1 byte) + MEG value (45 bytes)]

Table 2: CCM PDU Format in IEEE 802.1ag vs. ITU-T Y.1731 (Continued)

Bytes	IEEE 802.1ag CCM PDU Format	ITU-T Y.1731 PDU Format
Byte 59–74	Defined by ITU-T Y.1731	TxFCf (4 bytes) + RxFCb (4 bytes) + TxFCb (4 bytes) + reserved (4 bytes) Used in performance monitoring
Byte 75	End TLV or optional TLVs (Sender ID IP Address, Port Status, Interface Status, Organization Specific TLV) start here, if any	End TLV (0)
Byte 76 and onward	Optional TLVs continue, if present	—
Last byte	End TLV (if they are present prior optional TLVs)	—

To support Y.1731 CCMs, the association is provided with an additional ICC-based MEG name format which is user-configurable through the command line interface, as well as IEEE 802.1ag MIB. MD level is used as MEG level. All MEPs in this association transmit Y.1731 CCMs and are not allowed to be configurable for sender ID IP address, because optional TLVs are not supported.

Supported Platforms

- CFM in VLAN, VMAN is supported on all platforms
- Hardware MEPs on E4G-200/E4G- 400 cell site routers, and Summit X460 platforms can also support Y.1731 CCMs.

Limitations

- ISID-based association is not supported
- ISID-based MEP is not supported

Return-to-Normal Simple Network Management Protocol (SNMP) Notifications

This feature provides the ability for network management software (Ridgeline or NetSight) to receive “return-to-normal” SNMP notifications for CPU utilization or “overheated or too cold” from ExtremeXOS devices.

Supported Platforms

All platforms.



Simple Network Management Protocol (SNMP) Notification Logs

This feature provides a log of notifications sent by the SNMP agent that can be queried by a network management software (Ridgeline or NetSight). You can create multiple SNMP notifications logs, restrict what is added to the logs, age log entries, limit the maximum number of entries, and control these features using either the command line interface or SNMP.

Supported RFCs

- RFC 3413—SNMP Applications
- RFC 3014—Notification log MIB

Supported Platforms

All platforms.

Limitations

- Cannot query log entries by time duration (for example, list log entries from the last hour).
- You cannot create a log named “default”. The notification log name “default” is reserved to represent the default log in command line interface.
- Aging out of entries may occur sooner or later than the specified global age-out period if the current time of the device running ExtremeXOS changes.
- Notification log statistics (but not entries) is lost after a restart of the SNMP master process.
- Notification log statistics (but not entries) is lost on failover.
- Notification log entries and statistics are lost if the device running ExtremeXOS is rebooted.
- The following scalability constraints apply:
 - The maximum number of logs that can be created at a time is 16.
 - The maximum number of notifications that can be logged is 16,000 entries.



New CLI Commands

- `configure snmp notification-log [global-entry-limit <global_entry_limit> | global-age-out [none | <minutes>]]`
- `configure snmp add notification-log [default | <name> user <snmp_user_name> sec-model <sec_model> sec-level <sec_level>]`
- `configure snmp notification-log <name> [filter-profile-name [none | <filter_profile_name>] | entry-limit [system-managed | <entry_limit>]]`
- `configure snmp delete notification-log <name>`
- `[enable | disable] snmp notification-log [<name> | all]`
- `clear snmp notification-log [counters | entries] {<name>}`
- `show snmp notification-log`
- `show snmp notification-log <name>`
- `show snmp notification-log <name> entry <entry_index>`



Bidirectional Forwarding Detection (BFD) Up/Down Traps

ExtremeXOS now has read-only support for all BFD MIB tables and global objects. It supports BFD notifications as well.

Set operation is supported only for MIB object 'bfdSessNotificationsEnable' (knob to control up/down traps). The default value for this object is disabled state. No notification is sent in disabled state. Therefore, set operation is also supported for this MIB object to control emission of traps.

Supported Drafts

- *ETF Draft Bidirectional Forwarding Detection*
<http://www.ietf.org/internet-drafts/draft-ietf-bfd-base-09.txt>
- *ETF Draft Generic Application of BFD*
<http://www.ietf.org/internet-drafts/draft-ietf-bfd-generic-05.txt>
- *ETF BFD for Multihop Paths*
<http://www.ietf.org/internet-drafts/draft-ietf-bfd-multihop-07.txt>
- *BFD Management Information Base (draft-ietf-bfd-mib-14)*
<http://tools.ietf.org/html/draft-ietf-bfd-mib-14>
- *BFD textual conventions*
<http://tools.ietf.org/html/draft-ietf-bfd-tc-mib-02>

Supported Platforms

- BlackDiamond X8 and 8800 series switches
- Summit X770, X670, X480, and X460 series switches

New CLI Commands

- `[enable|disable] snmp traps bfd {session-down | session-up}`
- `configure snmp traps batch-delay bfd [none | <delay>]`
- `show snmp traps bfd`



Changed CLI Commands

The output of the `show bfd` command shows BFD up/down trap information (in bold).

```

Number of sessions                : 2
Sessions in Init State            : 0
Sessions in Down State            : 0
Sessions in Admin Down State     : 1
Sessions in Up State              : 1
SNMP Traps for Session Down    : Enabled
SNMP Traps for Session Up     : Enabled
SNMP Traps Batch Delay       : 1000 ms

```

Internet Protocol (IP) v6 Multi-cast Listener Discovery Protocol (MLD) Source-Specific Multi-cast (SSM) Map

This feature enables MLDv1 hosts to participate in Source Specific Multi-cast (SSM). MLD SSM mapping feature is the IPv6 equivalent of the IPv4 feature: IGMP SSM mapping. 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 a DNS name for a group/group range.

Supported Platforms

All platforms, except the Summit X430 series switches.

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

In the following command the keyword `dns group` is optional:

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



Access Control List (ACL) Enhancements

The following ACL enhancement are included in ExtremeXOS 15.5:

- ACL rule deletion improvement—“leave a hole”: Previously, deleting an ACL rule, caused TCAM shifting to occur to fill in all entries from the highest precedence slice to the lowest precedence slice. Now when a rule is deleted no shifting occurs, leaving a “hole” where the entry previously existed.
- ACL rule insertion improvement—attempt shifting up or down within a single slice: Previously, inserting an ACL rule at a given precedence caused all rules to shift down from rules in the lowest precedence slice to the slice mapped to the specified precedence to make an empty position for the new rule. Now shifting will occur either down or up until a “hole” is encountered only within a single slice.
- ACL rule insertion improvement—using virtual slice priorities: Previously, physical slice order dictates the precedence order of slices and rule/slice shifting occurs across the full spectrum of slice precedences as required by each operation. Now a new virtual slice precedence is inserted in between two existing slices to avoid rule shifting spanning multiple slices. Only in the worst case scenario, when all slices have been allocated for at least some rules, will new virtual slices not be available and TCAM rule shifting may have to span multiple slices.
- ACL rule insertion improvement—first attempt to find a “hole” at the same precedence level: Previously, deletions did not form “holes”. However, other ACL improvements listed above will create holes. Now we will attempt to fill an existing hole at the specified precedence level to avoid any shifting associated with the insertion.

Enhance Clear Counters Per Port Command

This feature enhances the `clear counters ports` command by supporting a `<port_list>` argument.

Supported Platforms

All platforms.

Changed CLI Commands

Changes are bolded.

```
clear counters ports {<port_list> | all}
```



Multi-switch Link Aggregation Group (MLAG) Out-of-Band Keep-Alive Protocol

Previously, MLAG peers exchanged keep-alive health check packets between themselves to verify the liveness of peers. When MLAG switches stop receiving health check packets from peers, it could be because of the following reasons:

- Failure of the ISC link when the remote peer is still active.
- Remote peer is down for some reason.

MLAG switches lack the ability to determine the reason they are not receiving the health check keep-alives from the peer switches.

If the ISC link alone goes down when the remote peer is alive, both the MLAG peers forward the south-bound traffic, duplicating traffic. However, this doesn't produce any traffic loops. This is because the remote node load shares to both the MLAG peers and does not forward the traffic received on one of the load-shared member ports to other member ports of the same load-shared group.

To solve the above mentioned issue, this feature directs health check messages to also be exchanged on an alternate path—typically management VLAN. If the peer is alive when the ISC link alone goes down, one of the MLAG peers disables its MLAG ports to prevent duplicate south-bound traffic to the remote node.

Supported Platforms

All platforms.

Limitations

- Checkpointing is not supported on the alternate VLAN used for health checks. For an ISC link failure, if an MLAG port that was active goes down, the corresponding MLAG port on the peer switch that was disabled is not enabled back. In other words, double failures (ISC link and MLAG link going down) are not handled.
- Bidirection Forwarding Detection (BFD) in hardware is not currently supported, and thus cannot be used over the alternate path for faster detection times.

New CLI Commands

```
configure {mLAG peer} <peer_name> alternate ipaddress  
[<ip_address> {vr <vr_name>} | none]
```



Changed CLI Commands

The output of the `show mlag peer` command shows alternate path information.

Multi-switch Link Aggregation Group (MLAG) Support for More than One Peer

ExtremeXOS now supports for MLAG switches to create two MLAG peers.

All of the basic MLAG functionality and the traffic forwarding rules that existed earlier also apply with this feature. An important point to note is that a port is an MLAG port only with respect to a particular MLAG peer switch.

Supported Platforms

All platforms.

Limitations

- To form an MLAG between a pair of switches, a directly connected ISC link is needed between the switches.
- The MLAG peers in a multi-peer setup cannot be looped; however, they can be extended as a linear daisy chain.
- The MLAG feature does not currently detect if the same VLAN is configured on different switches. In other words, MLAG doesn't detect if there are any loops in the network. The network design must ensure that there are no loops in the network.

Multi-switch Link Aggregation Group (MLAG) MD5 Hash for TCP Checkpointing Connecting

The checkpoint messages exchanged between MLAG peers over a TCP connection are sent in plain text and are vulnerable to spoofing. This feature secures the checkpoint connection against spoofing.

An authentication key must be configured on both MLAG peer switches. The checkpoint connection is not established if different keys are configured on the MLAG peer switches.

This key is used in calculating the authentication digest for the TCP messages. TCP_MD5SIG socket option is used for authentication, and so only MD5 authentication is supported. The configured key is used in setting up the TCP_MD5SIG option on the checkpoint socket.

Supported Platforms

All platforms.

Limitations

Only MD5 authentication is supported.

New CLI Commands

```
configure {mlog peer} <peer_name> authentication [md5 key  
{encrypted} <auth_key> | none]
```

Changed CLI Commands

Changes are bolded.

- `create mlog peer <peer_name> {authentication [md5 key {encrypted} <auth_key>]}`
- The output of the `show mlog peer` command displays information about the authentication key and alternative path.



Layer 2 Protocol Tunneling (L2PT) and Filtering

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

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

Supported Platforms

All platforms

Limitations

- L2PT and protocol filtering is implemented in software, thus limiting the number of frames that may be filtered or tunneled.
- Both L2PT and protocol filtering may be configured only through command line interface. Configuration through SNMP/XML is not supported.
- When tunneling protocols which are point-to-point in nature, you must ensure that there are only two tunnel endpoints for the protocol.
- If a protocol that is configured to be tunneled on a service interface cannot be uniquely identified by its DA and EtherType, then all packets with the same DA and EtherType of the protocol being tunneled but are not really PDUs of the protocol are slow path forwarded.
- Tagged protocol PDUs cannot be tunneled over VLANs. Tagged protocol PDUs can be tunneled only over VMANs (the VMAN can be the service VMAN for a VPLS/VPWS service or a standalone VMAN). Untagged protocol PDUs can be tunneled over both VLANs and VMANs (the VLAN/VMAN can be standalone or be the service VMAN for a VPLS/VPWS service).
- Untagged protocol PDUs cannot be bypassed if the ingress port is an untagged VMAN port with a default CVID. Untagged protocol PDUs may be bypassed if the ingress port is an untagged VMAN port without a default CVID.
- L2PT is not supported on VLAN ports that have a port specific tag.
- L2PT for VPLS/VPWS is not supported.



New CLI Commands

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

Changed CLI Commands

Changes are bolded.

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



Y.1731 Compliant Performance Monitoring SNMP MIBs

This feature adds support for Y.1731 performance measurement MIB defined by MEF-36 (http://metroethernetforum.org/PDF_Documents/technical-specifications/MEF_36.pdf).

This implementation the MIB supports GET operations for Frame Loss and Frame Delay.

Supported Platforms

All platforms.

Limitations

- A maximum of 1,800 frames of data are stored for each loss measurement message/delay measurement message (LMM/DMM) session.
- Each frame's data is about 60 bytes for LMM and uses about 44 MB of memory for 288 sessions.
- Each frame's data is about 80 bytes for DMM and uses about 59 MB of memory for 288 sessions.



Dynamic Host Configuration Protocol (DHCPv6) RFC4649 Relay Agent Remote-ID Option

This feature adds a remote-id option to the relay forward IPv6 message, if the specific VLAN has only the link local IPv6 address. The remote-ID is added to all the request packets received from the client on that VLAN, irrespective of whether it is configured on that VLAN. The remote-ID option is in the format defined in RFC 4649, and is added after the relay message header.

If the remote-ID is configured through the command line interface on a specified VLAN, it can have either a user-defined value or the system-name. If you don't configure or mark "none" for the remote-ID, the switch MAC address is used as the remote-ID. The value of the remote-ID present in the packet has the following format:

```
vlan_port_<remoteId>
```

Where <remoteId> is either the configured string, system name, or the switch MAC address. The enterprise number field contains the Extreme Networks enterprise number: 1916.

The remote-ID is added to the packets on a specified VLAN, only if:

- IPv6 bootprelay is enabled on the VLAN.
- The VLAN has only the link local address.

Supported Platforms

All platforms.



Multiprotocol Label Switching (MPLS) Pseudowire—Label-Switched Path (PW-LSP) MIB Counters

By implementing the tables in PW LSP sharing MIB, the SNMP manager can observe the transmit packet counters over each LSP that is configured for use by the PW, and aggregate transmit and receive packet counters over the PW itself.

There are no standard MIB tables and scalars objects that can be used for PW LSP sharing. MIB tables and objects for this feature are proprietary to Extreme Networks.

These three tables are implemented in `extrememplsmib.my`

- `extremePwPerfTable`: Contains the aggregated transmit and receive packet counters for in-service PWs.
- `extremePwLspOutboundMappingTable`: Provides the mapping between PWs and LSPs by providing an LSP index. LSP indexes are assigned uniquely for each PW. Entries in this table indicate that an LSP is being used by an in-service PW. An SNMP notification is sent when an entry is added or deleted.
- `extremePwLspPerfTable`: Contains the transmit packet and byte counters for traffic sent over a specific PW using a specific LSP.

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.



New Hardware Supported in ExtremeXOS 15.5.2

This section lists the new hardware supported in ExtremeXOS 15.5.:2:

- Summit X430-8p and X430-24p

New Hardware Supported in ExtremeXOS 15.5.1

This section lists the new hardware supported in ExtremeXOS 15.5.:1:

- BlackDiamond X8 100G4X blade (4 × 100Gb Ethernet ports)

ExtremeXOS Hardware and Software Compatibility Matrix

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

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

www.extremenetworks.com/documentation

Upgrading to ExtremeXOS

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

- Beginning with ExtremeXOS 12.1, an ExtremeXOS core image (.xos file) must be downloaded and installed on the alternate (non-active) partition. If you try to download to an active partition, the error message "**Error: Image can only be installed to the non-active partition.**" is displayed. An ExtremeXOS modular software package (.xmod file) can still be downloaded and installed on either the active or alternate partition.
- SummitX software is required for E4G cell site routers.
- Beginning with ExtremeXOS 15.4, a limited hitless upgrade procedure is supported on the BlackDiamond X8 and BlackDiamond 8800 series switches

Downloading Supported MIBs

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

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

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:

www.extremenetworks.com/support/service-notification-form



2 Limits

This chapter summarizes the supported limits in ExtremeXOS 15.5.5-Patch1-1.

Supported Limits

Table 3 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 3** 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 3** 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 3: Supported Limits

Metric	Product	Limit
AAA (local) —maximum number of admin and local user accounts.	All platforms	16
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 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	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 3: 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	8,192 ingress, 1,024 egress
	Summit X440, X430 group of 24 ports	1,024 ingress
	Summit X460, E4G-400	4,096 ingress, 512 egress
	Summit X480	(up to) 61,440 ingress, 1,024 egress
	Summit X670 VIM4-40G4x	2,048 ingress 1,024 egress
	Summit X480 VIM3-40G4X	2048 ingress 1024 egress
	Summit X770	4,096 ingress 1,024 egress

Table 3: 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	17 ^b
	8900-40G6X-xm	10 ingress, 4 egress
	BlackDiamond X8 a-series modules	10 ingress, 4 egress
	BlackDiamond X8-100G4X modules	16 ingress, 4 egress
	E4G-200	8 ingress, 4 egress
	Summit X440, X430	4 ingress
	Summit X460, E4G-400	16 ingress, 4 egress
	Summit X480	17 ^b ingress, 4 egress
	Summit X670 VIM4-40G4x	10 ingress, 4 egress
	Summit X480 VIM3-40G4X	10 ingress, 4 egress
Summit X770	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 Summit X670 Summit X430	1,024 4,096 100*
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 with Core license or higher	256

Table 3: Supported Limits (Continued)

Metric	Product	Limit
BGP (networks) —maximum number of BGP networks.	All platforms 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
	Summit X460, X670, X770	128*
	E4G-400, E4G-200	128*
	Summit X480	512
BGP (peer groups) —maximum number of BGP peer groups.	BlackDiamond 8900 series	128
	BlackDiamond X8 series	128
	Summit X480	128
	All platforms (except BlackDiamond X8 series, BlackDiamond 8900 series, and Summit X480) with Core license or higher	64
BGP (policy entries) —maximum number of BGP policy entries per route policy.	All platforms with Core license or higher	256
BGP (policy statements) —maximum number of BGP policy statements per route policy.	All platforms with Core license or higher	1,024
BGP multi-cast address-family routes —maximum number of multi-cast address-family routes.	BlackDiamond 8000 series	25,000
	BlackDiamond 8900 xl-series	524,256 (up to) ^b
	BlackDiamond X8 series	25,000
	Summit X460, X670, X770	25,000
	Summit X480	524,256 (up to) ^b
	E4G-400	25,000
BGP (unicast address-family routes) —maximum number of unicast address-family routes.	BlackDiamond 8000 series	25,000
	BlackDiamond 8900 xl-series	524,256 (up to) ^b
	BlackDiamond X8 series	25,000
	Summit X460, X670, X770	25,000
	Summit X480	25,000
	E4G-400	524,256 (up to) ^b
	25,000	

Table 3: Supported Limits (Continued)

Metric	Product	Limit
BGP (non-unique routes) —maximum number of non-unique BGP routes.	BlackDiamond 8000 series	25,000
	BlackDiamond 8900 xl-series	1,200,000
	BlackDiamond X8 series	25,000
	Summit X460, X670, X770	25,000
	Summit X480	1,000,000
	E4G-400, E4G-200	25,000
BGP ECMP —maximum number of equalcost multipath for BGP and BGPv6.	All platforms, except Summit X440	2, 4, or 8
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
	Summit X460	6,000
	Summit X480	20,000
	Summit X670, X770	8,000
E4G-400	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
	Summit X460	18,000
	Summit X480, X670, X770	24,000
E4G-400	18,000	
BOOTP/DHCP relay —maximum number of BOOTP or DHCP servers per virtual router.	All platforms	4
BOOTP/DHCP relay —maximum number of BOOTP or DHCP servers per VLAN.	All platforms	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.	All platforms	8
CFM —maximum number of CFM associations.	All platforms	256
CFM —maximum number of CFM up end points.	BlackDiamond 8000 series	32
	BlackDiamond X8 series	32
	Summit series	32

Table 3: Supported Limits (Continued)

Metric	Product	Limit
CFM —maximum number of CFM down end points.	BlackDiamond 8000 series	32
	BlackDiamond X8 series	32
	Summit series X460, E4G-200, E4G-400 (non-load shared ports)	256
	Summit series X460, E4G-200, E4G-400 (load shared ports)	32
	Summit series All other platforms	32 32
CFM —maximum number of CFM remote end points per up/down end point.	All platforms	2,000
CFM —maximum number of dot1ag ports.	All platforms	128
CFM —maximum number of CFM segments.	All platforms	1,000
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 8800 c-series	4,096
	BlackDiamond 8900 series	4,096
	BlackDiamond X8 series	4,096
	Summit X440	1,024
	Summit X670	2,048
Summit X480, Summit X770	4,096	
Data Center Bridging eXchange (DCBX) protocol Type Length Value (TLVs) —maximum number of DCBX application TLVs.	All platforms	8
Dynamic ACLs —maximum number of ACLs processed per second. NOTE: Limits are load dependent.	BlackDiamond 8800 with c-series MSM and I/O modules	8
	BlackDiamond 8900 series	8
	BlackDiamond X8 series	8
	Summit X480, X670	10
	with 50 DAACLs with 500 DAACLs	5
EAPS domains —maximum number of EAPS domains. NOTE: An EAPS ring that is being spatially reused cannot have more than four configured EAPS domains.	BlackDiamond 8000 series	64
	BlackDiamond X8 series	64
	Summit series (except X430), E4G-200, E4G-400	32
	Summit X430	4
EAPSV1 protected VLANs —maximum number of protected VLANs.	BlackDiamond 8000 series	2,000
	BlackDiamond X8 series	4,000
	Summit series, E4G-200, E4G-400	1,000

Table 3: Supported Limits (Continued)

Metric	Product	Limit
EAPsv2 protected VLANs —maximum number of protected VLANs.	BlackDiamond 8000 series	2,000
	BlackDiamond X8 series	4,000
	Summit series, E4G-200, E4G-400	500
ELSM (vlan-ports) —maximum number of VLAN ports.	BlackDiamond 8000 series	5,000
	BlackDiamond X8 series	5,000
	Summit series, E4G-200, E4G-400	5,000
ERPS domains —maximum number of ERPS domains without CFM configured	BlackDiamond 8806 series	32
	BlackDiamond X8 series	32
	Summit series (except X430), E4G-200, E4G-400	32
	Summit X430	4
ERPS domains —maximum number of ERPS domains with CFM configured.	BlackDiamond 8806 series	16
	BlackDiamond X8 series	16
	Summit series non-CSR platforms	16
	Summit X460	32
	E4G-200, E4G-400	32
ERPSv1 protected VLANs —maximum number of protected VLANs.	BlackDiamond 8806 series	2,000
	BlackDiamond X8 series	2,000
	Summit series, E4G-200, E4G-400	1,000
ERPSv2 protected VLANs —maximum number of protected VLANs	BlackDiamond 8806 series	2,000
	BlackDiamond X8 series	2,000
	Summit series, E4G-200, E4G-400	500
ESRP groups —maximum number of ESRP groups.	All platforms	7
ESRP domains —maximum number of ESRP domains.	BlackDiamond 8000 series	64
	BlackDiamond X8 series	64
	BlackDiamond 8900 series	128
	Summit series	64
ESRP VLANs —maximum number of ESRP VLANs.	BlackDiamond 8000 series	1,000
	BlackDiamond X8 and 8900 series	2,048
	Summit series	1,000
ESRP (maximum ping tracks) —maximum number of ping tracks per VLAN.	All platforms	8
ESRP (IP route tracks) —maximum IP route tracks per VLAN.	All platforms	8
ESRP (VLAN tracks) —maximum number of VLAN tracks per VLAN.	All platforms	1

Table 3: Supported Limits (Continued)

Metric	Product	Limit
Forwarding rate —maximum L2/L3 software forwarding rate.	BlackDiamond 8000 series	10,000 pps
	BlackDiamond X8 series	20,000 pps
	Summit series, except Summit X770	10,000 pps
	Summit X770	16,000 pps
FDB (blackhole entries) —maximum number of unicast blackhole FDB entries.	BlackDiamond 8800 c-series	32,000
	BlackDiamond 8000 e-series	8,000
	BlackDiamond 8900 series	
	8900 c-series	32,000
	8900 xl-series	524,288 (up to) ^b
	8900-40G6X-xm	128,000
	BlackDiamond X8 a-series modules	128,000
	BlackDiamond X8-100G4X modules	384,000
	E4G-200, E4G-400	32,000
	Summit X440, X430	16,000
	Summit X480	524,288 (up to) ^b
	Summit X460	32,000
Summit X670		
VIM4-40G4x	128,000	
Summit X770	288,000 ^d	
FDB (blackhole entries) —maximum number of multi-cast blackhole FDB entries.	BlackDiamond 8000 series	1,024
	BlackDiamond X8 series	1,024
	All Summit series switches, except X770	1,024
	Summit X770	4,000

Table 3: Supported Limits (Continued)

Metric	Product	Limit
FDB (maximum L2 entries) — maximum number of MAC addresses.	BlackDiamond 8000 c-series	32,768 ^e
	BlackDiamond 8000 e-series	8,192 ^e
	BlackDiamond 8000 (system), except 8900 xl-series	128,000 ^e
	BlackDiamond 8900 xl-series	524,488 (up to) ^b
	BlackDiamond X8 a-series modules	128,000 ^e
	BlackDiamond X8-100G4X modules E4G-200, E4G-400	384,000 ^e
	Summit X440, X430	32,000 ^e
	Summit X480	16,000 ^e
	VIM3-40G4X module	524,488 (up to) ^b
	Summit X460	128,000 ^e
Summit X670	32,768 ^e	
Summit X770	128,000 ^e	
Summit X770	288,000 ^d	
FDB (Maximum L2 entries) — maximum number of multi- cast FDB entries.	BlackDiamond X8	1,024
	BlackDiamond 8800	
	All Summit series switches, except X770	
	Summit X770	4,000
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	
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	512
Identity management — maximum number of Whitelist entries.	All platforms	512

Table 3: Supported Limits (Continued)

Metric	Product	Limit
Identity management — maximum number of roles that can be created.	All platforms	64
Identity management — maximum role hierarchy depth allowed.	All platforms	5
Identity management — maximum number of attribute value pairs in a role match criteria.	All platforms	16
Identity management — maximum of child roles for a role.	All platforms	8
Identity management — maximum number of policies/ dynamic ACLs that can be configured per role.	All platforms	8
Identity management — maximum number of LDAP servers that can be configured.	All platforms	8
Identity management — maximum number of Kerberos servers that can be configured.	All platforms	20
Identity management — maximum database memory- size.	All platforms	64-49, 152
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	100
Identity management — recommended number of ACL entries per identity. NOTE: Number of ACLs per identity based on system ACL limitation.	All platforms	20
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

Table 3: Supported Limits (Continued)

Metric	Product	Limit
IGMP sender —maximum number of IGMP senders per switch (IP multi-cast compression disabled). ^k NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 ^f
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900-10G24X-c modules	2,048 ^f
	BlackDiamond 8900-G96T-c modules	4,096 ^f
	BlackDiamond 8900-40G6X-xm	3,000 ^g
	BlackDiamond 8900 xl-series	4,096 ^f
	BlackDiamond X8 a-series modules	4,096 ^h
	BlackDiamond X8-100G4X	16,384 ^d
	E4G-200, E4G-400	2,048
	Summit X440	64
	Summit X480	4,096
	Summit X460	2,048
	Summit X670	
VIM4-40G4x	3,000 ^g	
Summit X770	4,000	
IGMP sender —maximum number of IGMP senders per switch (IP multi-cast compression enabled). ^k 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 58 • Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mixed-mode. on page 58 	BlackDiamond 8800 c-series	6,000 ^g
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 c-series	6,000 ^g
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond 8900-40G6X-xm	3,000 ^g
	BlackDiamond X8 a-series modules	6,000 ^{g h}
	BlackDiamond X8-100G4X modules	64,000 ^{g h}
	E4G-200	3,000 ^g
	E4G-400	6,000 ^g
	Summit X440	192 ^g
	Summit X460	6,000 ^g
	Summit X480	12,000 ^b
	Summit X670	
VIM4-40G4x	3,000 ^g	
Summit X770	16,000	

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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 modules	4,000
	E4G-200, E4G-400	1,000
	Summit X440	448
	Summit X460, X670	1,000
Summit X480	4,000	
Summit X770	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
IGMPv2 subscriber —maximum number of IGMPv2 subscribers per port. ¹	BlackDiamond 8800 c-series	2,000
	BlackDiamond 8900 c-series	2,000
	BlackDiamond X8 series	2,000
	Summit series (except Summit X460, X480, X770, and X670)	1,000
	Summit X460, X480, X670, E4G-400	2,000
Summit X770	3,000	
IGMPv2 subscriber —maximum number of IGMPv2 subscribers per switch. ¹	BlackDiamond 8800 c-series	20,000
	BlackDiamond 8900 c-series	20,000
	BlackDiamond X8 series	20,000
	Summit series (except Summit X480, X770, and X670)	10,000
	Summit X460, X480, X670, E4G-400	20,000
Summit X770	25,000	
IGMPv3 maximum source per group —maximum number of source addresses per group.	All platforms	250

Table 3: Supported Limits (Continued)

Metric	Product	Limit
IGMPv3 subscriber —maximum number of IGMPv3 subscribers per port. ¹	BlackDiamond 8800 e-series	1,000
	BlackDiamond 8800 c-series	2,000
	BlackDiamond 8900 series	5,000
	BlackDiamond X8 series	3,000
	Summit series (except Summit X460, X770)	1,000
	Summit X770	2,500
IGMPv3 subscriber —maximum number of IGMPv3 subscribers per switch. ¹	BlackDiamond 8800 e-series	10,000
	BlackDiamond 8800 c-series	20,000
	BlackDiamond 8900 series	30,000
	BlackDiamond X8 series	30,000
	Summit series (except Summit X460 and X770)	10,000
	Summit X460, E4G-400	20,000
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).	Summit X770	131,072 (up to) ¹
	BlackDiamond X8-100G4X modules	229,374 (up to) ¹
	All other platforms	20,480
IP ARP entries in software with distributed mode on —maximum number of IP ARP entries in software with distributed mode on.	BlackDiamond 8000 series with 8900-MSM128 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

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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 ^g
	BlackDiamond 8900 xl-series	16,000
	BlackDiamond X8 a-series	16,000
	BlackDiamond X8-100G4X modules	182,000(up to) ⁱ
	E4G-200	8,000
	E4G-400	16,000
	Summit X440	412
	Summit X670	8,000
	Summit X460, X480	16,000
Summit X770	108,000(up to) ⁱ	
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 ^g
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 xl-series	12,000 ^g
	BlackDiamond X8 a-series	12,000 ^g
	BlackDiamond X8-100G4X modules	172,000 (up to) ⁱ
	E4G-200	6,000 ^g
	E4G-400	12,000 ^g
	Summit X440	380
	Summit X460, X480	12,000 ^g
	Summit X670	6,000 ^g
Summit X770	98,000 (up to) ⁱ	
IP flow information export (IPFIX) —number of simultaneous flows.	BlackDiamond 8900 xl-series modules	4,096 ingress, 4,096 egress
	BlackDiamond 8900 c-series modules	4,096 ingress, 4,096 egress
	BlackDiamond X8-100G4X modules	2,048 ingress, 2,048 egress
	Summit X460, X460-G2	2,048 ingress, 2,048 egress
	Summit X480	4,096 ingress, 4,096 egress
	E4G-400	2,048 ingress, 2,048 egress

Table 3: 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 ^g
	BlackDiamond 8000 e-series	1,000 ^g
	BlackDiamond 8900 xl-series	40,000 ^b
	BlackDiamond 8900-40G6X-xm	22,000 ^g
	BlackDiamond X8 a-series	28,000 ^g
	BlackDiamond X8-100G4X modules	311,000 (up to) ⁱ
	E4G-200	18,000 ^g
	E4G-400	20,000 ^g
	Summit X440	448
	Summit X460	20,000 ^g
Summit X480	40,000 ^b	
Summit X670	22,000 ^g	
Summit X770	176,000 (up to) ⁱ	
IPv4 routes —maximum number of IPv4 routes in software (combination of unicast and multi-cast routes).	BlackDiamond 8900 xl-series with 8900-MSM128 or MSM-48c	524,256 (up to) ^b
	All other BlackDiamond 8000 series hardware	25,000
	BlackDiamond X8 series	25,000
	Summit X440	256
	Summit X460, X670, E4G-400, E4G-200	
	SummitStack or standalone	25,000
	Summit X480	
SummitStack or standalone	524,256 (up to) ^b	
Summit X770	25,000	
IPv4 routes (LPM entries in hardware) — number of IPv4 routes in hardware.	BlackDiamond 8800 c-series	12,000
	BlackDiamond 8000 e-series	480
	BlackDiamond 8900 xl-series	524,256 (up to) ^{b j}
	BlackDiamond 8900-40G6X-xm	16,000 ^e
	BlackDiamond X8 series	16,000 ^e
	E4G-200, E4G-400	12,000
	Summit X440	32
	Summit X460	12,000
	Summit X480	524,256 (up to) ^{b j}
	Summit X670	16,000 ⁱ
Summit X770	16,000	

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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 BlackDiamond X8 series E4G-200, E4G-400 Summit X440 Summit X460, X480 Summit X770, X670	512 2,048 512 254 512 2,048
IPv6 host entries in hardware — maximum number of IPv6 neighbor entries in hardware.	BlackDiamond 8800 c-, xm-series BlackDiamond 8000 e-series BlackDiamond 8900-10G24X-c modules BlackDiamond 8900-G96T-c modules BlackDiamond 8900 xl-series BlackDiamond X8 a-series BlackDiamond X8-100G4X modules E4G-200 E4G-400 Summit X440 Summit X460, X670 Summit X480 Summit X770	3,000 ^g 250 ^g 2,000 ^g 4,000 ^g 8,192 (up to) ^b 3,000 ^g 49,000 2,000 ^g 3,000 ^g 192 3,000 ^g 8,192 (up to) ^b 36,000
IPv6 route sharing in hardware — route mask lengths for which ECMP is supported in hardware.	Summit X460, X480, X670, X670V- 48t E4G-200, E4G-400 BlackDiamond 8800 (all I/O modules, except G48Te2), BlackDiamond X8 10G and 40G Summit X770 BlackDiamond X8 100G Summit X440, X430 BlackDiamond 8800 G48Te2	0-128 0-128 0-128 0-64 (> 64 single path only) 0-128 (> 64 single path only) N/A N/A

Table 3: 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), BlackDiamond X8 10G and 40G	0-128
	Summit X770	0-64 (> 64 single path only)
	BlackDiamond X8 100G	0-128 (> 64 single path only)
	Summit X440, X430 BlackDiamond 8800 G48Te2	N/A N/A
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
	E4G-200, E4G-400	6,000
	Summit X440	16
	Summit X460	6,000
	Summit X670	8,000
	Summit X480	245,760 (up to) ^b
Summit X770	8,000	
IPv6 routes with a mask greater than 64 bits in hardware —maximum number of such IPv6 LPM routes in hardware.	BlackDiamond 8000 c-, e-, xm-series	256
	BlackDiamond 8000 xl-series	245,760 (up to) ^b
	BlackDiamond X8 series	256
	E4G-200, E4G-400	256
	Summit X440, X460, X670, X770 Summit X480	256 245,760 (up to) ^b

Table 3: Supported Limits (Continued)

Metric	Product	Limit
IPv6 routes in software —maximum number of IPv6 routes in software.	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
	Summit X460, X670, X770, E4G-200, E4G-400, SummitStack, or standalone	25,000
	Summit X440	256
	Summit X480, SummitStack, or standalone	245,760 (up to) ^b
IP router interfaces —maximum number of VLANs performing IP routing-excludes sub VLANs (IPv4 and IPv6 interfaces).	Summit X670, X770, and BlackDiamond X8	2,048
	Summit X440	254
	All other platforms	512
IP multi-cast static routes —maximum number of permanent multi-cast IP routes.	All platforms	1,024
IP unicast static routes —maximum number of permanent IP unicast routes.	All platforms	1,024
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 X670 and BlackDiamond X8	2, 4, 8, 16, or 32
	Summit X670, BlackDiamond X8	2, 4, 6, 8, 16, 32, or 64

Table 3: Supported Limits (Continued)

Metric	Product	Limit
<p>IP route sharing (total destinations)—maximum number of unique destinations used by multipath OSPF, OSPFv3, BGP, IS-IS, or static routes.</p> <p>NOTE: For platforms with limit of 524,256, the total number of "destination+gateway" pairs is limited to 1,048,512. For example, if the number of unique destinations is 524,256, only 2 gateways per destination is supported.</p> <p>For other platforms, each limit is based on up to 8 gateways per destination for BGP and IS-IS routing protocols, up to 16 gateways per destination for OSPF, or up to 32 gateways per destination for static routes.</p>	BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 series Summit X460, E4G-200, E4G-400 Summit X480 Summit X670, X770 E4G-200, E4G-400	12,256 480 524,256 (up to) ^b 16,352 16,000 12,256 524,256 (up to) ^b 16,352 12,256
	<p>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.</p>	BlackDiamond 8800 c-, xl-, and xm-series default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 BlackDiamond 8000 e-series default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 BlackDiamond X8 series, Summit X670 default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 if maximum gateways is 64 Summit X460, X480, X770, E4G-200, E4G-400 default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32
<p>IP multinetting (secondary IP addresses)—maximum number of secondary IP addresses per VLAN.</p>	All platforms	64

Table 3: Supported Limits (Continued)

Metric	Product	Limit
IS-IS adjacencies —maximum number of supported IS-IS adjacencies.	BlackDiamond 8000 series	128
	BlackDiamond X8 series	128
	BlackDiamond 8900 xl-series	255
	Summit X460, X480, X670, X770, E4G-400	128
	E4G-200	256
IS-IS ECMP —maximum number of equal cost multipath for IS-IS.	All platforms, except Summit X440	2, 4, or 8
IS-IS interfaces —maximum number of interfaces that can support IS-IS.	All platforms	255
IS-IS routers in an area —recommended maximum number of IS-IS routers in an area.	Summit X480	128
	All other platforms	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 8900 xl-series	30,000
	Summit X480	30,000
	Summit X460, X670, X770, E4G-400	20,000
E4G-200	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 8900 xl-series	120,000
	Summit X480	50,000
	Summit X460, X670, X770, E4G-400	25,000
IS-IS IPv4 L2 routes —recommended maximum number of IS-IS Level 2 routes.	BlackDiamond 8000 series	25,000
	BlackDiamond X8 series	25,000
	BlackDiamond 8900 xl-series	120,000
	Summit X480	50,000
	Summit X460, X670, X770, E4G-400	25,000
IS-IS IPv4 L1 routes in an L1/L2 router —recommended maximum number of IS-IS Level 1 routes in an L1/L2 IS-IS router.	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond 8900 xl-series	20,000
	Summit X460, X480, X670, X770, E4G-400	20,000
	E4G-400	20,000

Table 3: 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 8900 xl-series	40,000
	Summit X480	25,000
	Summit X460, X670, 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 8900 xl-series	40,000
	Summit X480	25,000
	Summit X460, X670, X770, E4G-400	10,000
IS-IS IPv6 L1 routes in an L1/L2 router —recommended maximum number of IS-IS Level 1 routes in a L1/L2 router.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond 8900 xl-series	15,000
	Summit X480	15,000
	Summit X460, X670, X770, E4G-400	10,000
IS-IS IPv4/IPv6 L1 routes in an L1 router —recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond 8900 xl-series	60,000
	Summit X480	40,000
	Summit X460, X670, X770, E4G-400	20,000
IS-IS IPv4/IPv6 L2 routes in an L2 router —recommended maximum number of IS-IS Level 2 routes in a Level 2 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond 8900 xl-series	60,000
	Summit X480	40,000
	Summit X460, X670, X770, 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, X480, X670, X770, 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	16

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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
	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
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	1,000
	Summit 460	1,000
	Summit X480, X670, Summit X670V-48t, Summit X770	1,023
	Summit X480-40G VIM	1,023
L2 VPN: VPLS peers —maximum number of VPLS peers per VPLS instance.	Summit X480, Summit X770	64
	Summit X460, Summit X670	32
	BlackDiamond 8900 xl-series	64
	BlackDiamond 8900-40G6x-xm	64
	Summit X670V-48t, Summit X480-40G VIM	64
	BlackDiamond X8 series	64
	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 X460	1,000
	Summit X480	7,000
	Summit X670	3,000
	Summit X670V-48t	7,000
	Summit X480-40G VIM	3,000
Summit X770	7,800	

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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, Summit X480, Summit X670V-48t	7,116
	Summit X770	15,308
	Summit X480-40G, Summit X670	3,020
	E4G-200 E4G-400	2,764 6,860
L2 VPN: Virtual Private Wire Service (VPWS) VPNs — maximum number of virtual private networks per switch.	Summit X460	1,000
	Summit X480, Summit 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 E4G-200, E4G-400	4,000 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 and E4G-400 cell site routers, Summit X460, X480	8,000
	Summit X670	15,000
	Summit X440	4,000
	Summit X770	77,500 ⁱ
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	15,000
	E4G-200 and E4G-400 cell site routers, Summit X460, X480	8,000
	Summit X440	4,000
	Summit X770	77,500 ⁱ

Table 3: Supported Limits (Continued)

Metric	Product	Limit
Layer-3 IPMC forwarding caches —(PIM, MVR, PVLAN) in mixed-mode. ⁹ 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 modules	64,000
	E4G-200 cell site routers, Summit X670	3,000
	E4G-400 cell site routers, Summit X460, X480	6,000
	Summit X440	192
	Summit X770	77,500 ⁱ
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 3: 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 X670 (standalone)	32 * 16 **
	Summit X670 (stacked)	64 * 16 **
	Summit X770 (standalone)	32
	Summit X770 (stacked)	64
	All other Summit series, SummitStacks, 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
Mirroring (filters) —maximum number of mirroring filters. NOTE: This is the number of filters across all the active mirroring instances.	BlackDiamond 8000 series	128
	BlackDiamond X8 series	128
	Summit series	128
Mirroring (monitor port) —maximum number of monitor ports.	All platforms	1
Mirroring, one-to-many (filters) —maximum number of one-to-many mirroring filters. NOTE: This is the no. of filters across all the active mirroring instances	BlackDiamond 8000 series	128
	BlackDiamond X8 series	128
	Summit series	128
Mirroring, one-to-many (monitor port) —maximum number of one-to-many monitor ports.	All platforms	16

Table 3: Supported Limits (Continued)

Metric	Product	Limit
Maximum mirroring instances NOTE: Only two or four mirroring instance will be active at a time depending on the mirroring filter added to it. There are four hardware resource slots. Each single instance uses one such slot, while each ingress plus egress instance uses two slots. So this allows the you to use a total of four slots, while there are no more then two egress instances. The maximum possible combination for mirroring instances: 1 4 ingress 2 3 ingress + 1 egress 3 2 ingress + 2 egress 4 2 (ingress + egress) 5 1 (ingress + egress) + 2 ingress 6 1 (ingress + egress) + 1 egress + 1 ingress	All platforms	16 (including default mirroring instance)
MLAG ports —maximum number of MLAG ports allowed.	BlackDiamond 8800 series BlackDiamond X8 series Summit series	768 768 768
MLAG peers —maximum number of MLAG peers allowed.	All platforms	2
MPLS RSVP-TE interfaces —maximum number of interfaces.	All platforms	32
MPLS RSVP-TE ingress LSPs —maximum number of ingress LSPs.	All platforms	2,000
MPLS RSVP-TE egress LSPs —maximum number of egress LSPs.	All platforms	2,000
MPLS RSVP-TE transit LSPs —maximum number of transit LSPs.	All platforms	2,000
MPLS RSVP-TE paths —maximum number of paths.	All platforms	1,000
MPLS RSVP-TE profiles —maximum number of profiles.	All platforms	1,000
MPLS RSVP-TE EROs —maximum number of EROs per path.	All platforms	64
MPLS RSVP-TE fast reroute —MPLS RSVP-TE fast reroute (FRR) switching time.	E4G-200	50 msec

Table 3: Supported Limits (Continued)

Metric	Product	Limit
MPLS LDP peers —maximum number of MPLS LDP peers per switch.	Summit X460, Summit X670	32
	Summit X480, Summit X480-40G VIM, Summit X670V-48t, Summit X770	64
	BlackDiamond 8900 xl-series	64
	BlackDiamond 8900-40G6x-xm	64
	BlackDiamond X8 series	64
	E4G-400, E4G-200	32
	MPLS LDP adjacencies —maximum number of MPLS LDP adjacencies per switch.	BlackDiamond 8900 xl-series
BlackDiamond 8900-40G6x-xm		64
BlackDiamond X8 series		50
E4G-200, E4G-400		50
Summit X460, X480, X670		50
Summit X670V-48t, Summit X480-40G VIM, Summit X770		64
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, Summit X670V-48t, Summit X480-40G VIM, Summit X770	2,048
MPLS LDP-enabled interfaces —maximum number of MPLS LDP configured interfaces per switch.	Summit X460, X670	32
	Summit X480, Summit X480-40G VIM, Summit X670V-48t, Summit X770	64
	BlackDiamond 8900 xl-series	64
	BlackDiamond 8900-40G6x-xm	64
	BlackDiamond X8 series	64
	E4G-200, E4G-200	32
	MPLS LDP Sessions —maximum number of MPLS LDP sessions.	E4G-200
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, Summit X480, Summit X770, Summit X670V-48t	4,000
	Summit X670, Summit X480-40G	3,000

Table 3: 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, Summit X670V-48t	7,000
	Summit X670, Summit x480-40G VIM Summit X770	3,000 8,000
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, Summit X480, Summit X670V-48t	7,116
	Summit X480-40G, Summit X670	3,020
	Summit X770	8,000
	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, Summit X480	4,000
	Summit x480-40G, Summit X670, Summit x670V-48t, Summit X770	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, Summit X480, X670V-48t, Summit X770	4,000
	Summit X480-40G, Summit X670	3,000
	E4G-200	2,700
	E4G-400	4,000
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	16
	Summit X770	64

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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 X770	16,000
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	4
	Summit X770	16
Multi-cast listener discovery (MLD) IPv6 multi-cast data sender —maximum number of IPv6 multi-cast streams supported on a switch ^{k 9} 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 58 Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mixed-mode. on page 58 	BlackDiamond 8800 c-series	1,000
	BlackDiamond 8800 e-series	250
	BlackDiamond 8900 c-series	1,000
	BlackDiamond 8900-40G6X-xm	1,000
	BlackDiamond 8900 xl-series	3,000
	BlackDiamond X8 series	3,000
	E4G-200	1,500
	E4G-400	3,000
	Summit X440	90
	Summit X460	3,000
	Summit X480	3,000
	Summit X670	1,500
	Summit X770	4,096
Multi-cast 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 series	500
	E4G-400, Summit X460	1,000
	Summit X480	2,000
	Summit X440	250
	Summit X670, E4G-200	500
	Summit X770	1,200

Table 3: Supported Limits (Continued)

Metric	Product	Limit
Multi-cast listener discovery (MLD)v1 subscribers —maximum number of MLDv1 subscribers per port ¹	BlackDiamond 8800 c-series	500
	BlackDiamond xl-series	1,500
	BlackDiamond X8 Series	1,500
	Summit X440, SummitStack	750
	Summit X460, X480, X670, E4G-400	1,500
	Summit X770	4,000
Multi-cast listener discovery (MLD)v1 subscribers —maximum number of MLDv1 subscribers per switch ¹	BlackDiamond 8800 series	10,000
	BlackDiamond X8 series	10,000
	Summit X440, SummitStack	5,000
	Summit X460, X480, X670, E4G-400	10,000
	Summit X770	30,000
Multi-cast listener discovery (MLD)v2 subscribers —maximum number of MLDv2 subscribers per port ¹	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	2,000
	Summit X770	4,000
Multi-cast listener discovery (MLD)v2 subscribers —maximum number of MLDv2 subscribers per switch ¹	BlackDiamond 8800 series	10,000
	BlackDiamond xl series	10,000
	Summit X440, SummitStack	5,000
	Summit X460, X480, X670, E4G-400	10,000
	Summit X770	30,000
Multi-cast listener discovery (MLD)v2 maximum source per group —maximum number of source addresses per group	All platforms	200

Table 3: Supported Limits (Continued)

Metric	Product	Limit
Multi-cast VLAN registration (MVR) —maximum number of MVR senders per switch (IP multi-cast compression disabled). NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 ^f
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^f
	8900-G96T-c modules	4,096 ^f
	8900 xl-series	4,096 ^f
	8900-40G6X-xm	3,000 ^g
	BlackDiamond X8 a-series modules	4,096
	BlackDiamond X8-100G4X modules	16,000
	E4G-200	2,048
	E4G-400	500 ^g
	Summit X440	1,024
	Summit X480	2,048
Summit X460	2,048	
Summit X670		
VIM4-40G4x	3,000 ^g	
Summit X770	4,096	
Multi-cast VLAN registration (MVR) —maximum number of MVR senders per switch (IP multi-cast compression enabled). NOTE: Assumes source-group-vlan mode. For additional limits, see: Layer-3 IPMC forwarding caches —(PIM, MVR, PVLAN) in mixed-mode. ^g on page 59	BlackDiamond 8800 c-series	6,000 ^g
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 c-series	6,000 ^g
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 a-series modules	6,000 ^g
	BlackDiamond X8-100G4X modules	64,000
	8900-40G6X-xm module	3,000 ^g
	Summit X440	192 ^g
	Summit X460, E4G-400	6,000 ^g
	Summit X480	12,000 ^b
	Summit X670	
	VIM4-40G4x	3,000 ^g
Summit X770	6,300	
Network login —maximum number of clients being authenticated on MAC-based VLAN enabled ports.	BlackDiamond 8000 series (clients per module/per system)	1,024
	BlackDiamond X8 series	1,024
	Summit series	1,024
Network login —maximum number of dynamic VLANs.	All platforms	1,024
Network login VLAN VSAs —maximum number of VLANs a client can be authenticated on at any given time.	All platforms	10

Table 3: Supported Limits (Continued)

Metric	Product	Limit
OSPFv2/v3 ECMP —maximum number of equal cost multipath OSPFv2 and OSPFv3.	All platforms, except Summit X440	2, 4, 8, or 16
OSPFv2 areas —as an ABR, how many OSPF areas are supported within the same switch.	All platforms	8
OSPFv2 external routes —recommended maximum number of external routes contained in an OSPF LSDB.	BlackDiamond 8000 series BlackDiamond 8900 xl-series BlackDiamond X8 series Summit X460, X670, X770 Summit X480 E4G-400 E4G-200	20,000 130,000 20,000 5,000 130,000 5,000 5,000
OSPFv2 inter- or intra-area routes —recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB with one ABR in OSPF domain.	BlackDiamond 8000 series BlackDiamond 8900 xl-series BlackDiamond X8 series Summit X460, X670 E4G-400 Summit X480, X770	7,000 7,000 7,000 2,000 2,000 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	4
	All platforms with Core license or higher (active interfaces only)	400
OSPFv2 links —maximum number of links in the router LSA.	All platforms, except Summit X770 Summit X770	400 419
OSPFv2 neighbors —maximum number of supported OSPF adjacencies.	BlackDiamond 8000 series BlackDiamond 8900 xl-series BlackDiamond X8 Series Summit X460, X670, X770 Summit X440 Summit X480 E4G-400, E4G-200	128 255 255 128 128 255 128
OSPFv2 routers in a single area —recommended maximum number of routers in a single OSPF area.	BlackDiamond 8000 series BlackDiamond 8900 xl-series BlackDiamond X8 series Summit X460, X670, X770 Summit X480 E4G-400	100 200 100 50 200 50

Table 3: Supported Limits (Continued)

Metric	Product	Limit
OSPFv2 virtual links —maximum number of supported OSPF virtual links.	All platforms with Core license or higher	32
OSPFv3 areas —as an ABR, the maximum number of supported OSPFv3 areas.	All platforms with Core license or higher	16
OSPFv3 external routes —recommended maximum number of external routes.	BlackDiamond 8000 series BlackDiamond X8 series BlackDiamond 8900 xl-series Summit X460, X670, X770 Summit X480 E4G-400	10,000 10,000 60,000 10,000 60,000 10,000
OSPFv3 inter- or intra-area routes —recommended maximum number of inter- or intra-area routes.	BlackDiamond 8000 series BlackDiamond X8 series BlackDiamond 8900 xl-series Summit X460, X670, X770 Summit X480 E4G-400	6,000 6,000 6,000 3,000 6,000 3,000
OSPFv3 interfaces —maximum number of OSPFv3 interfaces.	NOTE: Active interfaces only, with Advanced Edge license. (See below for Core license limits.) All platforms	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 E4G-400	256 256 384 128 384 128
OSPFv3 neighbors —maximum number of OSPFv3 neighbors.	BlackDiamond 8000 series BlackDiamond X8 series BlackDiamond 8900 xl-series Summit X460, X670, X770 Summit X480 E4G-400	64 64 128 64 128 64
OSPFv3 virtual links —maximum number of OSPFv3 virtual links supported.	All platforms with Core license or higher	16

Table 3: Supported Limits (Continued)

Metric	Product	Limit
<p>PIM IPv4 snooping—maximum number of (S,G) entries programmed in the hardware (IP multi-cast compression disabled).</p> <p>NOTE: Assumes source-group-vlan mode.</p>	BlackDiamond 8800 c-series	2,048 ^f
	BlackDiamond 8000 e-series	500 ^f
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^f
	8900-G96T-c modules	4,096 ^f
	8900 xl-series	4,096 ^f
	8900-40G6X-xm	3,000 ^g
	BlackDiamond X8 a-series modules	4,096
	BlackDiamond X8-100G4X modules	16,384
	E4G-200	2,048
	E4G-400	2,048
	Summit X440	64
	Summit X460	2,048
	Summit X480	4,096
Summit X670		
VIM4-40G4x	3,000 ^g	
Summit X770	4,096	
<p>PIM IPv4 snooping—maximum number of (S,G) entries programmed in the hardware (IP multi-cast 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 58 • Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode. on page 58 	BlackDiamond 8800 c-series	6,000 ^g
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 c-series	6,000 ^g
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 a-series modules	6,000 ^g
	BlackDiamond X8-100G4X modules	64,000 ^g
	E4G-200	3,000 ^g
	E4G-400	6,000 ^g
	8900-40G6X-xm	3,000 ^g
	Summit X440	192 ^g
	Summit X480	12,000 ^b
	Summit X460	6,000 ^g
	Summit X670	
	VIM4-40G4x	3,000 ^g
Summit X770	66,500	

Table 3: Supported Limits (Continued)

Metric	Product	Limit
<p>PIM IPv4—maximum routes— maximum number of (S,G) entries installed in the hardware (IP multi-cast compression disabled).</p> <p>NOTE: Assumes source-group-vlan mode.</p>	BlackDiamond 8800 c-series	2,048 ^f
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^f
	8900-G96T-c modules	4,096 ^f
	8900 xl-series	4,096 ^f
	8900-40G6X-xm	3,000 ^g
	BlackDiamond X8 a-series modules	4,094
	BlackDiamond X8-100G4X modules	16,384
	E4G-200	2,048
	E4G-400	2,048
	Summit X440	64 ^g
	Summit X480	4,096
	Summit X460	2,048
Summit X670		
VIM4-40G4x	3,000 ^g	
Summit X770	4,096	
<p>PIM IPv4—maximum routes— maximum number of (S,G) entries installed in the hardware (IP multi-cast compression enabled).</p> <p>NOTE: Assumes source-group-vlan mode.</p> <p>For additional limits, see: Layer-3 IPMC forwarding caches—(PIM, MVR, PVLAN) in mixed-mode.^g on page 59</p>	BlackDiamond 8800 c-series	6,000 ^g
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 c-series	6,000 ^g
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 a-series modules	6,000 ^f
	BlackDiamond X8-100G4X	64,000 ^f
	E4G-200	3,000 ^g
	E4G-400	6,000 ^g
	8900-40G6X-xm modules	3,000 ^g
	Summit X440	192
	Summit X480	12,000 ^b
	Summit X460	6,000 ^g
	Summit X670	
	VIM4-40G4x	3,000 ^g
Summit X770	66,500	

Table 3: Supported Limits (Continued)

Metric	Product	Limit
PIM IPv4-SSM (maximum SSM routes) —maximum number of (S,G) entries installed in the hardware with PIM SSM configuration (IP multi-cast compression disabled). NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 ^f
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^f
	8900-G96T-c modules	4,096 ^f
	8900 xl-series	15,000
	8900-40G6X-xm	3,000 ^g
	BlackDiamond X8 a-series modules	4,094
	BlackDiamond X8-100G4X modules	16,384
	E4G-200	2,048
	E4G-400	2,048
	Summit X440	64
	Summit X480	4,096
	Summit X460	2,048
Summit X670		
VIM4-40G4x	3,000 ^g	
Summit X770	4,096	
PIM IPv4-SSM (maximum SSM routes) —maximum number of (S,G) entries installed in the hardware with PIM SSM configuration (IP multi-cast compression enabled). NOTE: Assumes source-group-vlan mode. For additional limits, see: Layer-3 IPMC forwarding caches —(PIM, MVR, PVLAN) in mixed-mode. ^g on page 59	BlackDiamond 8800 c-series	6,000 ^g
	BlackDiamond 8000 e-series	500 ^g
	BlackDiamond 8900 c-series	6,000 ^g
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 a-series modules	6,000 ^g
	BlackDiamond X8-100G4X modules	64,000 ^g
	E4G-200	3,000 ^g
	E4G-400	6,000 ^g
	8900-40G6X-xm	3,000 ^g
	Summit X440	192 ^g
	Summit X480	12,000 ^b
	Summit X460	6,000 ^g
	Summit X670	
	VIM4-40G4x	3,000 ^g
Summit X770	66,500	

Table 3: Supported Limits (Continued)

Metric	Product	Limit
PIM IPv6 (maximum routes) —maximum number of (S,G) entries installed in the hardware. NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	1,000
	BlackDiamond 8800 e-series	250
	BlackDiamond 8900 c-series	1,000
	BlackDiamond 8900-40G6X-xm	1,000
	BlackDiamond 8900 xl-series	3,000
	BlackDiamond X8 a-series modules	3,000
	BlackDiamond X8-100G4X modules	64,000 ^d
	E4G-200	1,500
	E4G-400	3,000
	Summit X440	90
	Summit X460	3,000
	Summit X480	3,000
	Summit X670	1,500
Summit X770	4,096	
PIM IPv4 (maximum interfaces) —maximum number of PIM active interfaces.	All platforms	512
PIM IPv4 (maximum interfaces) —maximum number of PIM snooping enabled interfaces.	All platforms	512
PIM IPv4 Limits —maximum number of multi-cast groups per rendezvous point	All platforms	180
PIM IPv4 Limits —maximum number of multi-cast sources per group	All platforms	175
PIM IPv4 Limits —maximum number of dynamic rendezvous points per multi-cast group	All platforms	145
PIM IPv4 Limits —static rendezvous points	All platforms	32
PIM IPv6 (maximum interfaces) —maximum number of PIM active interfaces	All platforms	512
PIM IPv6 Limits —maximum number of multicast group per rendezvous point	All platforms	70
PIM IPv6 Limits —maximum number of multicast sources per group	All platforms	43
PIM IPv6 Limits —maximum number of dynamic rendezvous points per multicast group	All platforms	64

Table 3: Supported Limits (Continued)

Metric	Product	Limit
PIM IPv6 Limits —maximum number of secondary address per interface	All platforms	70
PIM IPv6 Limits —static rendezvous points	All platforms	32
Policy-based routing (PBR) redundancy —maximum number of flow-redirects.	All platforms	256 ^m
Policy-based routing (PBR) redundancy —maximum number of next hops per each flow-direct.	All platforms	32 ^m
Port-specific VLAN tags —maximum number of port-specific VLAN tags	All platforms	1,023
Port-specific VLAN tags —maximum number of port-specific VLAN tag ports	BlackDiamond X8 and BlackDiamond 8800 xl-series Summit X480 Summit X460-48t Summit X460-24x, X670-48x Summit X670V-48t Summit X670v-48t stack Summit X770 E4G-400 E4G-200	8,090 3,800 7,200 3,400 3,600 7,200 6,400 3,400 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 8900-G96T-c modules BlackDiamond X8 series Summit series	383 767 767 One less than the number of available user ports
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, X460, X480 Summit X670 Summit X440 All other platforms	1,024 2,046 256 512

Table 3: 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	2,046
	Summit X670	597
	Summit X460 Summit X770	820 1,282
PTP/1588v2 Clock Ports	Summit X770 and E4G cell site routers	32 for boundary clock 1 for ordinary clock
PTP/1588v2 Clock Instances	Summit X770 and E4G 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 and E4G cell site routers	40 entries per clock port
PTP/1588v2 Unicast Static Masters	Summit X770 and E4G 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 neighbors —maximum number of RIP neighbors.	E4G-200	256
RIP interfaces on a single router —recommended maximum number of RIP routed interfaces on a switch.	BlackDiamond 8000 series	256
	BlackDiamond X8 series	256
	BlackDiamond 8900 xl-series	384
	Summit X440	128
	Summit X460	256
	Summit X480	384
	Summit X670, X770 E4G-400	256 256

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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	3,000
	Summit X770	10,000
	E4G-400	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	256
	Summit X460, X480, X440	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
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.	All platforms (except Summit X460, X440, X430, and E4G-400)	500
	Summit X460 and E4G-400	600
	Summit X440	250
	Summit X430	100
Spanning Tree —maximum number of VLANs on all MSTP instances.	All platforms (except Summit X460, X440, X430, X770, and E4G-400)	1,000
	Summit X460, X770, 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

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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 and 8900 series	1,024
	Summit X770	1,024
	Summit X460 and E4G-400	600
	Summit X430	128
	All other platforms	560
SSH (number of sessions) —maximum number of simultaneous SSH sessions.	All platforms	8
Static MAC multi-cast FDB entries —maximum number of permanent multi-cast MAC entries configured into the FDB.	BlackDiamond 8000 c-, e-, xl-series	1,024
	BlackDiamond X8 series	1,024
	Summit X460, X480, X670, X430	1,024
	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

Table 3: Supported Limits (Continued)

Metric	Product	Limit
TRILL —nodes	BlackDiamond X8	256
	Summit X670, X770	256
TRILL —links	BlackDiamond X8	2,000
	Summit X670, X770	2,000
Virtual routers —maximum number of user-created virtual routers that can be created on a switch. NOTE: Virtual routers are not supported on Summit X440 series switches.	BlackDiamond 8000 c-, xl-, xm-series	63
	BlackDiamond X8 series	63
	E4G-200, E4G-400	63
	Summit X460, X480, X670, X770	63
Virtual router forwarding (VRFs) —maximum number of VRFs that can be created on a switch. NOTE: * Subject to other system limitations.	BlackDiamond 8000 c-, xl-, xm-series	960 *
	BlackDiamond X8 series	
	Summit X460, X480, X670, X770	
	E4G-400, E4G-200	
Virtual router protocols per VR —maximum number of routing protocols per VR.	All platforms	8
Virtual router protocols per switch —maximum number of VR protocols per switch.	All platforms	64
VLAN aggregation —maximum number of port-VLAN combinations on any one superVLAN and all of its subVLANs.	All platforms (except Summit X440)	1,000
	Summit X440	256
VLANs —includes all VLANs.	All platforms	4,094
	Summit X670, X670V-48t, BlackDiamond 8800 xl-series only, BlackDiamond X8 series	1,023
	Summit X460, X770, X480, E4G-400, BlackDiamond X8 xl-series	4,093
VLANs —maximum number of port-specific tag VLANs.	E4G-200	2,047
	Summit X460, X670, X670V-48t, BlackDiamond 8800 xl-series only, BlackDiamond X8, E4G-400, E4G-200	4,096
	BlackDiamond X8 xl-series	32,767
	Summit X770	8,192
VLANs —maximum number of port-specific tag VLAN ports	Summit X480	16,383
	All platforms	4,094
VLANs (Layer 2) —maximum number of Layer 2 VLANs.	All platforms	4,094
VLANs (Layer 3) —maximum number of Layer 3 VLANs.	All platforms	2,000

Table 3: Supported Limits (Continued)

Metric	Product	Limit
VLANs (maximum active port-based) —(Maximum active ports per VLAN when 4,094 VLANs are configured with default license)	Summit X670, X480, X460, X770	32
	E4G-400	32
	Summit X440	16
	E4G-200	12
	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 c-, e-, xl-series with eight modules of 48 ports 8900-G96T-c modules	383 767
	BlackDiamond X8 series	767
	Summit series	One less than the number of available user ports
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, X460, X480	1,024
	Summit X670	2,046
	Summit X440	256
	All other platforms	512
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	512
	Summit X480, X670, X770	2,046
Summit X430	100	
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	511
	E4G-200, E4G-400	255
	Summit X460	255
	Summit X480	511
	Summit X440	30

Table 3: Supported Limits (Continued)

Metric	Product	Limit
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	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	7
VRRP (v2/v3-IPv4/IPv6) (maximum VRIDs per VLAN) —maximum number of VRIDs per VLAN.	All platforms with Advanced Edge license or higher	7
VRRP (v2/v3-IPv4/IPv6) (maximum ping tracks) —maximum number of ping tracks per VLAN.	All platforms with Advanced Edge license or higher	8
VRRP (maximum ping tracks) —maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher.	All platforms, except the Summit X440	8 (20 centisecond or 1 second hello interval)
	Summit X440 Hello interval: 20 centiseconds	2
	Summit X440 Hello interval: 1 second	4
VRRP (v3-IPv6) (maximum ping tracks) —maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher.	All platforms, except the Summit X440	8 (20 centisecond or 1 second hello interval)
	Summit X440 Hello interval: 20 centiseconds	1 (IPv6)
	Summit X440 Hello interval: 1 second	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	8
VRRP (v2/v3-IPv4/IPv6) —maximum number of VLAN tracks per VLAN.	All platforms with Advanced Edge license or higher	8

Table 3: Supported Limits (Continued)

Metric	Product	Limit
XML requests —maximum number of XML requests per second. NOTE: Limits are dependent on load and type of XML request. These values are dynamic ACL data requests.	BlackDiamond 8800 c-series with 100 DACLs with 500 DACLs	10 3
	BlackDiamond 8900 series with 100 DACLs with 500 DACLs	10 3
	Summit X480, X670 with 100 DACLs with 500 DACLs	4 1
XNV authentication —maximum number of VMs that can be processed (combination of local and network VMs).	All platforms	2,048
XNV database entries —maximum number of VM database entries (combination of local and network VMs).	All platforms	16,000
XNV database entries —maximum number of VPP database entries (combination of local and network VPPs).	All platforms	2,048
XNV dynamic VLAN —Maximum number of dynamic VLANs created (from VPPs /local VMs)	All Platforms	2,048
XNV local VPPs —maximum number of XNV local VPPs.	All platforms (except Summit X430) Ingress Egress	2,048 512
	Summit X430 Ingress	1,024
XNV policies/dynamic ACLs —maximum number of policies/dynamic ACLs that can be configured per VPP. ⁿ	All platforms (except Summit X430) Ingress Egress	8 4
	Summit X430 Ingress	8
XNV network VPPs —maximum number of XNV network VPPs. ⁿ	All platforms (except Summit X430) Ingress Egress	2,048 512
	Summit X430 Ingress	1,024

- The table shows the total available.
- Limit depends on setting configured for `configure forwarding external-tables`.
- When there are BFD sessions with minimal timer, sessions with default timer should not be used.
- Based on forwarding internal table configuration "I2-only".
- Effective capacity varies based on actual MAC addresses and VLAN IDs used and hash algorithm selected.

- f. Applies only if all enabled BlackDiamond 8000 I/O modules are BlackDiamond 8000 c-, xl-, or xm-series modules.
- g. 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, X480, and X670 switches compared to BlackDiamond 8000 e-series modules.
- h. 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.
- i. Based on forwarding internal table configuration "I3-and-ipmc".
- j. The limit depends on setting configured with `configure iproute reserved-entries`.
- k. The IPv4 and IPv6 multi-cast entries share the same hardware tables, so the effective number of IPv6 multi-cast entries depends on the number of IPv4 multi-cast entries present and vice-versa.
- l. If IGMP and MLD are simultaneously configured on the switch, the number of effective subscribers supported would be appropriately lessened.
- m. Sum total of all PBR next hops on all flow redirects should not exceed 1024.
- n. 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.



NOTE

Extreme Networks is transitioning to a new software defect numbering system. Previously, software defect ID numbers were prefaced with the letters “PD”; they will now be prefaced with “XOS.” During this transition period, some software defects will have the old format ID and some will have the new one.

This chapter contains the following sections:

- [Open Issues on page 84](#)
- [Known Behaviors on page 89](#)
- [Resolved Issues in ExtremeXOS 15.5.5-Patch1-1 on page 90](#)
- [Resolved Issues in ExtremeXOS 15.5.5 on page 93](#)
- [Resolved Issues in ExtremeXOS 15.5.4-Patch1-5 on page 96](#)
- [Resolved Issues in ExtremeXOS 15.5.4-Patch1-4 on page 97](#)
- [Resolved Issues in ExtremeXOS 15.5.4 on page 101](#)
- [Resolved Issues in ExtremeXOS 15.5.3-Patch1-6 on page 104](#)
- [Resolved Issues in ExtremeXOS 15.5.3-Patch1-5 on page 105](#)
- [Resolved Issues in ExtremeXOS 15.5.3-Patch1-2 on page 107](#)
- [Resolved Issues in ExtremeXOS 15.5.3 on page 109](#)
- [Resolved Issues in ExtremeXOS 15.5.2-Patch1-5 on page 111](#)
- [Resolved Issues in ExtremeXOS 15.5.2-Patch1-1 on page 113](#)
- [Resolved Issues in ExtremeXOS 15.5.2 on page 115](#)
- [Resolved Issues in ExtremeXOS 15.5.1 on page 120](#)

Open Issues

The following are the open issues for supported features in ExtremeXOS 15.5.5-Patch1-1.

Table 4: Open Issues, Platform-Specific and Feature issues

ID Number	Description
General	
xos0057679	Account and password information is not encrypted while configuring in ExtremeXOS.
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.
xos0055311	IPARP entries are cleared immediately after they are learned from source VLAN port when IPARP timeout is disabled.
xos0054261	It can take up to 7 seconds for ISC status to come up after removing MD5 authentication from both peers. This also occurs when MD5 keys are added on both peers.
PD4-4246500781	EMS logs are not available for DHCPv6 relay agent remote ID option feature.
PD4-4246500792	CLI counters are not available for DHCPv6 relay agent remote ID option feature.
PD4-4502231710	Device manager is reporting incorrect slot type after clearing a slot while that slot is disabled.
PD4-4496325110	Cannot create an IPv6 SNMP target address using the <code>configure snmpv3 add target-addr</code> command. The command does not recognize a correctly formatted IPv6 address for the target address. The following error appears: <pre>%% Ambiguous command: "configure snmpv3 add targetaddr t1 param p1 ipaddress 2001::1"</pre> Other forms of the command are accepted.
PD4-4519100637	Cannot configure I/O slot module type on disabled slots.
BlackDiamond 8800 Series Switches	
PD4-4332910219	Switches receive more packets than are sent over VPLS/MLAG connection after master switch fabric module failover.
PD4-4497435810	Issuing the command <code>clear slot</code> on a disabled slot does not clear the pre-existing port numbers associated with the previous module type in that slot. I/O module stays in VLAN sync.
PD4-4534887951	L2PT is not tunneling PVST+ packets received on BD8900 XL modules. The issue is specific to XL modules.

Table 4: Open Issues, Platform-Specific and Feature issues (Continued)

ID Number	Description														
BlackDiamond X8 Series Switches															
PD4-4539773451	<p>Rx PFC is not working for some priorities. When all QoS profiles QP1 through QP8 are configured and mapped to dot1p 0 to 7 in order, PFC packets pause traffic according to the following table:</p> <table> <tr> <td>PFC priority</td> <td>Pauses traffic priority</td> </tr> <tr> <td>0,1,2</td> <td>no effect</td> </tr> <tr> <td>3</td> <td>0</td> </tr> <tr> <td>4</td> <td>1,5</td> </tr> <tr> <td>5</td> <td>2,6</td> </tr> <tr> <td>6</td> <td>5,7</td> </tr> <tr> <td>7</td> <td>4</td> </tr> </table>	PFC priority	Pauses traffic priority	0,1,2	no effect	3	0	4	1,5	5	2,6	6	5,7	7	4
PFC priority	Pauses traffic priority														
0,1,2	no effect														
3	0														
4	1,5														
5	2,6														
6	5,7														
7	4														
PD4-4480905341	<p>On BlackDiamond X8 series switches configured with OSPF, MPLS, RSVP, LSPs, and RSVP-TE paths, clearing counters produces the following error:</p> <pre>"MM-A: ILM instance 872677376 not found 03/13/2014 09:26:26.68 <Error:HAL.MPLS.Error> MM-A: ILM instance 889454592 not found"</pre>														
PD4-4520528421	<p>PSTag: FDB entries are not learned correctly when you delete all ports from the VLAN, and then add the ports back.</p>														
PD4-4431208163	<p>The command <code>show identity-management entries</code> shows the IP addresses in reverse order.</p>														
PD4-4440553141	<p>Unable to configure sharing port-based key using wildcard on 100G4X or the 40G24X modules.</p>														
PD4-4356518461	<p>The BlackDiamond X8-100G4X module does not support TRILL (known limitation), and the module continuously reboots if TRILL is enabled on that slot.</p>														
PD4-4398103442	<p>Incorrect frame delay measurements occur on BlackDiamond X8 and 8800 series switches since kernel time stamping for software end points isn't supported.</p>														
PD4-4400091092	<p>Configuring, and then unconfiguring, large ACL policy files two or three few times causes the ACL to fail during subsequent installs.</p>														
PD4-4411209261	<p>Running I/O diagnostics on a dual management modules, produces <code>chassispiBConduitMasterRcvOneSlot</code> error messages.</p> <p>NOTE: The error message can be ignored.</p>														
Summit Family Switches															
PD4-4147204751	<p>Untagged L2 traffic stops when one-to-many remote mirroring is enabled, and then it doesn't resume when the mirroring is disabled</p> <p>Workaround: Remove affected ports from the VLAN, and then add them back.</p>														
PD4-4458472124	<p>The command <code>clear license-info all</code> does not work.</p> <p>Workaround: Use <code>clear license-info</code>.</p>														

Table 4: Open Issues, Platform-Specific and Feature issues (Continued)

ID Number	Description
SummitStack	
PD4-4484035633	On Summit X480-24x(SSV80)/X670v stacks, MPLS process ends unexpectedly with signal 6 when enabling the stacking by using the command <code>config stacking easy-setup</code> in the switch with core and MPLS feature pack license. This issue does not occur with Summit X460 and X440 stacks.
PD4-4364521713	IPv6 ping fails in stacks and chassis in PStag configuration.
Summit X430 Series Switches	
xos0057691	AVB does not work on the Summit X430-8p and X430-24p switches.
Summit X440 Series Switches	
PD4-3797993347	Running failover more than two or three times on a Summit X440 stack ceases to update the counter in the new master. You must restart the IdMgr process to update the counter. Even if the IDM roles are mapped correctly, you sometimes receive spurious values after restarting the netlogin process. Workaround: Restart the idmgr process.
PD4-3918240371	Executing a MAC address move on a role-based ID-enabled port, and then changing it back to the original port, causes a loop as log messages are logged continuously, which are particular to the role-based VLAN, even though one port is inactive. NOTE: Issue resolves without intervention in approximately 10 minutes.
Summit X480 Series Switches	
PD4-4507367531	Layer 2 VPN (VPWS and VPLS) tunnels are broken when routed path is changed on the transit node by disabling ports.
Summit X670 Series Switches	
xos0054901	Traffic fails to flow when two links between two rbridges have different speeds (and metrics).
PD4-4521256640	On Summit X670 switches with BlackDiamond X8 switches as neighbors, the L2PT received tunneled packets are not tunneled towards customer edge. This issue occurs only in VMANs.
PD4-4394772440	Traffic loss occurs after creating the same TRILL nickname on two switches, The frequency of occurrence is at least 70%.
PD4-4323714132	On Summit X670v-48x switches, port comes up in active state when a BASET optics is inserted without copper cable and the port is disabled, switch is saved and rebooted, and then the port is enabled.
PD4-4363911031	On Summit X670-48x switches, ports 47 and 48 with stacking support enabled do not come up when an SFP+ transceiver (optics) module is removed, and then reinserted. Workaround: Reboot the switch.

Table 4: Open Issues, Platform-Specific and Feature issues (Continued)

ID Number	Description
Summit X770 Series Switches	
xos0055686	VPLS: On Summit X770 series switches, IGMP hello packets received by pseudowires are not forwarded to service ports.
xos0055657	On Summit X770-32q switches. partitioned 10G ports are not becoming active when disabling/enabling ports with mirroring configuration. Ports 3 and 4 only become active state after 25-35 seconds.
PD4-4464975979	On Summit X770 series switches, partitioned 10G ports 77, 78, 81, and 82 do not appear in output of the command <code>edp ports all</code> . As a result, bi-directional traffic across these ports does not flow, even though these ports are up with speed 10G. Workaround: Use ports 49-52 for 40G, unless ports 1-53 are all partitioned for 10G. Use ports 101-104 for 40G, unless ports 53-104 are all partitioned for 10G. When stack support is enabled on ports 103-104, ports 101-102 should be configured in 40G mode.
ACLs	
PD4-4377098911	When ACLs are applied in both ingress and egress directions, you cannot see egress direction via SNMP. When a policy has more than one counter, via SNMP, you can only check the updates from the first counter, and subsequent counters do not appear.
BGP	
PD4-4442759045	With BGP session established, and with routes advertised, between switches, and then after using a policy to block the routes, route refresh capability (<code>configure bgp neighbor <neighbor-id> soft-reset in</code>) does not work, and the routes persist. Workaround: Refresh policy.
ERPS	
PD4-4501394621	ERPS rings do not move to Idle state after disabling, and then enabling revertive mode.
IP Protocols	
PD4-4445294484	Partial traffic loss occurs after the primary PE goes down in MPLS/BGP network where CE device is DUAL homing to PEs.
PD4-4476964651	Multi-peer MLAG-PIM: PIM process ends unexpectedly on an MLAG peer when one of its neighbors is rebooted. Traffic is running for 2,000 (S;G)s.
PD4-4454371832	PIM-SSM: It takes ~30 seconds until the (*;G)s are created when PIM-SSM is disabled/no SSM range is configured. IGMPv3 using record type 2 are sent.
PD4-4345409499	Multi-peer MLAG-PIM: (S;G)s for a source located behind an edge switch never expires.
PD4-4388000494	Multi-peer MLAG-PIM: After assert is triggered, the first-hop router forwards traffic using two paths. This state lasts for 210 seconds. The (S;G) from an edge MLAG peer is not asserted.

Table 4: Open Issues, Platform-Specific and Feature issues (Continued)

ID Number	Description
MPLS	
PD4-4576305291	On BlackDiamond XB-100G4X and Summit X770 switches, VPLS pseudowire Tx packet counters are incrementing incorrectly when unknown uni-cast, multi-cast or broadcast traffic is received. This traffic should only increment the Rx counters.
PD4-4326073541	With eight static LSPs configured, the command <code>show iproute mpls</code> shows only seven routes from MPLS.
PD4-4502231831	If MPLS VPLS/VPWS pseudowire ports are set to 30, traffic is not forwarded.
OSPFv3	
xos0057538	OSPFv3 fails to select the best cost external route.
Security	
xos0061068	In Netlogin post-client authentication, IP addresses are not synced with backup Master Switch Fabric Module.
xos0061008	Static FDB configured on Netlogin ports is lost after rebooting the switch.
PD4-4503475386	NetloginMac: After executing the commands <code>clear netlogin state</code> and <code>clear fdb</code> , netlogin MAC address authenticated entries are not getting blackholed. Workaround: Re-create FDB blackhole entries for netlogin authenticated MAC addresses.
PD4-4503482301	Netlogin Dot1x: Changing the re-authentication value after the client is authenticated is not reflected in the output of the <code>show netlogin dot1x</code> command. A new re-authentication timer value is not taken accepted immediately, and only logging the client off, and then back on again, causes the new re-authentication timer value to be taken into account.
PD4-4396321375	Log messages are not generated when SSH2 access is rejected by access-profile.
PD4-4521639362	Authentication fails for a netlogin client in dot1x mode, since the port added as untagged in one VLAN cannot be moved to another VLAN.
ScreenPlay	
PD4-4519257691	SNMPv1/v2 communities and SNMPv3 users do not appear. The following error message appears: "Data could not be retrieved for SNMP stats. Either switch is not accessible or web service is not enabled."

Known Behaviors

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

Table 5: Known Behaviors, Platform-Specific and Feature Issues

ID Number	Description
General	
xos0053752	PSTag VLAN-VPLS port is accepting 802.1ad frames instead of dropping them when the same port is added to tagged VMANs.
PD4-4490673726	USM users created via SNMP with 3des/aes192/aes256 privacy keys are lost when upgrading from ExtremeXOS 15.4 to 15.5.
PD4-4249647562	Configuring multiple protocol filters on single ports does not produce any warning or error message, but it replaces the previous configurations.
BlackDiamond Series Switches	
PD4-4231603031	The BlackDiamond X8 100G4X modules cannot learn a stream of unique MAC addresses at line rate. Traffic with various MAC addresses or received at below line rate are correctly learned.
PD4-4379392431	On BlackDiamond 8800 series switches, after disabling/removing I/O slots, incorrect neighborship appears in the new ESRP master.
PD4-4405960220	On BlackDiamond X8 series switches, I/O modules remain operational while running diagnostics on all fabric module slots simultaneously. This could potentially cause issues on the network as other switches still see ports as active and continue to try and pass traffic. Disabling all the fabric module slots puts all the I/O modules into the failed state until at least one fabric module is enabled again. Workaround: Run diagnostics on fabric module one at a time.
Summit Series Switches	
PD4-4498257471	On Summit X670 series switches, multi-cast traffic is completely dropped on customer edge port (added to VPLS) when another VLAN is also configured on that port with IGMP snooping enabled. No traffic is dropped if you disable IGMP snooping for that VLAN. Workaround: Disable IGMP snooping for that VLAN on all ports.
PD4-4350056231	When downgrading from ExtremeXOS 15.5.1 to 15.4.1 after some configuration (QoS configuration), Summit X770 series switches take a long time to load configuration in ExtremeXOS 15.4.1 image.
PD4-4376759909	On Summit X460 series switches, when executing the commands <code>disable learning</code> or <code>disable learning drop-packets port all</code> when executing the command <code>incremental source mac unknown destination mac L2 traffic</code> to flood traffic to other VLAN ports, egress ports in the same VLAN take more than 30 seconds to stop forwarding packets.

Resolved Issues in ExtremeXOS 15.5.5-Patch1-1

The following issues were resolved in ExtremeXOS 15.5.5-Patch1-1. ExtremeXOS 15.5.5-Patch1-1 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.4.2, and ExtremeXOS 15.5.5. For information about those fixes, see the release notes for the specific release.

Table 6: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.5-Patch1-1

ID Number	Description
General	
xos0050590	Multicast cache entries are not programmed in hardware due to false resource full condition due to IPv6 cache entries.
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.
xos0059042	Kernel panic occurs during FDB flush operations caused by deleting private VLANs.
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.
xos0061507	SNMPget on EXTREME-SOFTWARE-MONITOR table returns value with incorrect OID.
xos0061781	Identity manager entries become stale when clients are moved from one port to another in sub-VLANs.
xos0061788	The process devmgr ends unexpectedly during snmpwalk when continuous EMS logs are sent to the switch console.
xos0061855	Configured OSPF neighbor is not retained after rebooting.
xos0062017	DHCP trusted port configuration is lost after disabling, and then re-enabling LAG.
xos0062240	Port that was administratively disabled becomes up after enabling rx pause.
xos0062537	HAL crash occurs when redirect-port-list action contains more than 64 ports.
xos0062618	ELRP forgets the disabled port information if the port is deleted from another VLAN that also has ELRP enabled. As a result, the disabled port stays disabled unless manually enabled.

Table 6: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.5-Patch1-1 (Continued)

ID Number	Description
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.
xos0063257	Saving configuration fails/times-out when VLANs added to a mirror filters are renamed.
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 address in CFM domain does not return information about destination switch.
xos0057538	OSPFv3 fails to select the best cost external route.
xos0057574	After multiple disable/enable OSPFv3, OSPFv3 routes are not advertised to route manager.
xos0057575	OSPFv3 external routes are not updated in the routing table after link flap events.
xos0057583	In OSPFv3, after adding an ASBR, new route is received, but it is not added to routing table even though it is a best route.
xos0057584	When there are equal cost ASBR routes, OSPFv3 is sharing only one path to route manager.
xos0058555	Inconsistent behavior allowed in secondary EtherType settings on LAG ports when present on both tagged VLAN and tagged VMAN.
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.
xos0062494	Source MAC addresses learned through MVRP packets on a blocked port (STP) cause traffic to be dropped.
xos0062508	When a port is added in a loopback VLAN, OSPFv3 route is not advertised with /128 mask.
xos0062754	VPLS traffic egresses out with dot1q tag when secondary EtherType is configured.

Table 6: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.5-Patch1-1 (Continued)

ID Number	Description
xos0063207	Error occurs while adding LAG ports as tagged in one VMAN and untagged in another VMAN, even though the VMAN EtherType is primary.
xos0062367	ACL process ends unexpectedly on repeated refresh of ACL policy with clear-flow action.
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.
BlackDiamond X8 Switches	
xos0054860	Configuring and unconfiguring large ACL policy files two or three times causes the ACL to fail during subsequent installs.
Summit X440 Series Switches	
xos0062621	On Summit X440-8p switches, the <code>show fan</code> command output displays that the fan is unsupported.
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.

Resolved Issues in ExtremeXOS 15.5.5

The following issues were resolved in ExtremeXOS 15.5.5. ExtremeXOS 15.5.5 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, and ExtremeXOS 15.5.4.2. For information about those fixes, see the release notes for the specific release.

Table 7: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.5

ID Number	Description
General	
xos0055923	Traffic not forwarded to other ports in VMAN after deleting service VMAN from VPLS instance.
xos0056191	Executing the command <code>show temperature</code> from a user-created account returns an error.
xos0056885	IGMPv3 packets with invalid checksum are not dropped.
xos0057234	SNMP Informs are not being sent when there are unsent informs remaining in the queue.
xos0057264	Informs are not sent when there is only inform receiver configured on the switch.
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.
xos0058135	Tagged VLAN traffic is inappropriately grabbed by VMANs after removing service VLANs from VPLS instances. This happens when the same port is tagged in the service VLAN and untagged in the VMAN.
xos0058775	snmpMaster process ends unexpectedly while fetching system MAC and boot count information simultaneously from switch.
xos0060622	The following error message appears after disable/enable sharing: <code>extreme_mpls_del_vp_nh_xref: MPLS ECMP NH vp : 0xddf13intf = 0xebf map entry not found in btree rv = -7 (Entry not found)</code>
xos0061092	Traffic forwarding on VPLS-serviced VMAN stops after link flap.
xos0061198	Disabling of one VPN-VRF instance affects traffic flowing on another VPN-VRF
xos0062128	L3VPN traffic is not forwarded after executing <code>disable port</code> and <code>enable port</code> in MPLS core network.

Table 7: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.5 (Continued)

ID Number	Description
xos0062427	EDP process ends unexpectedly when CDP packets without portId TLV are received.
xos0062499	Multicast packets are dropped in Layer 2 bridged VLANs.
xos0062719	Allow use of 3rd-party optics without any addition license.
xos0062914	The process mcmgr ends unexpectedly after receiving corrupted IGMPv3 join packets on MLAG ports.
xos0058808	Rarely, MAC addresses of authenticated clients learned on NetLogin-enabled ports are not programmed in hardware.
xos0059942	SSH connection ends when show commands produce lengthy output.
xos0060103	SNMP walk does not return all VLANs under extremePortVlanStatsTable.
xos0062260	BGP process ends unexpectedly when local address or password is changed for BGP neighbor, and then you immediately execute a BGP show/configuration command.
xos0062271	CLI memory leak occurs when executing show commands with include option through script.
xos0062380	Switch rejects incorrect LSP configurations as expected, but this operation still uses LSP indexes in hardware.
xos0062472	Source MAC addresses learned through CDP packets received on EAPS-blocked ports cause traffic to be dropped.
xos0062710	On BlackDiamond 8800 or BlackDiamond X8 series switches, with Distributed IP ARP mode on, ECMP routes are sometimes not installed when gateways flap.
xos0061886	SNMP master process ends unexpectedly with signal 6 with certain sequence of snmpbulkget and snmpget.
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.
E4G-200 Cell Site Routers	
xos0056078	Rx packets counter is not incrementing correctly as shown in the output of the command <code>show mpls ldp lsp prefix detail</code> .
Summit X670 Series Switches	
xos0061167	Links become active without a connection with tri-speed Base-T SFP installed.

Table 7: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.5 (Continued)

ID Number	Description
SummitStack	
xos0057130	Mirroring instances are disabled in SummiStacks after a second or third failover.
xos0061067	In 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.
BlackDiamond X8 Series Switches	
xos0056773	Resetting a connected I/O fabric module, produces the following error messages: <pre>04/07/2014 10:09:10.07 <Erro:Kern.Card.Error> Slot-7: pca9506_read, not enough buffer provided (need 5 bytes, have 1 bytes) 04/07/2014 10:09:10.57 <Erro:Kern.Card.Error> Slot-7: i2c-1: bus busy, addr=0x23 rw=1 cmd=0x0 size=2 retry=0 04/07/2014 10:09:10.57 <Erro:Kern.Card.Error> Slot-7: Failed to read data from INPUT_PORT_0_COMMAND (-145)</pre>
xos0060426	"Unable to read QSFP info" error messages appear when reading EEPROM from optics.
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.
BlackDiamond 8800 Series Switches	
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).

Resolved Issues in ExtremeXOS 15.5.4-Patch1-5

The following issues were resolved in ExtremeXOS 15.5.4-Patch1-5. ExtremeXOS 15.5.4-Patch1-5 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8 and ExtremeXOS 15.5.4.2. For information about those fixes, see the release notes for the specific release.

Table 8: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4-Patch1-5

ID Number	Description
General	
xos0057464	TLS protocol is impacted by CRIME (Compression Ratio Info-Leak Made Easy) vulnerability.
xos0061385	EAPS process ends unexpectedly after deleting EAPS shared-port configuration.
xos0062018	For IPv6 routes with mask lengths greater than 64-bits, IPv6 unicast packets destined for the switch CPU can be dropped if another IPv6 route is present with a matching prefix and mask length less than or equal to 64-bits. This issue affects Summit X440, X460-G2, X670-G2, X770 switches, BlackDiamond 8800 G48Te2 I/O modules, and BlackDiamond X8 100G I/O modules.
BlackDiamond 8800 Series Switches	
xos0062060	FDB entries are not programmed in hardware even though hardware resources have sufficient capacity.
BlackDiamond X8 Switches	
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.
Summit X460 Series Switches	
xos0059688	Default debounce timer value (in ms) for alternate stacking port should be "0" instead of "150" in Summit X460-48p switches.

Resolved Issues in ExtremeXOS 15.5.4-Patch1-4

The following issues were resolved in ExtremeXOS 15.5.4-Patch1-4. ExtremeXOS 15.5.4-Patch1-4 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8 and ExtremeXOS 15.5.4.2. For information about those fixes, see the release notes for the specific release.

Table 9: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4-Patch1-4

ID Number	Description
General	
xos0053821	IPv6 neighbor advertisements for VRRP virtual IP address uses virtual MAC address as source MAC address instead of switch MAC address.
xos0055376	After executing the commands <code>clear netlogin state</code> and <code>clear fdb</code> , NetLogin MAC address authenticated entries are not blackholed.
xos0055514	The process "bcmRX" consumes 25% of CPU with one-to-many mirroring feature enabled.
xos0055680	Counters in output of <code>show ipstat</code> command are not incremented for ingress traffic.
xos0055945	Error message "exosmc: ip_mc_handle_msdp_data:2038: MC: Ingress vif not found" appears when sending multicast tagged packets to member VLAN of translation-VLAN/ subscriber VLAN of private VLAN.
xos0056243	OSPF process ends unexpectedly during frequent route re-calculation caused by switch reboot or MSM failover or BFD flap events.
xos0056340	Unknown Layer 2 traffic from Isolated subscriber VLANs are forwarded to the remote MLAG ports, even though local MLAG ports are up.
xos0056553	The output of the command <code>show fdb netlogin all</code> does not show 's' or 'd' flags for netlogin entries.
xos0056874	Error message "radDecodeVsa: Unknown vendor 311" appears even though RADIUS and dot1x authentication is successful.
xos0057030	ACL process ends unexpectedly after executing the command <code>refresh access-list network-zone</code> .
xos0058188	LACP member ports in the remote end are removed from the aggregator when back-to-back MSM failovers are executed in the local end.

Table 9: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4-Patch1-4 (Continued)

ID Number	Description
xos0058671	In VPLS, traffic destined to unknown MAC addresses is duplicated after changing dot1q tag include or exclude configuration.
xos0059266	VPLS traffic is not switching over to alternate RSVP LSP after disabling the active LSP.
xos0059655	Error "Unable to delete permanent entry" appears while deleting static blackhole FDB entries.
xos0059730	The process mcmgr ends unexpectedly after removing a slot from existing SummitStack and unconfiguring that slot using the command <code>unconfigure slot slot_number</code> .
xos0059924	The output of the command <code>show access-list meter ports</code> displays additional meter name when only one meter is applied using ACL policy.
xos0059989	Configuring non-persistent command using UPM script shows dirty bit(*) in the prompt.
xos0060075	Snmpwalk on <code>extremePortLoadShare2Status</code> returns incomplete information after adding a port in sharing group with a port number that is smaller than any of member ports.
xos0060176	The process rtmgr ends unexpectedly with signal 11 because of segmentation fault. This occurs only when default route is exported from BGP neighbor.
xos0060214	Process netTools ends unexpectedly during reboot of switch with VRRP track-ping configuration.
xos0060449	NetLogin MAC-based mode does not work as expected with PVLANS.
xos0060716	Need support for new ACL action "redirect-vlan" to redirect matched packets to all ports in specified VLANs.
xos0060780	With VRRP enabled, local VLAN's direct route is not installed in hardware after reconfiguring the VLAN's IP address.
xos0060794	VRRP advertisement interval configuration changes after upgrading ExtremeXOS from 12.6 to 15.4 or later releases. Issue occurs only when interval is configured in milliseconds.
xos0060909	In UPM profiles the variable <code>EVENT.TIME</code> incorrectly has the current time rather than the time when the event was queued/triggered.
xos0061009	The output of the command <code>show netlogin MAC</code> output displays username for unauthenticated client.
xos0061038	Loops occur in EAPS-protected VLANs, after peer reboot, if a VLAN's port is also protected by ELSM.
xos0061069	In Netlogin ISP mode, client MAC addresses configured as static FDBs are removed after reboot.
xos0061085	Kernel oops occurs while deleting VR with enable BGP export and IPARP proxy configurations.

Table 9: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4-Patch1-4 (Continued)

ID Number	Description
xos0061178	Dynamic ACL for gratuitous ARP violation on LAG member ports are incorrectly getting installed on LAG master ports.
xos0061222	Gratuitous ARP packets for VRRP virtual IP addresses have ARP sender addresses as physical MAC addresses, instead of VRRP virtual MAC addresses.
xos0061565	The TCL function, "clock scan," generates errors with default time zone configuration.
xos0061699	Traffic is dropped when moving idmgr client from one port to another with role-based authentication.
xos0061835	The process exsh causes excessive CPU utilization after performing continuous Telnet/SSH for the switch.
xos0061922	Dynamic ACLs applied as "any" fail to install in hardware after upgrading ExtremeXOS from any release other than EXOS 15.3.
xos0061730	Traffic is affected when unconfiguring ACL on another port.
xos0053869	When ARP requests are sent from an isolated VLAN on one switch to the network VLAN on another switch, no ARP reply occurs after the IP address of the network VLAN is unconfigured on the first switch.
xos0055398	Authentication fails for a Netlogon client in dot1x mode, since the port added untagged in one VLAN cannot be moved to another VLAN.
xos0056263	The following error message occurs when deleting VPLS instances: <Error:Kern.MPLS.Error> : bcm_custom_extr_vfp_tagged_vlan_port_add, Entry exists, rv 0
xos0058611	OSPFv2 external routes are not updated in routing tables after disabling uplink ports in peer switches. During this condition, the route to Advertising Router (external routes) is available by OSPF neighbors, but external routes are not updated dynamically in the routing tables.
xos0059575	Process "exsshd" ends unexpectedly after initiating an SFTP connection using Filezilla.
xos0060407	After disabling, and then enabling EBGP, new routes from different autonomous systems (ASs) are not considered for best path calculation.
xos0061656	Nodes remain in the "FDBSync" state due to temp-flooding while rebooting the stack.
xos0061965	Configuring ESRP member VLANs (VRRP-enabled) produces errors.
Summit X430 Series Switches	
xos0061864	FDB process consumes more than 20% utilization when Summit X430 switches are configured with 100+ VLANs.

Table 9: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4-Patch1-4 (Continued)

ID Number	Description
BlackDiamond 8800 Series Switches	
xos0060891	While sending continuous join and leave, as well as multicast streams multiple times allowing the streams to age out, kernel error messages appear for add/delete operation with reason "Entry not found".
xos0061338	Packets are not switched on port-specific, tag-enabled ports on XL series I/O modules.
xos0061796	Error message "aspenSmlpmcAddEgressPort: group does not exist" appears during switch reboot or MSM failover.
xos0061822	HAL process ends unexpectedly during failover when switches have ACL policies without meter action.
xos0057624	Traffic loss occurs on PVLAN after restarting VRRP process in VRRP Master switch.
BlackDiamond X8 Switches	
xos0057827	Layer 2 multicast traffic is not forwarded to IGMP receivers after MM (management module) failover.
xos0060264	The output of the <code>show port transceiver info</code> command for optics inserted in 40G/100G ports might be abnormally lengthy if the same command is executed from two different CLI sessions simultaneously.
xos0061186	Bytes counter associated with show port utilization command output displays inaccurate value for 100G ports when utilization exceeds 40%.
xos0061902	BlackDiamond X8 series switches use VLAN instance as index instead of router interface (rtif) for ARP entries.
SummitStack	
xos0061957	HAL process ends unexpectedly during failover when switches have ACL policy without meter action.
Summit Family Switches	
xos0060965	Netlogin process ends unexpectedly when web-based users log out and then refresh the logout window a few times.
xos0061191	In SummitStacks with PVLAN configuration, after stack failover, VRRP advertisement has incorrect VLAN ID.
xos0057438	Memory depletion occurs in Backup/Standby nodes of SummitStack with highly scaled IPFIX flow records.
xos0055035	Slot reboot occurs on the stacks due to HAL deadlock while querying media-type.

Table 9: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4-Patch1-4 (Continued)

ID Number	Description
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.
Summit X460 Series Switches	
xos0060517	Untagged ports in service VMANs take other VLANs' traffic after events such as add/delete port in VLANs/VMANs, change tag in VLANs/VMANs, etc.
xos0061180	In Summit X460 stack with mixed alternate and native stacking enabled slots, traffic ingressing one specific slot is not forwarded to other slots.
Summit X670 Series Switches	
xos0061770	Detected parity errors may cause a kernel crash.

Resolved Issues in ExtremeXOS 15.5.4

The following issues were resolved in ExtremeXOS 15.5.4. ExtremeXOS 15.5.4 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8 and ExtremeXOS 15.5.4.2. For information about those fixes, see the release notes for the specific release.

Table 10: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4

ID Number	Description
General	
xos0051961	Unable to block IPv6 traffic from SSH/Telnet/Web interface by access-profile policy.
xos0053634	MAC-lockdown-timeout on user ports does not work as expected if Netlogin is enabled on those ports.
xos0053828	HAL process ends unexpectedly when all entries from network-zone are deleted and associated ACL is refreshed.
xos0054348	Cannot delete flow names after deleting, and then creating, the flow while the ACL is installed.

Table 10: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4 (Continued)

ID Number	Description
xos0056254	Management port remains in down state when peer switch has "auto-neg off" configuration. Issue occurs with Summit X460-G2, X670-G2, X450a, X450e, and BlackDiamond 8800 series switches.
xos0057211	Traffic gets forwarded for blackholed MAC address when limit learning is enabled.
xos0057407	Hops fields in DHCP packets are not incremented when processed by Bootprelay.
xos0058221	Rarely, OSPFv3 process ends unexpectedly with signal 11 when link flaps occur.
xos0058994	PoE is not delivering power to several model phones when legacy mode is enabled.
xos0059288	Cannot install multiple XMOD files in BlackDiamond 8800 and BlackDiamond X8 switches due to lack of space.
xos0059574	OSPF packets larger than 8,192 are dropped even with jumbo frame enabled.
xos0060088	Kernel oops triggered rarely during continuous addition/deletion of ARP entries for long duration in presence of high CPU utilization.
xos0060119	Changing the primary TACACS server configuration locks out TACACS-authenticated users.
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.
xos0060825	Double tagged CFM frames are dropped by kernel in VMAN environment.
Summit Family Switches	
xos0056230	SNMP query on "extremeMemoryMonitorsystemTable" does not show backup information, if slot2 is master and slot1 is backup.
xos0057089	Slot rebooting due to kernel oops after stack failover.
xos0059462	Timezone configuration is not applied to standby nodes after stack reboot.
xos0059950	In Summit series switches, you cannot download bootROM images from memory card.
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.
Summit X430 Series Switches	
xos0059934, xos0057028	Downloading BootROM corrupts the bootloader.

Table 10: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.4 (Continued)

ID Number	Description
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.
BlackDiamond X8 Series Switches	
xos0060210	Rarely, HAL process may end unexpectedly due to buffer overflow condition while running diagnostics for Management Module.
Summit X460 Switches	
xos0059320	CCM is dropped for "Hardware Down MEPs" when they are received on ports that are blocked by ERPS.
xos0059671	On Summit X460 series switches with 750 W power supplies installed, log messages "Power usage data unknown" appear.
Summit X670-G2 Series Switches	
xos0059445	Link flaps occur when stacks are firmed with 3 m/5 m QSFP+ passive copper cables.
Summit X770 Series Switches	
xos0055746	Stacking port link flap occurs on Summit X770 series switches when using 3-meter QSFP+ cables.

Resolved Issues in ExtremeXOS 15.5.3-Patch1-6

The following issues were resolved in ExtremeXOS 15.5.3-patch1-6. ExtremeXOS 15.5.3-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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.4.6-Patch1-14, ExtremeXOS 15.4.2.8 and ExtremeXOS 15.5.3.4. For information about those fixes, see the release notes for the specific release.

Table 11: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3-Patch1-6

ID Number	Description
General	
xos0055795	The <code>show fdb stats vlan</code> command output does not show the number of MAC addresses learned over VPLS pseudowires.
xos0059077	Getting error after executing the <code>upload log</code> command multiple times.
xos0059789	Dos-Protect ACL is not cleared after continuous DOS attack occurs.
xos0059945	Memory corruption occurs rarely due to packet buffer overrun while sending/receiving control packets between slots.
xos0056913	OSPFv3 process ends unexpectedly when link goes down on set of ports in the switch.
xos0058391	Switches allow installing ACL policy with meter even though the corresponding meter is not yet created.
xos0059757	With per-VLAN IGMP snoop filter, backup VRRP does not forward hellos messages of the same VRIDs. This occurs only when the VRRP backup is master for a different VLAN using the same VRID and the IPMC traffic is being slow-path forwarded.
xos0060100	Kernel oops occurs due to memory corruption caused by slow-path forwarded traffic.
E4G-200 Cell Site Routers	
xos0058239	In E4G-200 cell site routers, power supply status displays incorrect value in the output of the <code>show power</code> command.
Summit X440 Series Switches	
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.

Table 11: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3-Patch1-6 (Continued)

ID Number	Description
BlackDiamond 8800 Series Switches	
xos0059648	Static ARP entries are not properly synced with new Master Switch Fabric Module after failover.

Resolved Issues in ExtremeXOS 15.5.3-Patch1-5

The following issues were resolved in ExtremeXOS 15.5.3-patch1-5. ExtremeXOS 15.5.3-patch1-5 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-8, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.4.6-Patch1-10, ExtremeXOS 15.4.2.8 and ExtremeXOS 15.5.3.4. For information about those fixes, see the release notes for the specific release.

Table 12: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3-Patch1-5

ID Number	Description
General	
xos0054949	With VLAN aggregation, local multi-cast packets received on one sub-VLAN are flooded to other sub-VLANs.
xos0056323	On failed stack nodes, running any show commands produces an error.
xos0056857	Configuring invalid addresses for RADIUS netlogin server causes CLI to stop responding.
xos0059030	ARP entries incorrectly point to ISC port after MLAG peer is rebooted.
xos0059243	The process <code>exsh</code> ends unexpectedly after executing a show command with a port list followed by invalid letters (for example, <code>show port 1:1,1:2ab</code>), and then pressing TAB .
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.
xos0059330	With dual master switch fabric module (MSM) installed, clear-flow ACL intermittently fails.
xos0059581	Rtmgr process ends unexpectedly when OSPF external routes are deleted from the route table.
xos0059584	OSPFv3 intra-area routes are rarely not added to routing tables due to timing issue.
xos0059733	LSP load sharing does not occur on Summit X460-G2 and BlackDiamond X8-100G4X switches.

Table 12: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3-Patch1-5 (Continued)

ID Number	Description
xos0056342	Misleading power supply unit (PSU) traps are sent when PSUs are inserted or powered on/off.
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.
BlackDiamond X8 Switches	
xos0057352	Kernel crash occurs when there is a Layer 2 loop in the network.
xos0059343	The process snmpMaster might end unexpectedly during upgrade from ExtremeXOS 15.3 to 15.5 for some SNMP community names.
xos0059603	Management Modules fail to complete booting when there is a failed Fabric Module in the chassis.
Summit Family Switches	
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.
Summit X430 Series Switches	
xos0059524	Link status is incorrect when auto-polarity setting is off.

Resolved Issues in ExtremeXOS 15.5.3-Patch1-2

The following issues were resolved in ExtremeXOS 15.5.3-patch1-2. ExtremeXOS 15.5.3-patch1-2 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-8, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-44, ExtremeXOS 15.3.4.6-Patch1-8, ExtremeXOS 15.4.2.8, and ExtremeXOS 15.5.3.4. For information about those fixes, see the release notes for the specific release.

Table 13: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3-Patch1-2

ID Number	Description
General	
xos0056994	Unable to add EAPS shared ports to VLANs even after disassociating them from VPLS domains.
xos0057235	Switch crashes when the command <code>restart process mpls</code> is executed repeatedly within a short time interval.
xos0057435	Packets are dropped when learning is disabled in a VLAN when its associated ports are configured with limit learning in another VLAN.
xos0057785	STP domain tag is removed when all ports are deleted from STP auto-bind enabled-VLANs.
xos0058849	Jumbo frames are fragmented on LAG ports after re-configuring port sharing even though jumbo frame is enabled on those ports.
xos0058873	FDB entries are learned incorrectly on VMANs in MLAG peers when the MLAG ports are CEP ports for multiple VMANs with different CVIDs.
xos0058968	Error log "Function Pointer Database is not fully initialized" appears during bootup on non-Summit platforms.
xos0059002	Checkpoint errors occur during execution of STP debug command if switch contains many STP-enabled VLANs.
xos0059037	Pre-emphasis <code>show</code> command displays incorrect values for non-Summit X460 series switches' slots in mixed stacks.
xos0055814	The process <code>mcmgr</code> ends unexpectedly when processing corrupted MLDv2 report packets.
xos0057808	Packet loss occurs due to hardware convergence during RSVP fallback scenario.
xos0058683	RIP packets are dropped when another VLAN has a secondary IP address configured.
xos0058717	The message "Warn:MPLS.RSVPTE.InternalProb" appears after disabling ports in RSVP secondary path.

Table 13: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3-Patch1-2 (Continued)

ID Number	Description
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.
xos0058880	Packets are not switched to primary path after recovering from path failure in MPLS RSVP-TE.
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.
xos0059222	SFLOW-sampled packets are flooded out of VLANs when these same packets are software learned.
xos0059305	OSPF consumes a large amount of memory when a large number of Link State Acknowledgment packets are queued up for transmission.
xos0057643	On Summit X460-G2, X770, and BlackDiamond switches, slow learning rates occur for scaled multi-cast traffic in l2-and-l3-and-ipmc mode.
BlackDiamond X8 Switches	
xos0058375	ACLs to match VALN-ID, CVID parameters do not work for slow path forwarded packets.
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.
SummitStack	
xos0057767	Static FDB associated with a VPLS service VLAN is not programmed in hardware after reboot when "disable learning" is configured.
xos0057230	During failovers in SummitStacks, the backup/standby nodes go to failed state occasionally and get rebooted.
Summit X440 Series Switches	
xos0058300	Packets are dropped on combo ports when the preferred medium is configured as copper force.
xos0058547	In Summit X440-24t switches, the maximum hotspot temperature should be changed to 70 C.
Summit X460 Series Switches	
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.
Summit X670 Series Switches	
xos0059128	In Summit X670 series switch, all LEDs are blinking at a faster rate.

Resolved Issues in ExtremeXOS 15.5.3

The following issues were resolved in ExtremeXOS 15.5.3. ExtremeXOS 15.5.3 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-Patch1-8, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-44, ExtremeXOS 15.3.4.6-Patch1-5, ExtremeXOS 15.4.2.8, and ExtremeXOS 15.5.2.9-Patch1-5. For information about those fixes, see the release notes for the specific release.

Table 14: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3

ID Number	Description
General	
xos0056423	The command <code>show access-list meter port</code> does not display the meters applied on the port via policy.
xos0057281	Retries and new informs are not sent to non-responding inform receivers.
xos0057328	ACL rule to match IPv6 packets with arbitrary mask does not work correctly.
xos0057336	Time interval between inform and retried inform is greater than the configured value when continuous informs are triggered.
xos0057964	ACL process ends unexpectedly when accessing AVB related dynamic rule from ExtremeXOS ScreenPlay.
xos0058056	OSPF-opaque, LSA-related configurations do not appear in output of the <code>show configuration</code> command.
xos0058397	Unknown unicast traffic via VPLS is duplicated after back-to-back LSP failover.
xos0058464	In ERPS rings, blocking the control channel by deleting the ports from the control VLAN causes a short loop in the ring.
xos0058578	Sync Receipt Timeout Interval appears incorrectly in the output of the <code>show network-clock gptp ports</code> command.
xos0058603	PTP follow up does not happen correctly when correction field is greater than 32 bits.
xos0058695	Process <code>emsServer</code> ends unexpectedly with signal 6 when multiple VRRP messages are logged.
xos0054199	Ingress traffic stalls on port when switches receive continuous 802.3x pause frames on egress ports for that traffic stream.
xos0055347	Temperature reported in log messages is different than the output of the <code>show temperature</code> command.
xos0055870	Output of <code>show configuration</code> command should show deletion of default SNMP communities.
xos0057179	ESRP feature is not enabled immediately after the installing an Advance Edge license.

Table 14: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3 (Continued)

ID Number	Description
xos0057672	The process rtmgr ends unexpectedly with signal 11 after disabling a GRE tunnel.
xos0057647	Packets are forwarded to CPU after deleting the VLAN with disable learning.
Summit Family Switches	
xos0058001	Multi-cast traffic is dropped in the last hop router even though egress ports are correctly associated with multi-cast group.
xos0058050	gPTP: Noise introduced to audio.
xos0058306	gPTP: Propagation delay measurement shows negative value.
xos0058537	Switches become unresponsive and drop traffic when they have a high number of traffic streams and AVB enabled ports.
SummitStack	
xos0057255	Multi-cast entries are not programmed in hardware intermittently for certain multi-cast groups in stacking setup.
xos0057562	PoE initialization fails on certain SummitStack nodes with SSH enabled.
xos0058218	Need commands to tune debounce timer for stacking port.
xos0058589	SummitStack reboots due to temperature out of range messages.
xos0055846	VPLS traffic is not forwarded to some service VLANs in SummitStack.
xos0056129	In SummitStacks with a high number of VLANs and protocols enabled, the control traffic might get looped back to stacking port after failover.
BlackDiamond X8 Switches	
xos0055433	The process tDiag occasionally ends unexpectedly when the command show debug system-dump MM B is executed from the Management Module.
xos0055802	The command debug ha1 show congestion does not work on switches with BDXB-100G4X modules.
xos0058568	Some front panel ports cannot be enabled after rebooting the I/O module.
Summit X440 Series Switches	
xos0058068	Summit X440-24tDC switches reporting maximum temperature limit 60°C under normal condition.s
xos0058301	In Summit X440 series switches, error message "mounting /dev/hda4 on /data failed" appears during bootup.
Summit X460 Series Switches	
xos0058043	MSRP: Packets are dropped when bandwidth is increased to 611M.
xos0058217	Need commands to tune pre-emphasis settings for stacking ports.

Table 14: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.3 (Continued)

ID Number	Description
Summit X770 Series Switches	
xos0057772	IPv6 unicast traffic destined to local IPv6 host that has a switch interface mask greater than 64-bits is not forwarded after clearing FDB.
xos0058540	Packet loss occurs for more than 10 seconds during MLAG fallback.

Resolved Issues in ExtremeXOS 15.5.2-Patch1-5

The following issues were resolved in ExtremeXOS 15.5.2-Patch1-5. ExtremeXOS 15.5.2-Patch1-5 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-8, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-44, ExtremeXOS 15.3.4.6, ExtremeXOS 15.4.2.8, and ExtremeXOS 15.5.1.6-Patch1-1. For information about those fixes, see the release notes for the specific release.

Table 15: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2-Patch1-5

ID Number	Description
General	
xos0055958	When using CLI scripting error occurs while executing <code>set var seconds \$TCL(clock seconds)</code> command.
xos0047180	EMS process ends unexpectedly when high amount of logs were targeted to console
xos0056004	Traffic loss occurs when CEP VPLS and untagged VMAN VPLS is configured on same port.
xos0056339	FDB learning does not occur after deleting/adding subscriber VLAN from PVLAN.
xos0056977	Client authenticated using netlogin dot1x gets incorrect IP address due to delay in moving port to success VLAN.
xos0056995	IPv6 traffic for routes with mask length greater than 64 are not forwarded after clearing FDB.
xos0057013	IPv6 traffic for routes with mask lengths greater than 64 characters is slowpath forwarded if switch has tunnels configured or when destination MAC address of the IPv6 packet is a virtual MAC address.
xos0057043	With MLAG and PVLAN configured, after disabling MLAG port, IP ARP entries continue to point to disabled MLAG port, instead of ISC port, causing traffic loss.

Table 15: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2-Patch1-5 (Continued)

ID Number	Description
xos0057088	Cannot log on using SSH with a 32-character or greater password. After eight logon attempts, no more SSH connections are permitted.
xos0057481	Fragment Reassembly Time Exceeded messages not sent for IPv6.
xos0057045	Unable to match ARP packets using ACL match criteria "arp-sender-address" and "arp-target-address" in c-series IO cards.
xos0057384	Hardware learning is not enabled on the sharing member ports on deleting a MLAG peer
xos0057771	Packet loss seen during RSVP fallback scenario.
xos0057804	If same port is used for CEP-VPLS + Untagged-VMAN(L2) , then all CEP-VPLS traffic goes via Untagged-VMAN(L2).
xos0057492	GTP: Latency timestamp for VIM XGM2S-2sf is invalid.
Summit X430 Series Switches	
xos0058013	Summit X430 series switches should support identity management and HealthLAG features.
xos0057677	Summit X430-24p switches total PoE budget should be 370W instead of 380W.
xos0057837	MSRP stream propagation fails after link flap.
xos0057691	AVB does not work on the Summit X430-8p and X430-24p switches.
BlackDiamond 8800 Series Switches	
xos0057354	Kernel gets stuck after issuing the command <code>clear fdb</code> , followed by MSM failover when switch has highly scaled FDB and ARP entries.
xos0057422	Need to limit the kernel error log messages when packets are dropped.
xos0057561	Enabling mirroring on BlackDiamond 8800 series switches with MSM-48c cause VRRP/LACP flapping.
BlackDiamond X8 Series Switches	
xos0050424	Slot goes into failed state after running extended diagnostics.
Summit X460 Series Switches	
xos0057916	Default debounce timer should be set to zero on stacking ports.
xos0057024	Remote ports connected to 10G ports from XGM3 module experience link flap, and stacking link formed using ports from XGM3 module experience link flap.
Summit X480 Series Switches	
xos0054362	After learning around 10K FBB entries, MPLS FDBs are not cleared from hardware if the <code>clear fdb</code> command is executed.
xos0057454	FDB entries are not aged out after two continuous MAC address moves.

Table 15: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2-Patch1-5 (Continued)

ID Number	Description
xos0057602	Ping across VPLS domain does not succeed if the ingress port in CPE switch is part of both service VLAN (tagged) and service VMAN (untagged).
xos0057752	On Summit X480 series switches when untagged VLAN packets are received on tagged VMAN ports, which in turn receive from untagged VMAN ports bursts of more than 500 packets in line rate, some FDB entries are not learned in hardware.
Summit X670 Series Switches	
xos0055074	On Summit X670v-48x series switches, links go down after rebooting if the ports are configured with auto negotiation off and the speed is 1G on both peers when using a mini Gigabit interface converter.
Summit X770 Series Switches	
xos0057290	Few IPv4 unicast/multicast streams get slow path forwarded when switch has more than 40K IP ARP entries.

Resolved Issues in ExtremeXOS 15.5.2-Patch1-1

The following issues were resolved in ExtremeXOS 15.5.2-Patch1-1. ExtremeXOS 15.5.2-Patch1-1 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-8, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-41, ExtremeXOS 15.3.4.6, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.1.6-Patch1-1, and ExtremeXOS 15.5.2.9. For information about those fixes, see the release notes for the specific release..

Table 16: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2-Patch1-1

ID Number	Description
General	
xos0056339	FDB entries are not learned after deleting/adding subscriber VLAN from PVLAN.
xos0057043	With MLAG and PVLAN configured, after disabling MLAG port, IP ARP entries continue to point to disabled MLAG port, instead of ISC port, causing traffic loss.
xos0057088	Cannot log on using SSH with a 32-character or greater password. After eight logon attempts, no more SSH connections are permitted.
xos0057481	Fragment Reassembly Time Exceeded messages not sent for IPv6.

Table 16: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2-Patch1-1 (Continued)

ID Number	Description
xos0056004	Traffic loss occurs when CEP VPLS and untagged VMAN VPLS are configured on the same port.
xos0056995	IPv6 traffic for routes with mask lengths greater than 64 are not forwarded after clearing FDB.
BlackDiamond 8800 Series Switches	
xos0057045	Unable to match ARP packets using ACL match criteria "arp-sender-address" and "arp-target-address" in c-series I/O modules.
Summit Series Switches	
xos0057492	GTP: Latency timestamp for VIM XGM2S-2sf is invalid.
Summit X460 Series Switches	
xos0057024	Remote ports connected to 10G ports from XGM3 module experience link flap. Also, stacking link formed using ports from XGM3 module experience link flap.
Summit X480 Series Switches	
xos0054362	After learning around 10,000 FDB entries, MPLS FDB entries are not cleared from hardware after executing <code>clear fdb</code> .
xos0057454	FDB entries are not aged out after two contiguous MAC moves
xos0057602	Ping across VPLS domain does not succeed if the ingress port in CPE switch is part of both service VLAN (tagged) and service VMAN (untagged).
Summit X770 Series Switches	
xos0057290	Some IPv4 unicast/multicast streams are slow path forwarded when switches have more than 40,000 IP ARP entries.

Resolved Issues in ExtremeXOS 15.5.2

The following issues were resolved in ExtremeXOS 15.5.2. ExtremeXOS 15.5.2 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, and ExtremeXOS 15.5.1. For information about those fixes, see the release notes for the specific release.

Table 17: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2

ID Number	Description
General	
xos0056109	Traffic loss occurs for 15–30 seconds, when PIM non-designated router changes to designated router.
xos0055662	Delays in triggering a cyclic UPM profile, cause the next FireTime to be calculated using the delayed current time.
xos0055486	Traffic that doesn't match the protocol filter configured with L2PT profile is dropped in the provider edge. The traffic should tunnel through the VMAN and reach the other end.
xos0055482	L2PT traffic in the provider network is dropped in transit node if L2PT is configured in it.
xos0055526	The command "set var emp3 \$tcl)" produces the error "Error: Unsupported TCL Command """, but it should show "% Missing close parenthesis at '^' marker.".
xos0055477	Switches don't allow configuring L2PT protocol filter with dest-mac, etype, etc.
xos0055476	DHCP decline packets are dropped if the client address field within the DHCP decline packet is 0.0.0.0.
xos0055234	Enhance the output of the command <code>debug ha1 show sys-health-check</code> to include I/O card memory information and Async queue counters.
xos0055194	With option-82 enabled, DHCP request packets contain incomplete circuit-id information when the string exceeds 15 characters.
xos0055055	CliMaster process ends unexpectedly when the telnet session disconnects from the switch while it is printing a lengthy debug output.
xos0056364	In BlackDiamond X8 and 8800 series switches, with distributed ARP turned on, L3 traffic doesn't egress after a link failover, even though ARP entries have been learned and they are pointing to the correct ports. The traffic starts flowing after executing the command <code>clear iparp</code> .
xos0055849	By enabling ACL log filters, Admin user can see the dynamic ACL binding/unbinding information of Lawful Intercept user.

Table 17: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2 (Continued)

ID Number	Description
xos0055257	After multiple reboots, VRRP instances remain in INIT state when ELSM and VRRP are running over the same single port.
xos0056020	L3 traffic ingressing on the ISC link from an MLAG peer is not egressing the appropriate egress MLAG port after L3 forwarding. Occurs only with jumbo frames enabled on a few of the ports.
xos0052365	Unconfiguring L3VPN and then configuring it again results in BGP failing to advertise the routes.
xos0053543	Switch responds to a GARP attack for an IP address that is configured as a static IP ARP entry with the switch MAC address, but the switch should respond with the MAC address listed in the static IPARP entry.
xos0054120	Enabling IPv6 smart relay at VR level does not forward the relay-forward solicit from having multiple BOOTPV6 relay enabled agents.
xos0055585	Multicast traffic can take up to 60 seconds to recover when an ingress port on a first-hop router (FHR) is disabled.
xos0055611	PIM does not failover to alternate source received from the MSDP peers when the primary source fails.
xos0055685	UPM process ends unexpectedly after executing the command <code>show config upm</code> repeatedly for a prolonged period of time.
xos0055754	Login authentication event is generated for Lawful Intercept user if the user logs on using SSH.
xos0055783	Switches drop packets and display "Invalid MAC Binding" error when you remove an existing client and connect a different client with the same IP address.
xos0055810	OSPF process remains in STOPPED state after terminating the OSPF process and deleting the OSPF protocol from the user VR.
BlackDiamond 8800 Series Switches	
xos0055744	BlackDiamond 8800 series switches allow you to write to the <code>/usr/local</code> directory.
xos0052294	A port in the ISC VLAN is incorrectly allowed to be configured as an MLAG port.
xos0055588, xos0057196	On BlackDiamond 8800 c-series I/O modules, hardware programming of a large batch of IPMC cache entries takes several minutes.
xos0052349	Process <code>rtmgr</code> ends unexpectedly with signal 11 when disabling an I/O slot in a BGP configuration.
xos0056329	Process <code>vlan</code> ends unexpectedly with signal 11 when configuring ELRP-client for single (non-load sharing port); works for load sharing port.
BlackDiamond X8 Series Switches	
xos0055803	Backup management cards do not enter "In Sync" state if port-specific tags are configured.
xos0054748	When BlackDiamond X8 QSFP+ LR4 optic is disabled, the link goes to disabled state as expected. When it is enabled, the link remains in ready state instead of going back to active state.

Table 17: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2 (Continued)

ID Number	Description
xos0053879	The error "<Error:HAL.MPLS.Error> MM-A: ILM instance 872677376 not found" appears after issuing the command <code>clear counters mpls rsvp-te</code> .
PD4-4484429777, xos0055282	Strict-Priority (SP) and Weighted-Deficit-Round-Robin (WDRR) scheduling are not supported on the BlackDiamond XB-100G4X in this release. Weighted-Round-Robin (WRR) scheduling is supported on the BlackDiamond X8-100G4X switches in this release.
xos0055644	Add Extreme Network's serial number to the "easily readable" section of CFP2 optic command <code>debug hal show optic-info slot <slot> port <port></code> .
xos0053683	DHCP packets from sub-VLANs are not egressing through the super VLAN ports
xos0055730	On BlackDiamond X8 switches with BDX8-100G4X cards, MAC addresses are no longer learned on EAPS ring ports after ports are disabled and enabled multiple times.
xos0055733	VPLS: Multicast traffic is dropped on BDX8-100G4X cards when the service port is part of a VLAN with IGMP snooping enabled.
Summit Series Switches	
xos0057026	PoE initialization fails in Summit X440 and X460 series switches when upgrading from ExtremeXOS 15.2 to 15.5.
xos0055459	Maximum number of private VLANs in an L2-only environment limit not attainable.
xos0055623	MLAG ports configured with LACP go down after performing a stack failover.
xos0054086	VLAN aggregation configurations are not removed after rebooting the switch when configuring a dynamic VLAN as the sub-VLAN.
xos0053227	VLAN statistics are not working after a VLAN with base PSTAG is removed and added back to ports.
xos0056434	MLAG ports are not added to LACP aggregator in all the MLAG peers.
SummitStack	
xos0055724	CDP PDUs are encapsulated and sent, even when CDP profiles are not binded to any access ports on either side of the L2PT tunnel.
xos0053970	When rebooting Summit X670-X670-X670-X770 stack, process ntp ends unexpectedly with signal 11. "Process ntpd pid 4970 died with signal 11 Code: 2ac14e18 00e43821 addu a3,a3,a0 2ac14e1c 01003021 addu a2,t0,zero 2ac14e20 24840008 addiu a0,a0,8 2ac14e24 <ac85fff8>sw a1,-8(a0) 2ac14e28 1487fffd bne a0,a3,0x2ac14e20 2ac14e2c ac85fffc sw a1,-4(a0) 2ac14e30 30c80004 andi t0,a2,0x4 2ac14e34 11000003 beq t0,zero,0x2ac14e44 2ac14e38 00c83023 subu a2,a2,t0"

Table 17: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2 (Continued)

ID Number	Description
xos0055980	In Summit-stack, traffic passing through a port from backup node stops after adding that port to another VLAN with port specific tag configuration.
Summit X460 Series Switches	
xos0052683	Stacking does not come up when trying to use one port in alternate and another port in native mode for stacking.
xos0055580	After unconfiguring the switch and loading the default.xsf, script, sending HTTP SOAP/XML produces the error: "Process thttpd exceeded pre-configured or default crash rate" and the switch reboots.
xos0057040	Due to an incorrect latency time stamp adjustment, with gPTP enabled on ports corresponding to ports slot 2 ports, a large negative propagation delay value appears in the output of the <code>show network-clock gptp ports</code> command.
xos0056033	CFM fault is not detected by hardware down MEP when CCM transmission is disabled at its RMEP.
xos0055416	On Summit X460 series switches, after enabling jumbo frames on 10G ports on XGM modules, doing either a save and reboot or changing port speed settings removes the jumbo frames setting.
xos0055748	VPLS: On Summit X460 series switches, the command <code>clear fdb</code> does not clear the MAC addresses learned on service VLAN (local and pseudowire MAC addresses). The command <code>clear fdb <service vlan name></code> does clear the MAC addresses.
Summit X480 Series Switches	
xos0056976	On Summit X480 series switches with 40G4X VIM cards, FDB entries are not in sync across units when FDB entries are learned and aged out frequently.
Summit X670 Series Switches	
xos0055369	On Summit X670 series switches with BlackDiamond X8 series switches as neighbors the L2PT received tunneled packets are not being tunneled towards customer edge. This defect is in VMANs only.
xos0054490	ACL rule with match condition <code>igmp-msg-type</code> does not work when packet contains <code>ip-option</code> .
xos0056125	MLAG ports move to R state after multiple executions of the command <code>restart process vsm</code> .
xos0056115	Some VRRP VLANs undergo state change when ports belonging to VRRP VLANs that are VLAN/route tracked are disabled/enabled.
xos0048730	After MLAG is established, enabling, and then disabling MLAG ports produces error messages: <pre><Crit:VSM.ParmInv> : Argument Ingress Port Instance has an invalid value 16777217 07/30/2012 22:22:37.27 <Erro:Kern.Error> : exvlan: handleVsmKernelRequest:7311: handleVsmKernelRequest Invalid Ingress port: 1000001 got</pre>

Table 17: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2 (Continued)

ID Number	Description
Summit X770 Series Switches	
xos0055632	VPLS: On Summit X770 series switches, traffic is dropped after configuring LAG on service port (untagged VLAN port). The issue is with high port numbers (port 101, 102) on Summit X770 series switches.
xos0055095	Summit X770 series switches' port up/down activity can occasionally result in erroneous entries in the L3 tunnel hardware table, causing L3 slowpath traffic. This does not affect Layer-2 configurations, only Layer-3 configurations.
ACLs	
xos0056054	Disabling user-created mirror instance with ACL filter is not working when Default mirror instance is enabled with ACL filter.
ESRP	
xos0053647	Rebooting after a failover from master to backup works fine, but after performing another failover, the ports in some ESRP member VLANs in slave are unblocked, but ports in ESRP master VLAN are blocked. This creates a situation where ports in ESRP member VLAN are unblocked in both ESRP master and slave, causing traffic loops.
xos0056018	In ERPS protection state, FDB flush occurs every 5 seconds without any topology change.
xos0055959	ERPS ring does not move to protection state after disabling a ring-port.
MPLS	
xos0055671	When untagged service port from service VLAN of VPLS is removed, the following errors occur: <pre><Error:Kern.MPLS.Error> (ems_main.c:447) bcm_custom_extr_vfp_action_mpls_inport_set(remove) failed for unit = 0 vpn = 0x1, port = 0x800000f vp = 0x18000002, rv = -13 (Invalid identifier)</pre>
xos0056114	Packets are not passing over VPLS when PS tag is configured in the service VLAN.
xos0055857	FDB entries are not learned after changing VMAN tag when the VMAN is associated with VPLS.
xos0055914	LSP-specific transmit counters incorrectly display a zero value in the output of the command <code>show mpls statistics l2vpn</code> after failover to secondary RSVP path.
xos0055991	Traffic is dropped after disabling learning on VPLS service VLAN ports with L2VPN sharing enabled.
OpenFlow	
xos0055315	When sending 300 packets to hit the ofDefault_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.

Table 17: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.2 (Continued)

ID Number	Description
Security	
xos0044258	Uninstalling SSH image fails with error: <code>mount: mounting /dev/mtdblock.</code>
VLANs	
xos0056253	FDB entries are learned despite learning being disabled on ports when port specific tag added to the VLAN.
VRRP	
xos0055515	VRRP backup node is moving to master state even despite connectivity between master and backup nodes.

Resolved Issues in ExtremeXOS 15.5.1

The following issues were resolved in ExtremeXOS 15.5.1. ExtremeXOS 15.5.1 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, and ExtremeXOS 15.4.1. For information about those fixes, see the release notes for the specific release.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1

ID Number	Description
General	
xos0055488	OpenSSL versions 1.0.1 through 1.0.1f contain a flaw in its implementation of the TLS/DTLS heartbeat functionality (RFC6520).
xos0055513	With CDP enabled, EDP process ends unexpectedly if malformed CDP packets are received on a port.
PD4-4445539969	CDP neighborhood is not detected when TLV app is enabled on peer switches.
PD4-4365231211	switches display only one ELRP log message even with two loop conditions.
PD4-4485052528	ERPS status stays in IDLE state even though CFM detects port link down event.
PD4-4269336974	In SummitStacks, VRRP MAC is not checkpointed to other slots after those slots are rebooted.
PD4-4448380078	When jumbo frames are enabled on some ports, the LACP error message "pibL3MTUExceededInstallFilterLag" appears after disable sharing on one port, and then enable sharing on other ports.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
PD4-2983188951	With approximately five LACP LAGs with active member ports in each LAG and sharing enabled on multiple ports, executing the command <code>restart ports all</code> produces the following warning message for each LAG: <Warn:LACP.UnackTimerExp> Unack timer expired for LAG 25 262148 Instance 4 index 4
PD4-4341189061	When STP port priority is set to zero, the <code>show configuration stp</code> output does not reflect the configured priority value of zero.
PD4-4195557040	STP ports from the STP domain are removed when you remove untagged ports from a VLAN which is auto-bound to an STP domain.
PD4-4371661882	OID <code>radiusAuthServerTable</code> does not match the standard OID, and the <code>snmpwalk</code> on <code>radiusAuthServerTable</code> is not complete.
PD4-1237324487	Error messages should appear when the <code>dhcp-address-range</code> is invalid (0.0.0.0 - 10.1.1.1), but this does not occur.
PD4-4272356613	BootpRelay forwarding is not working when client sends the Bootp request.
PD4-4315137254	Configuring <code>dhcp-address-range</code> for management VLANs should not be allowed.
PD4-4325909054	Switches stop responding after deleting user-created virtual-routers when SNTF is enabled for that virtual router.
PD4-4366873466	When switches are synchronized with the NTP server for four or more days, NTP process ends unexpectedly with signal 6 when disabling NTP-enabled VLANs.
PD4-4217943426, PD4-4240054737	LLDP detect and undetect events are not properly triggered when <code>netlogin</code> , <code>identity management</code> and <code>UPM</code> are running on the same port.
PD4-3820903431	With multiple MIPs associated with an Up-MEP, CCM received on one MIP is forwarded on other MIPs of that Up-MEP, producing a loop.
PD4-4476473377	The length in the CCM MA header is different for MA Primary VID. This creates a problem if an S/k/TOR/ExtremeXOS CFM operation uses a MA VID name. When this occurs, the following type of error appears: 03/05/2014 10:27:57.46 <Noti:CFM.RxPktInv> Received invalid cfm on port 14, VLAN "Video2101"; MA Identifier does not match configured MEP's Identifier
PD4-3674095924	CFM: <code>AvgFlr</code> is always greater than <code>maxFlr(100000)</code> when doing a SNMP walk on MIB objects: <code>mefSoamLmCurrentAvailStatsBackwardAvgFlr</code> <code>,mefSoamLmCurrentStatsForwardAvgFlr,mefSoamLmHistoryAvailStatsForwardAvgFlr,</code> <code>mefSoamLmHistoryStatsForwardAvgFlr.</code>
PD4-4429113125	The help text next to <code>internal-memory</code> and <code>memory card</code> start with a lower case letter while all the other help text starts with capital letters.
PD4-4483766120	Quitting or exiting the BCM shell with at least one uppercase letter causes the <code>bcm.shell</code> thread to end.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
PD4-3387133462	The command <code>show cli commands</code> displays an infinite list of repeated commands.
PD4-4448380163	The command <code>upload debug</code> produces an error when the same policy is applied for SNMP and Telnet as for access-profiles.
PD4-4490946011	With VPLS enabled, ARP packets are dropped when ingress and egress ports are present in different units.
PD4-4227389365	<code>extremeSlotModuleInsertedType</code> enumeration not defined for Summit x450-24t
PD4-4472590563	The help description for the command <code>show lacp member-port <port></code> should state "port" instead of "port-list".
BlackDiamond 8800 Series Switches	
xos0055389	In BlackDiamond 8800 series switches, optimize I/O card memory usage to increase free memory after bootup.
PD4-4368267946	With master switch management module B as master, EPM process ends unexpectedly when performing an SNMP query to check download image status.
PD4-4235009619	With two master switch fabric modules having different versions of ExtremeXOS, error messages like the following appear, but can be ignored until an upgrade to bring both modules up to the same version of ExtremeXOS is performed: "<Info:HAL.Port.RateLimit> MSM-B: Flood Rate Limiting activated on Port 7:193"
PD4-4370250981	Ports stop forwarding traffic when egress mirroring is configured on some other port.
PD4-4170411471	Mirroring hardware filter counter is not recalculating after deleting ports from VLANs.
PD4-4332594144	LAG ports are not added to the aggregation group after a master switch fabric module failover followed by a port restart.
PD4-4268336458	On BlackDiamond 8800 and BlackDiamond X8 series switches, slowpath forwarding of IPv4 packets can occur if there is contention for IP hardware table resources.
PD4-4215712519	Executing the command <code>show iproute mpls</code> along with disabling/enabling MPLS causes process <code>rtmgr</code> to end unexpectedly with signal 11.
PD4-4461173469	System ends unexpectedly with process <code>rtmgr pid 1480</code> with signal 11 error while detecting tunnels in the IPv6 services in E4G-200-12x cell site routers.
BlackDiamond X8 Series Switches	
PD4-4176106657	When running diagnostics on the I/O modules, fabric module modules, and on the backup management modules simultaneously, the backup management module reports all I/O modules as remaining in the "RT sync" state.
PD4-3282517920	CFM_MIB - <code>snmp get of dot1agCfmMaNetName</code> failed for one display parameter. Expected result is "Hex-STRING: 00 00 11 FF FF FF FF" but received "Hex-STRING: 22 11 00 FF FF FF FF" instead.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
PD4-4176234031	VPLS; When load sharing (LAG) ports are part of both CNP (untagged VMAN port)s and tagged VLAN ports, the tagged VLAN traffic is flowing through untagged VMAN ports.
PD4-4332910254	The error message "Could not derive slot/port from modid/port" appears after doing master switch fabric module failover with VPLS traffic.
PD4-4462714410	With bi-directional traffic via VPLS/VPWS session., statistics appear correctly while executing <code>show mpls statistics l2vpn</code> command. After clearing counters, the show command produces error messages.
PD4-4459526211	BlackDiamond X8 series switches fail to learn source MAC addresses on pseudowires.
PD4-3956394996	Get bulk on last OID on MIB tree is supposed to respond with a single varbind containing endOfMibView exception.
PD4-4238911703	Slot goes to failed state after configuring TRILL.
E4G Cell Site Routers	
PD4-4459291184	CFM fails to recover from a fault state after the control VLAN is disabled/enabled in a transit switch. As a consequence, ERPS stays in protection state.
PD4-4184442680	BFD session does not come up when ExtremeXOS software image is downgraded from 15.4 version to 15.1.
PD4-4210700243	Cannot create LAG groups on ports where ERPS was configured with CFM.
PD4-4388433741	VRRP trackip route feature fails to detect all invalid routes when the routes are tracked across multiple VRRP instances. Only one VRRP instance is in INIT state, while the remaining VRRP instances are in MASTER state.
PD4-4237445510	After restarting all the nodes in SVT-MBH topology, A2 node [E4G-400] that is configured to act as both PTP master/slave fails to sync with 2.1.1.101 GM.
Summit Family Switches	
xos0054357	On Summit X250 and 450 series switches, MAC addresses learned on port 1 do not appear in <code>show fdb</code> command.
PD4-4172052424	MVRP propagation fails when disabling/enabling MVRP transit flag for a specific VLAN tag.
PD4-4444562981	MAC address learning limit with stop action is not working for pstag, as unlimited learning is configured for other pstag when both pstags are added on same port.
PD4-4524379524	MRP: data packets are dropped when two streams are configured and LV is received on one of the streams.
PD4-4236108493	Need to automatically adjust for ingress and egress time stamp latency.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
SummitStack	
xos0055481	<p>L2PT encapsulated packet counters in VPLS are not cleared after executing the following commands:</p> <pre>clear counters vpls vp1233 clear counters ports all clear counters</pre>
xos0055478	<p>On SummitStacks, the command <code>show vpls vp1233 l2pt</code> fails with and displays the following error message:</p> <pre>Error:can't read \$xmlData(reply.0.message.0.l2ptConfig.- 1.serviceIfType.0)": no such element in array</pre>
PD4-4445793058	<p>SNMP query on SummitStack for <code>extremeMemoryMonitorsystem</code> does not return memory usage details of backup node.</p>
PD4-4192894881	<p>In SummitStacks, temporary loops occur briefly in EAPS rings when processes end unexpectedly on backup/standby nodes or during slot failover.</p>
PD4-4309090083	<p>Control packets (VRRP, OSPF, or BGP, etc.) originating from Summit X670/X770 stack nodes and egressing on 40G ports on backup or standby nodes are prepended with an extra 4-byte VLAN header with VLAN-ID 0 after failover. Effectively, egressing packets have two VLAN tags with tag-id 0x8100. This causes ping failures, VRRP flaps, and OSPF neighbors in down state.</p>
PD4-4395902048	<p>Cannot downgrade a Summit 160/320G stack to 80G. The command does not show the option <code>v80</code>:</p> <pre>Slot-1 Stack.2 # configure stacking-support stack- ports all selection native <cr> Execute the command V160 Select the V160 stacking mode V320 Select the V320 stacking mode</pre>
PD4-4193985055	<p>The command <code>show access-list dynamic rule</code> does not show the mirror instance name for a non-permanent dynamic ACL with a mirror action modifier after executing failover twice in a SummitStack or with BlackDiamond switches.</p> <p>Also, traffic is mirrored even if after disabling or deleting the mirror instance:</p> <pre>Slot-1 Stack.6 # sh access-list dynamic rule "adminaccess" entry adminaccess { if match all { source-address 10.1.1.1/255.255.255.0 ; } then { mirror ; ==> no mirror instance } }</pre>
PD4-4200357397	<p>Process <code>thttpd</code> pid 1658 ends unexpectedly with signal 11 in SummitStacks and standalone switches when connecting with the web interface:</p> <pre>Process thttpd pid 1658 died with signal 11 Code: 2ac583ac 3c1c0003 lui gp,0x3 2ac583b0 279c7bf4 addiu gp,gp,31732 2ac583b4 0399e021 addu gp,gp,t9 2ac583b8 <8c830020>lw v1,32(a0) 2ac583bc 1060000c beq v1,zero,0x2ac583f0 2ac583c0 8f878030 lw a3,-32720(gp) 2ac583c4 8c830030 lw v1,48(a0)</pre>

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
PD4-4238911763	Summit X67, X670, X670, X770 stack is not sending out multi-cast data traffic that belongs to few VLANs on the LAG port when some combination of member and master ports in the LAG are enabled.
PD4-4318329872	In SummitStacks, when clearing IP multi-cast entries the following error message appears in the backup node: " <code><Warn:HAL.IPv4Adj.Warning> Slot-2: adj 0.0.0.0: # L3 hash table entries already 0.</code> "
PD4-4488867973	In four-node Summit X670, X670, X670, X770 stacks, with a LAG bundle comprising ports from X770 and X670 nodes, traffic belonging to some multi-cast destinations is dropped in the stack node, when ports from the Summit X770 node along with ports from the Summit X670 are enabled.
PD4-4286668270	Netlogin process ends unexpectedly when you log on through netlogin web-based authentication if session refresh is disabled.
Summit X430 Series Switches	
PD4-4247133750	For Summit x430-24p and 8p, when ports 3 and 4 are connected to TFM slave, expected LED behavior for this PoE port is slowly blinking amber. Instead, connecting PD to port 3 causes the LED of port 4 to light amber, and vice versa.
PD4-4247133742	For Summit X430-24P, connecting a device to port 15 causes the LED for port 16 to light, and vice versa. Also, removing cable from port 15 produces a link down log message for port16, and vice versa.
PD4-4179279354	The following error log occurred: <code><Erro:Kern.Error> ems: (ems_main.c:471) BUG: scheduling while atomic: fdb/1220/0x10000200</code>
PD4-4324838766	Unable to ping an IPv6 interface created on Summit X430 series switches.
PD4-4439482778	MVR CLI commands should be blocked for Summit X430 series switches, since this feature is not supported.
PD4-4323778971	The L2 IP multi-cast cache entries are not programmed in hardware, and therefore the multi-cast traffic is software-forwarded. This issue does not occur in the Summit X430-24t.
Summit X440 Series Switches	
PD4-3572812521	Summit X440-8t switches show the error "Fan module 1 is not present" even though it does not support fan.
PD4-4188456050	After configuring ACL with ARP-sender-address and target address, enabling mirror, and then disabling mirror, hal process ends unexpectedly with signal 11.
PD4-4300989281	With OSPF configured between two switches, with one acting as a DHCP server, and a third switch (without OSPF configured) acting as a DHCP client, rebooting either OSPF-configured switch causes a duplicate IP address to be assigned to two VLANs on the third (DHCP client) switch. This causes an OSPF failure on the third switch, which doesn't have OSPF configured on it.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
Summit X460 Series Switches	
PD4-4502209879	Summit X460-G2-48t switches stop responding and display the error message "Process tBcmxAsync pid 1702 died with signal 6" after saving, and then rebooting, with IPv6 forwarding enabled.
PD4-4280254806	ACL with match condition "snap-8192", matches multi-cast traffic as well.
PD4-4175991191	With protected VPLS with ESRP redundancy, process MPLS ends unexpectedly with signal 6 after executing the command <code>restart process mpls</code> .
Summit X670 Series Switches	
PD4-3897318945	When meters are configured for several egress ACL rules, the hardware might be mis-configured when ACL rules are forced to be moved in the TCAM. This causes meters to be applied to the wrong ACL rule.
PD4-4247635211	On Summit X670v-48x SummitStacks, kernel error message appears while enabling/disabling diffserv on all ports.
PD4-4163132141	Egress ACL CVID match criteria is not working properly on CNP (untagged VMAN ports) when CEP is also configured on the same ports. It is working on Ingress side.
PD4-4200778453	TRILL Rbridges are not forwarding traffic from the dynamic VLANs created by XNV.
PD4-3717316170	EAPS convergence time is greater than 110 milliseconds.
PD4-3906172771. PD4-4198495910	With correlated QoS profiles (QP2-QP6) created and mapped on the Rbridges, when the native traffic has priority value in the 802.1Q tag set to value other than 0 (between 1-5), the Rbridge is only replacing the 802.1Q tag in the inner tag, but not in the outer tag for the ingress unknown uni-cast or multi-cast traffic. It works correctly for known uni-cast traffic.
PD4-4198675163	Disabling LAG between two Rbridges and then re-enabling it twice with the same configuration causes traffic to fail to resume. The issue does not re-occur with VLAN pruning disabled.
Summit X770 Series Switches	
PD4-4301019255	Hybrid scheduling is not working when QoS Scheduler is SP.
PD4-4432525731	Summit X770 series switches are not producing warning messages when unsupported 450W power supplies are inserted.
PD4-4409964727	LED is not blinking amber when a fan is removed from slots 2, 3, or 4, and LEDs are green with no fan inserted in those modules.
PD4-4211343935	Two-node, high speed stack of Summit X770 is dismantled after partitioning 40G ports (all) in backup slot to 10G. NOTE: In standalone mode, partitioning all ports to 10G cause the switch to stop forwarding or sending slow-path packets (like stacking or EDP) to the CPU in this state.
PD4-4488867731	Summit X670 stacks fail to egress VPLS encapsulated traffic when Summit X770 ports are part of the egress LAG.
PD4-4200778353	TRILL data traffic is not de-capsulated using LAG links on network ports.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
PD4-4200778321	When TRILL network ports are configured with L2 static LAGs, traffic is not evenly distributed over the LAG ports.
PD4-4211303616	Soft watchdog expires when enabling an EAPS ring port after disabling this port.
PD4-4390906721	Non-TCP any configuration overwrites non-TCP red drop parameters in hardware table. As a result, non-TCP any color with non-TCP red is not working correctly.
PD4-4249042211	Traffic received on LAG member port (service VLAN) is not forwarded to VPLS peers, and FDB learning stops.
PD4-4200778383	In Summit X770-320G stacks, "Process hal pid 1439 died with signal 11" error appearing when enabling MPLS license in standby slot. License updates correctly in master and backup slots.
PD4-4477383056	Summit X770 series switches produce kern card error messages because of difficulty allocating memory for DAD message when trying to scale to 2,000 L3 IPv6 VLANs.
PD4-4176268416	IPv6 traffic matching with long mask route is not forwarded after the command <code>clear fdb</code> .
PD4-4176268381	IPv6 traffic across slots is not forwarded at line rate during run failover "master slot" and traffic is forwarded through software. Traffic resumes running at line rate after master and backup nodes are "In sync" after run failover.
PD4-4238470671	Jumbo traffic is software-switched despite having ingress and egress ports enabled for jumbo frames. This issue happens when using the command <code>enable jumbo-frame ports <port no.s></code> , but works when using the command <code>enable jumbo-frame ports all</code> .
PD4-4237445629	"PTP Rx" process shows more than 40% CPU utilization by default, even when no network-timing license is enabled.
PD4-4237445447	When two ports are configured as PTP slave ports, with one port configured as active, then the output of the command <code>show port net ptp ports</code> , shows the port status as "passive" when it should show "slave".
PD4-4237445563	PTP domain number does not change when the PTP clock is deleted, and then re-created with a different domain number.
PD4-4287760650	Routing table does not get updated after resetting RIP neighborhood when two RIP routers are exporting the same routes to two different ASBRs.
PD4-3865986001	OSPF state moves from "waiting" to "DR" immediately after configuring the priority value on the OSPF-enabled VLAN.
PD4-3927422131	Changing instance ID along with timer causes OSPFv3 never comes up.
ACLs	
PD4-4181588324	ACL rule to match all IPv6 packets is incorrectly matching all other packets when using match condition "source-address ::/0".
PD4-4174814835	In Summit X480 series switches, you cannot create 4,096 ACL meters.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
PD4-4214618905	Memory leak occurs after creating, and then deleting ACL network-zones with IP/MAC attributes.
PD4-4214618709	Memory leak occurs after configuring, and then unconfiguring ACL policies that have network-zones.
PD4-4214074302	The access-list policy is not applied when using destination-zone attribute in the rule entry for the first time.
PD4-4213889990	ACL process signal 6 ends unexpectedly when refreshing zone policies.
PD4-4214074143	ACL that permits a specific TCP port range does not work if the destination-zone attribute is present in the policy file.
BGP	
PD4-4456405581	Process BGP ends unexpectedly with signal 11,when sending a route with next hop as "0.0.0.0".
PD4-4489667001	BGP export policy with match condition "next-hop" does not work as expected.
FDB	
PD4-4433761990	With a stack of two Summit X440-48T series switches with dynamic VLAN enabled, after performing a failover, FDB information on the VLAN on the backup node is not in sync with the master.
IP Protocols	
PD4-4552704903, PD4-4552704871, PD4-4552418866	In BlackDiamond X8 and BlackDiamond 8800 series switches, with distributed ARP turned on, L3 traffic doesn't egress after a link failover even though ARP is learned and is pointing to the correct port. The traffic starts flowing after executing the command <code>clear iparp</code> .
PD4-4472675365	Multi-cast cache creation for reserved multi-cast addresses (224.0.0.x/8) for VLANs with 100+ active ports, can fail. This is only applicable for VLANs with IP address configured.
PD4-4347955207	ISIS hello packets are dropped when their IP-MTU size is greater than 1,500.
PD4-4477383021	Deleting member VLANs from ERSP domains, and then running the command <code>run ip dad</code> produces the following error messages: BD-8806.47 # run ipv6 dad esrpv15 Ignored 'run dad' command for fe80::204:96ff:fed:2f0, interface esrpv15: DAD failed for link-local prefix address on interface esrpv15 Ignored 'run dad' command for 3a02:2028:ff00::f:300, interface esrpv15: DAD failed for link-local prefix address on interface esrpv15 BD-8806.45 # run ip dad esrpv15 Ignored 'run dad' command for 4.1.15.1, interface esrpv15: DAD failed for link-local prefix address on interface esrpv15
PD4-4341188919	When a VLAN's IPv6 address is configured with mask 97 and above, OSPFv3 sends the full IPv6 address, instead of its prefix in the Link LSA address prefix field.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
PD4-4341188919	Unconfiguring, and then reconfiguring an IPv6 addresses on loopback VLAN causes OSPFv3 process to end unexpectedly with signal 11.
PD4-4200853291	After a failover, RTMgr is not updating the route table to reflect the use of a RSVP-TE LSP.
PD4-4196582666	Improve ECMP results by adding a CLI command to choose the IP route sharing hardware hash algorithm.
PD4-4068090296	<p>After issuing the command <code>configure iparp distributed-mode</code> on data traffic to destinations via tunnel is slow-path forwarded.</p> <p>Also, the following error messages appear after rebooting the switch or gateway to tunnel endpoint changes:</p> <pre>10/01/2013 15:00:30.91 <Warn:Kern.IPv4FIB.Warning> Slot-2: dest 0x02036400 / 24 nexthop 0x06010301: Unable to add route to unit 1, rc Entry not found. Shadow problem. 10/01/2013 15:00:30.96 <Warn:Kern.IPv6FIB.Warning> Slot-2: dest</pre>
PD4-4196582631	IPv6 OSPF traffic loss occurs between switches because traffic is sent with source MAC address as 00:00:00:00:00:00.
MPLS	
PD4-4333899744	L2 VPN sharing with spoke pseudowires does not work.
PD4-4444563018	Untagged VMAN-VPLS traffic is getting egress filtered when it has the same ports as CEP-VPLS, and egress filtering is enabled on the port.
PD4-4378481260	Ping fails from customer edge to provider edge access VLAN when access VLAN is associated with VPLS.
PD4-4375495944	Traffic is not load balanced when backup LSP is active using RSVP-TE.
PD4-4200906536	CEP-VPLS egress filtering isn't working after booting.
PD4-4179247910	With LSP Loadsharing, pending ECMP LSPs are not installed in hardware when the current LSP is deleted from hardware.
PD4-4238911733	With loadsharing enabled on VPLS ports on PE nodes when L2 VPN sharing is enabled and VPLS is up, hal process ends unexpectedly.
PD4-4289363711	IP routing tables are cleared and not updated again after an management module failover with OSPF graceful restart enabled along with the MPLS RSVP-TE protocol.
PD4-4184397758	MPLS process isn't aware of CEP attachment circuit after rebooting.
OpenFlow	
PD4-4194081467	Repeated enabling, and then disabling OpenFlow causes OpenFlow to not be enabled on any VLANs that have been configured for OpenFlow, and the controller status becomes inactive.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
OSPF	
xos0053410	OSPFv3 is sending IPv6 address instead of its prefix in the address prefix field of Link LSA, when the IPv6 address of vlan is configured with mask 97 and above.
ScreenPlay	
PD4-4345139476	Unable to disable/enable ports using ScreenPlay.
PD4-4207566174	ScreenPlay shows incomplete details for ports and VLANs.
Security	
PD4-3475078057	The command <code>create netlogin local-user</code> used with any option other than <code>password</code> for encrypted returns an error.
PD4-4484220106	Issuing the command <code>enable ip-security source-ip-lockdown</code> on a few ports causes ACL resources to be used inefficiently resulting in ACL slice full condition.
PD4-4203427942	Dynamic ACL entries are not cleaned up when load sharing is disabled on DHCP snooping enabled ports that are a LAG/trunk.
PD4-4193006645	Process <code>ipSecurity</code> ends unexpectedly when 64-byte TCP/IP frames with flags TCP FIN, URG, and PSH bits are received.
PD4-4399596381	With a VLAN with IDM-enabled ports with multiple FDB identities, enabling LLDP on both ports, and then verifying the detected IDM entries (FDB to LLDP correlation), IDM process ends unexpectedly with signal 6.
PD4-4483766224	The output from <code>show ssl</code> does not display complete information.
PD4-4345139603	Uninstalling SSH XMOD. produces the spurious error: "Error: Failed to install image - mount: mounting /dev/mtdblock2 on /exos failed: Device or resource busy"
PD4-3511925162	TACACS+/RADIUS: After configuring shared secret key in encrypted form with characters "&" and "<", the output of the command <code>show configuration aaa</code> shows a different secret key from what was actually configured.
PD4-4341188991	DoS protect in simulated mode logs that an ACL is added, but no traffic is blocked accordingly.
SNMP	
xos0053182	Help text for <code>configure snmp add community readonly</code> command is incorrect.
PD4-3714872811	SNMP set on atTable is not working.
PD4-4356959193	Get bulk on last OID on MIB tree should respond with a single <code>varbind</code> containing <code>endOfMibView</code> exception.

Table 18: Resolved Issues, Platform-Specific and Feature Issues in ExtremeXOS 15.5.1 (Continued)

ID Number	Description
TRILL	
PD4-4240009635	Disabling, and then enabling TRILL on two switches connected by a third switch causes the third switch to stop forwarding traffic even though the other two switches are still transmitting traffic to the third switch.
PD4-4356948681	TRILL process ends unexpectedly with signal 11 upon executing <code>configure trill maintenance-mode enable</code> command. Switch does not reboot, but. <code>show trill</code> shows maintenance-mode is still disabled.
VLANs	
PD4-4362588226	VLAN pid 1506 ends unexpectedly with signal 6 when performing an SNMP walk for OID 1.3.6.1.4.1.1916.1.4.17.1.
PD4-4172052360	MVRP fails to propagate XNV-created VLANs to other switches after a certain timeout, though traffic continues.
PD4-4172052318	Layer 2 traffic fails to recover if on of the MVRP-enabled nodes is rebooted in AVB setup.
VRRP	
PD4-4092277711	Enabling, and then disabling VRRP causes switches to become briefly unresponsive (13-18 secs). Switches still works fine executing VRRP command by opening another session of the switch.
PD4-4254552212	VRRPv2 goes to dual-master state when it is configured with sub-second advertisement intervals.

4 ExtremeXOS Documentation Corrections

This chapter lists corrections to the *ExtremeXOS 15.5 User Guide*.

This chapter contains the following sections:

- [ACL Egress Counters Limitation on page 136](#)
- [ACL Egress Counters Limitation on page 136](#)
- [ACL Policy Redirect on page 137](#)
- [Configure Access-List VLAN-ACL-Precedence Command Usage Guidelines on page 137](#)
- [Configure IP-MTU VLAN Command Syntax Description on page 138](#)
- [Debounce Commands on page 139](#)
- [ELRP on page 141](#)
- [ELRP and QoS on page 142](#)
- [L2VPN Sharing Commands on page 142](#)
- [Link Aggregation \(LAG\) Limit for Multiprotocol Label Switching \(MPLS\) Terminated Packets on page 143](#)
- [Mirroring on page 143](#)
- [MLAG on page 144](#)
- [MLAG PIM-SM *,G Forwarding Limitation on page 144](#)
- [NetLogin on page 144](#)
- [PoE Power Delivery on page 145](#)
- [Policies and Security on page 145](#)
- [Rate Limiting/Meters on page 146](#)
- [Routing Policies on page 147](#)
- [Synchronize Command on page 148](#)
- [TACACS Server on page 149](#)
- [Unconfigure Switch Erase Command on page 151](#)
- [VRRP on page 152](#)
- [VRRP Guidelines on page 153](#)

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.

ACLs General

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 Policy Redirect

In the *ExtremeXOS User Guide*, under *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.”

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 VLANbased 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-aclprecedence` 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 VLANbased 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.”

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 stackport 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:

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

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

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

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

In this example,

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

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

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

or

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



NOTE

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

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.

L2VPN Sharing Commands

ExtremeXOS Command Reference, L2VPN commands

xos0062252

Remove the following commands:

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

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

- enable l2vpn sharing
- disable l2vpn sharing

Platform Availability:

This command is available only on the following switches

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

Link Aggregation (LAG) Limit for Multiprotocol Label Switching (MPLS) Terminated Packets

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

xos0061631

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



NOTE

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

Mirroring

Basic Switch Operation ExtremeXOS User Guide

Chapter 8: “Configuring Slots and Ports on a Switch” > “Mirroring” > “Guidelines for Mirroring”

xos0058665

The following text should appear:

Under “Summit Family Switches”:

“One-to-many remote mirroring does not work as expected where ‘mirror-to’ ports could receive double-tagged packets. This is due to hardware limitation and applies to the following platforms: Summit X150, X250e, X350, X450, X450e, and X450a”.

Under “BlackDiamond X8, BlackDiamond 8800 Series Switches and SummitStack”:

“One-to-many remote mirroring does not work as expected where ‘mirror-to’ ports could receive double-tagged packets. This is due to hardware limitation and applies to the following platforms: BlackDiamond G48T, G48P, 10G4X, G24X, a-series, e-series, c-series (except 8900 modules), and 8500 series modules.”

MLAG

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

In the *ExtremeXOS User Guide*, under chapter *Netlogin > Configuring Network Login > Exclusions and Limitations*

xos0062943

Add the following additional limitation:

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

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.

Policies and Security

ExtremeXOS Concepts Guide

Chapter 20: "ACLs" > "Policy-Based Routing" > "Layer 2 Policy-Based Redirect"

xos0057861

The following note should appear



NOTE

"redirect-port" or "redirect-port-list" does not work for L3-switched packets matching ACL, if distributed IP ARP feature is turned on.

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."

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 tacacsauthorization” (if enabled), before disabling TACACS+.

Unconfigure Switch Erase Command

ExtremeXOS Command Reference for the `unconfigure switch` command

xos0059832

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

“`unconfigure switch`

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

Description

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

Syntax Description

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

SummitStack-specific parameters, and the switch rebooted.

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

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

VRRP

Layer-3 Unicast Protocols ExtremeXOS User Guide, Chapter 1: "VRRP"

xos0056279

The following VRRP guidelines should appear:

- 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.
- VRRP and other L2 redundancy protocols can be simultaneously enabled on the same switch.
- We do not recommend simultaneously enabling VRRP and ESRP on the same switch.
- When VRRP and BOOTP/DHCP relay are both enabled on the switch, the relayed BOOTP agent IP address is the actual switch IP address, not the virtual IP address.
- VRRP and ESRP cannot be configured on the same VLAN or port. This configuration is not allowed.
- RFC 5798 describes a situation where a master VRRP router takes on a duplicate IP address due to interaction with the duplicate address detection (DAD) feature. To prevent such duplicate addresses, the DAD feature is disabled whenever a VRRP router is configured for IPv6 or IPv4.
- A VRRP router instance can be configured with multiple IP addresses on the same subnet or on different subnets, provided that all virtual IP addresses match the subnet address of a VLAN on the switch. For example, if a host switch has VLAN IP addresses in the 1.1.1.x and 2.2.2.x subnets, then that VRRP router instance can contain virtual IP addresses in both those subnets as well.
- If a VRRP router instance is assigned priority 255, then the host router must own all the IP addresses assigned to the VRRP router instance. That is, each virtual IP address must match an IP address configured for a VLAN on the router.
- When a VRRPv2 instance spans routers using ExtremeXOS version 12.6 and earlier and routers using ExtremeXOS version 12.7 and later, routers using ExtremeXOS version 12.6 and earlier log packet-size warning messages.
- VRRP scaling numbers differs based on the license and hardware used; please refer the release notes for individual scaling limits.
- Seven unique VRIDs can be configured on a router.
- The maximum number of VIPs supported for a single VRRP instance is 255.

VRRP Guidelines

ExtremeXOS Concepts Guide

Chapter 31: “VRRP” under the heading “VRRP Guidelines”

xos0056279, xos0062263

The VRRP guidelines should change to the following:

“The following guidelines apply to using VRRP:

The maximum number of supported VRIDs per interface is seven.

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

- When a VRRPv2 instance spans routers using ExtremeXOS version 12.6 and earlier and routers using ExtremeXOS version 12.7 and later, routers using ExtremeXOS version 12.6 and earlier log packet-size warning messages.
- VRRP scaling numbers differs based on the license and hardware used; please refer the release notes for individual scaling limits.
- The maximum number of VIPs supported for a single VRRP instance is 255.”