



ExtremeXOS Release Notes

Software Version ExtremeXOS 15.3

Published: Feb 2013
Part Number: 120807-00 Rev 03



Copyright © 2001-2013 Extreme Networks

AccessAdapt, Alpine, Altitude, BlackDiamond, Direct Attach, EPICenter, ExtremeWorks Essentials, Ethernet Everywhere, Extreme Enabled, Extreme Ethernet Everywhere, Extreme Networks, Extreme Standby Router Protocol, Extreme Turbodriven, Extreme Velocity, ExtremeWare, ExtremeWorks, ExtremeXOS, Go Purple Extreme Solution, ExtremeXOS ScreenPlay, ReachNXT, Ridgeline, Sentriant, ServiceWatch, Summit, SummitStack, Triumph, Unified Access Architecture, Unified Access RF Manager, UniStack, XNV, the Extreme Networks logo, the Alpine logo, the BlackDiamond logo, the Extreme Turbodriven logo, the Summit logos, and the Powered by ExtremeXOS logo are trademarks or registered trademarks of Extreme Networks, Inc. or its subsidiaries in the United States and/or other countries.

MD5C.C - RSA Data Security, Inc., MD5 message-digest algorithm

Copyright (C) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved. License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest

Algorithm" in all material mentioning or referencing this software or this function. License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind. These notices must be retained in any copies of any part of this documentation and/or software.

Copyright (c) 1995-1998 by Cisco Systems, Inc.

Permission to use, copy, modify, and distribute this software for any purpose and without fee is hereby granted, provided that this copyright and permission notice appear on all copies of the software and supporting documentation, the name of Cisco Systems, Inc. not be used in advertising or publicity pertaining to distribution of the program without specific prior notice be given in supporting documentation that modification, permission, and copying and distribution is by permission of Cisco Systems, Inc.

Cisco Systems, Inc. makes no representations about the suitability of this software for any purpose. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

sFlow is the property of InMon Corporation.

Specifications are subject to change without notice.

All other registered trademarks, trademarks, and service marks are property of their respective owners.

Extreme Networks, Inc.

3585 Monroe Street

Santa Clara, California 95051

(888) 257-3000

(408) 579-2800

<http://www.extremenetworks.com>

Table of Contents

Chapter 1: Overview 5

New Features and Functionality in ExtremeXOS 15.3	5
Multiple Stream Registration Protocol (MSRP) for Audio Video Bridging (AVB)	6
Limitations	6
Supported Platforms	7
Multiple Registration Protocol (MRP)/Multiple VLAN Registration Protocol (MVRP)	7
Limitations	7
Multi-protocol Label Switching (MPLS) Layer 3 Virtual Private Network (L3 VPN)	8
Limitations	8
Supported Platforms	9
Network Time Protocol (NTP)-Virtual Router Redundancy Protocol (VRRP) Virtual IP	9
Requirements	9
Port Isolation	10
New CLI Command	10
Limitations	10
Supported Platforms	10
Ethernet Ring Protection Switching (ERPS) G.8032 Enhancements	11
Limitations	11
ExtremeXOS Network Virtualization (XNV) Per Virtual Machine (VM) Statistics	12
Limitations	12
Ethernet Automatic Protection Switching (EAPS) License Change	12
Identity Management (IDM) OR Operation and Active Directory (AD) Group Attribute Support	13
Limitations	13
IPv6 Equal-Cost Multi-Path (ECMP)	13
Supported Platforms	13
CLI Commands	14
Protocol Independent Multicast (PIM) IPv6	14
Limitations	14
Link Aggregation Group (LAG) Scaling Enhancements	15
Limitations	15
Service Verification Tool	15
CLI Commands	15
Limitations	15
Supported Platforms	16
OpenFlow	16
CLI Comamnds	17
Limitations	17
Supported Platforms	17
Generic Routing Encapsulation (GRE) Tunnel Support	18
Limitations	18
Synchronous Ethernet (SyncE) to Derive Timing for Precision Time Protocol (PTP)	18
CLI Commands	19
Multi-switch Link Aggregation Groups (MLAG)-Link Aggregation Control Protocol (LACP)	19
CLI Commands	19
Multi-session Mirroring	20
Limitations	20
ExtremeXOS Network Virtualization (XNV) Dynamic VLAN	20
Limitations	21
New CLI commands	22
OpenStack	22
Use Cases	23
Layer 2 Multicast Scaling	23
Limitations	23
CLI Commands	24

255-Character Port Description String	24
Limitations	24
CLI Commands	24
Protocol Independent Multicast (PIM) Register Filtering	24
CLI Commands	24
Flow Redirects Increased from 32 to 256	25
Command to Locate a Switch Using Front Panel LEDs	25
CLI Commands	25
Supported Platforms	25
New Hardware Supported in ExtremeXOS 15.3	25
ExtremeXOS Hardware and Software Compatibility Matrix	25
Upgrading to ExtremeXOS	26
Downloading Supported MIBs	26
ExtremeXOS Command Line Support	26
Tested Third-Party Products	27
Tested RADIUS Servers	27
Tested Third-Party Clients	27
PoE Capable VoIP Phones	27
Extreme Switch Security Assessment	28
DoS Attack Assessment	28
ICMP Attack Assessment	28
Port Scan Assessment	28
Service Notifications	28

Chapter 2: Limits 29

Supported Limits	29
------------------	----

Chapter 3: Open Issues, Known Behaviors, and Resolved Issues 69

Open Issues	69
Known Behaviors	93
Resolved Issues in ExtremeXOS 15.3	97



1 Overview

These release notes document ExtremeXOS® 15.3, which includes software feature enhancements, new hardware platform support, and resolves software deficiencies. This chapter contains the following sections:

- [New Features and Functionality in ExtremeXOS 15.3 on page 5](#)
- [New Hardware Supported in ExtremeXOS 15.3 on page 25](#)
- [ExtremeXOS Hardware and Software Compatibility Matrix on page 25](#)
- [Upgrading to ExtremeXOS on page 26](#)
- [Downloading Supported MIBs on page 26](#)
- [ExtremeXOS Command Line Support on page 26](#)
- [Tested Third-Party Products on page 27](#)
- [Extreme Switch Security Assessment on page 28](#)
- [Service Notifications on page 28](#)

New Features and Functionality in ExtremeXOS 15.3

ExtremeXOS 15.3 includes the following features:

- [Multiple Stream Registration Protocol \(MSRP\) for Audio Video Bridging \(AVB\) on page 6](#)
- [Multiple Registration Protocol \(MRP\)/Multiple VLAN Registration Protocol \(MVRP\) on page 7](#)
- [Multi-protocol Label Switching \(MPLS\) Layer 3 Virtual Private Network \(L3 VPN\) on page 8](#)
- [Network Time Protocol \(NTP\)-Virtual Router Redundancy Protocol \(VRRP\) Virtual IP on page 9](#)
- [Port Isolation on page 10](#)
- [Ethernet Ring Protection Switching \(ERPS\) G.8032 Enhancements on page 11](#)
- [ExtremeXOS Network Virtualization \(XNV\) Per Virtual Machine \(VM\) Statistics on page 12](#)
- [Ethernet Automatic Protection Switching \(EAPS\) License Change on page 12](#)
- [Identity Management \(IDM\) OR Operation and Active Directory \(AD\) Group Attribute Support on page 13](#)
- [IPv6 Equal-Cost Multi-Path \(ECMP\) on page 13](#)
- [Protocol Independent Multicast \(PIM\) IPv6 on page 14](#)
- [Link Aggregation Group \(LAG\) Scaling Enhancements on page 15](#)
- [Service Verification Tool on page 15](#)
- [OpenFlow on page 16](#)

- Generic Routing Encapsulation (GRE) Tunnel Support on page 18
- Synchronous Ethernet (SyncE) to Derive Timing for Precision Time Protocol (PTP) on page 18
- Multi-switch Link Aggregation Groups (MLAG)-Link Aggregation Control Protocol (LACP) on page 19
- Multi-session Mirroring on page 20
- ExtremeXOS Network Virtualization (XNV) Dynamic VLAN on page 20
- OpenStack on page 22
- Layer 2 Multicast Scaling on page 23
- 255-Character Port Description String on page 24
- Protocol Independent Multicast (PIM) Register Filtering on page 24
- Flow Redirects Increased from 32 to 256 on page 25
- Command to Locate a Switch Using Front Panel LEDs on page 25

Multiple Stream Registration Protocol (MSRP) for Audio Video Bridging (AVB)

Stream Reservation Protocol (SRP) enables bandwidth reservation for data streams from a talker end-station to one or more listener end-station(s) for Audio Video Bridging (AVB). SRP uses three MRP-based protocols to complete the end-to-end reservation. MVRP adds the ports to the VLAN where the data stream resides, MMRP optionally learns the listeners that are interested in the data streams, and MSRP reserves bandwidth for the stream. MSRP in turn uses Forwarding and Queuing for Time-Sensitive Streams (FQTSS) to manage scheduling.

This feature includes:

- Support for SRP in accordance with IEEE 802.1Qat-2010 including support for:
 - IEEE 802.1Qat-2010 MSRP
 - IEEE 802.1Qav-2009 Forwarding and Queuing for Time-Sensitive Streams (FQTSS)
- Declaration and registration of MSRP domain discovery and reservation
- Reserve link capacity for the user of SRP streams

Limitations

- Stacking is not supported.
- IEEE8021-SRP-MIB is not currently supported.
- Talker pruning is not available, since MMRP is not available.
- 802.11 designated MSRP node (DMN) support is not available.
- LAG and MLAG are not supported.



- The only supported layer 2 topology protocols are STP and RSTP. In particular, MSTP, EAPS, and ERPS are not supported, and MVRP/MSRP should not be enabled on the same ports as MSRP, EAPS or ERPS.
- When using AVB with STP or RSTP, VLAN “default” must be the carrier VLAN for the STP domain on the ports where AVB is enabled.
- Maximum number of active streams:
 - For the Summit X440 and X460: 1024.
 - For the Summit X670: 8,192.

Supported Platforms

- Summit X460 series switches
- Summit X670/X670V series switches
- Summit X440 series switches

Multiple Registration Protocol (MRP)/Multiple VLAN Registration Protocol (MVRP)

Multiple Registration Protocol (MRP) is a simple, fully distributed, many-to-many protocol, that supports efficient, reliable, and rapid declaration and registration of attributes by multiple participants on shared and virtual shared media. MRP allows a participant of a given MRP application to make or withdraw declarations of attributes, which results in registration of those attributes with the other MRP participants for that application. MRP does not do any work on its own, but rather provides the framework and state machines for implementing MRP applications such as Multiple VLAN Registration Protocol (MVRP), Multiple MAC Registration Protocol (MMRP), and Multiple Stream Registration Protocol (MSRP).

This feature includes:

- MRP and MVRP implementation as per IEEE 802.1ak-2007
- Support for Secure Remote Password Protocol (SRP) for Audio Video Bridging (AVB)
- Support for XNV Dynamic VLAN propagation
- Support IDM role-based VLAN propagation
- Support for MVRP over individual Ethernet ports

Limitations

- MIB support is not available.
- MVRP support for EAPS is not available.
- MVRP over MLAG is not available.
- MMRP (part of IEEE 802.1ak-2007) is not supported.
- VMAN (S-VLAN) creation is not available.



Multi-protocol Label Switching (MPLS) Layer 3 Virtual Private Network (L3 VPN)

L3 VPNs provide the ability to interconnect IP networks across a shared MPLS BGP backbone. Networks interconnected using the same provider edge (PE) device may have overlapping IP addresses. Customer-specific IP addresses are separately managed using unique Virtual Routing and Forwarding (VRF) domains. Each VRF instance maintains a separate IGP topology and separate routing tables associated with each customer. This is also commonly referred to as light-weight L3 VRs.

This feature includes:

- Multiple VRF support
- Multi-instance (RIB) BGP protocol
- VPN-IPv4 address family support in BGP
- BGP carrying labeled VPN-IPv4 routes
- BGP for PE-CE peering routing protocol
- Static routes per VRF
- Two layer MPLS label stack for data traffic forwarding
- Ping and Traceroute network diagnostics tools under a VPN scope
- SNMP support to control and monitor parameters of all VR/VRFs
- BGP/MPLS VPN MIB access (read only), as per RFC-4382 (except two tables)
- OSPFv2 and ISIS as core (SP's backbone) IGP routing protocol
- RFC 1657 MIB support

Limitations

The following are not supported:

- Static L3VPN
- RIP for PE-CE peering routing protocol
- IP Multicast BGP/MPLS VPN
- OSPFv2 and ISIS for PE-CE peering routing protocol
- IPv6 VPN
- Graceful restart mechanism for BGP with MPLS (RFC-4781)
- Constraint Route distribution for BGP/MPLS VPN (RFC-4684)
- Carrier of carriers BGP/MPLS VPN configuration (RFC 4364, Section 9)
- XML support to configure BGP/MPLS VPN parameters
- VR/VRF Management Account
- BGP Outbound Route Filtering (ORF)



- Inter-AS/inter-provider VPNs (RFC 4364, Section 10)
- Route leaking of internet default routes into the VRFs
- BGP-related MIBs, other than RFC 1657

Supported Platforms

MPLS L3 VPN feature is supported on all platforms that can support MPLS L3 VPNs:

- Summit X480 series switches
- Summit X460 series switches
- Summit X670 series switches
- BlackDiamond 8800 series switches with XL modules
- BlackDiamond X8 series switches
- E4G-200 and E4G-400 cell site routers

Network Time Protocol (NTP)-Virtual Router Redundancy Protocol (VRRP) Virtual IP

This feature adds the ability for switches to configure the Virtual Router Redundancy Protocol (VRRP) virtual IP as a Network Time Protocol (NTP) server address. The NTP server when configured on the VRRP master monitors the physical and virtual IP address for NTP clients.

Requirements

- For this feature to work correctly, you need to enable “accept” mode in VRRP using the following command:


```
configure vrrp vlan <vlan_name> vrid <vridval> accept-mode [on | off]
```

 Enabling accept mode allows the switch to process non-ping packets that have destination IP set to the virtual IP address.
- Summit switches configured as NTP clients need to have the following bootrom version:

- Summit X480, X460, X440, X670: 2.0.1.7
- Summit X150,250e, X350, X450a, X450e, X650, NWI-E450A:1.0.5.7

- The use of FHRP Virtual IPs is usually not recommended for NTP configuration since it can cause undesirable behavior when the NTP servers themselves are not in sync or if the delay is asymmetric. Therefore, ensure that both servers derive their clock information from the same source.

The problem may be more acute if there is a node connected to VRRP peers using MLAG and VRRP is in active-active mode. In this case, there is a theoretical possibility that every other packet can be sent to a different switch due to LAG hashing at the remote node.



Port Isolation

This feature blocks accidental and intentional inter-communication between different customers residing on different physical ports. Previously, this kind of security was obtained through the access-list module, but this can be complicated to manage and can be resource intensive. This feature provides a much simpler, more elegant, blocking mechanism without the use of ACL hardware.

You select a set of physical ports or load share ports which are deemed isolated. Once a physical port or load share port is isolated, it cannot communicate with other isolated ports, but can communicate with any other port in the system.

Blocked traffic types include:

- L2 unicast
- L2 multicast
- L2 unknown unicast
- L2 broadcast,
- L3 unicast
- L3 multicast

New CLI Command

The following command is available to enable isolation mode on a per-port basis (default = off):

```
configure ports <port_list> isolation [on | off]
```

The above command can be issued either on a single port or on a master port of a load share group. If the command is issued on a non-master port of a load share group, the command fails. When a port load share group is formed, all of the member ports assume the same 'isolation' setting as the master port.

Limitations

Port isolation is not allowed on the mirror-to port.

Supported Platforms

- BlackDiamond X8 series switches
- BlackDiamond 8800 series switches
- All Summit family switches



Ethernet Ring Protection Switching (ERPS) G.8032 Enhancements

The basic concept of G.8032/ERPS is that traffic may flow on all links of a ring network except on one link called the Ring Protection Link (RPL). The RPL owner is the node that blocks the RPL and the other node of the RPL is called the RPL neighbor node. All other nodes are called non-RPL nodes. When a link fails, the RPL owner unblocks the RPL to allow connectivity to the nodes in the ring. The G.8032/ERPS rings use a channel (dedicated path) for carrying their control traffic which is the R-APS messages.

The following enhancements have been added to the ERPS feature in ExtremeXOS:

- G.8032 version 2 with “no virtual channel support”
- Support for attaching to a Connectivity Fault Management (CFM) down-maintenance end point (MEP) configured external to ERPS
- Multiple failure protection for subrings using up-MEP (as per Appendix X.3 of G.8032 standard)

Limitations

- Backup master switch fabric module (MSM) failover and checkpointing for both v1 and v2 not available.
- In platforms that do not have hardware OAM, the recommended CFM interval is 1 second for link monitoring, which produces approximately 3 seconds of overhead in convergence times.
- Optimizations done in EAPS for Virtual Private LAN Service (VPLS) and any other within EAPS are not available.
- No interoperability with Spanning Tree Protocol (STP).
- Simple Network Management Protocol (SNMP) is not available.



ExtremeXOS Network Virtualization (XNV) Per Virtual Machine (VM) Statistics

For ExtremeXOS 15.3, per virtual machine (VM), for each direction, you can install counters to count ingress and egress traffic using the command:

```
configure vm-tracking vpp <vpp_name> counters [ingress-only | egress-only | both | none]
```

- For each local and network VPP, you can now specify whether counters need to be installed to count traffic matching VM MAC, which gets this VPP mapping.
- You can collect statistics on ingress traffic, egress traffic, or both. You can disable counters at any time.
- You can view the list of packets/bytes counts of this counter using CLI command `show access dynamic-counter`.
- The counter is un-installed when the VM MAC is deleted on the switch or VPP gets mapped to the VM MAC, or the counter option is set to none.
- If the VM MAC move happens, then the counter installed on previous port is un-installed and the counter is installed on new port. The counter values of old port are not maintained during the MAC move.

Limitations

During VM MAC moves, the values dynamic counters installed on the previous port are not maintained.

Ethernet Automatic Protection Switching (EAPS) License Change

The Ethernet Automatic Protection Switching (EAPS) “Full” (multiple rings with multiple interconnect points) feature is now part of the Advanced Edge license option, rather than the Core license option.



Identity Management (IDM) OR Operation and Active Directory (AD) Group Attribute Support

ExtremeXOS now supports the following enhancements to Identity Management (IDM):

- OR operation in match criteria of user roles.

Previously, ExtremeXOS supported only AND in the match criteria of user roles. Now OR is also supported—the user can have either AND or OR in the match criteria, but not both. That is, for a particular role, you can have all match criteria with AND, or have all the match criteria with OR. For role hierarchy and match criteria inheritance, there is no restriction across roles. You can have the parent role with AND, and the child role with OR, or vice versa. The inheritance of match criteria to the child role from the parent uses AND.

- Added support for Lightweight Directory Access Protocol (LDAP) group attributes.

Network users can be mapped to a role based on group membership (distribution list) information. When a user is detected by identity manager, it retrieves the groups that the detected user is member of from LDAP server. Identity manager then places the user under the appropriate role, based on group information and existing 8 LDAP attributes as supported today.

Limitations

Mix of AND and OR is not supported in the match criteria definition of the role.

IPv6 Equal-Cost Multi-Path (ECMP)

This feature adds IPv6 Equal-Cost Multi-Path (ECMP) support. Also, support is added for 16-way and 32-way ECMP for both IPv4 and IPv6, using static routes. Previous releases were limited to 2, 4, or 8-way ECMP.

Sharing of ECMP gateway sets now applies to IPv6 as well as IPv4. Sharing of ECMP gateway sets for IPv6 means the entire IPv6 Longest-Prefix Match (LPM) hardware capacity can use ECMP, across up to 32 gateways.

Supported Platforms

- Summit X460, X480, X650, X670 (stack or standalone)
- E4G-200, E4G-400 cell site routers
- BlackDiamond 8800 series switches with all I/O modules
- BlackDiamond X8 with all I/O modules



CLI Commands

The CLI command to enable/disable IPv6 ECMP is:

```
[enable | disable] iproute ipv6 sharing {{{vr} <vrname>} | { {vr} all}}}
```

The existing CLI command to configure the maximum number of gateways in each IPv4 or IPv6 gateway set now accepts 16 and 32 as acceptable values, along with 2, 4 and 8. As in prior releases, changing the value of max-gateways requires a save and reboot.

```
configure iproute sharing max-gateways <max_gateways>
```

Protocol Independent Multicast (PIM) IPv6

This feature provides Protocol Independent Multicast (PIM) support for IPv6. PIM is the de-facto standard for routing multicast traffic over the Internet.

PIM has two types, sparse and dense mode, meant for deployment in different topologies. These two flavors, called PIM-SM and PIM-DM, are entirely different in operation. PIM-SM is an explicit join protocol, where traffic is not forwarded on a segment, unless explicit an request come from the network segment (typically through IGMP). In contrast, PIM-DM is based on the flood-and-prune mechanism, where everybody receives the traffic until they explicitly inform (through PIM-DM prune mechanism) that they don't want to receive that particular stream. Thus, PIM-DM is mainly meant for topologies where listeners are very densely populated. PIM-SM should be deployed where the receivers are sparsely populated over the network, so that most of the network segments' bandwidth is conserved.

This new feature includes support for:

- Secondary address list

This is added to the V6 hello messages sent. The list includes all addresses assigned to an interface, including the link local addresses. The receiving router must process these addresses and must associate them with the neighbor that sent the message.

- Tunnel interface

This is very similar to a VLAN interface. You now receive IP address configuration for a tunnel, etc., from the VLAN manager client.

Limitations

The following features are not available:

- Embedded RP support
- Anycast RP support



Link Aggregation Group (LAG) Scaling Enhancements

For *BlackDiamond X8 series switches*: Load sharing groups configured for more than 16 aggregator ports per group no longer need the address-based “custom” algorithm.

For *Summit X670 family switches*: For Summit X670s in a stack, the maximum number of ports per group is increased from 8 to 64.

Limitations

- For Summit X670 switches and X670 stacking, load sharing groups configured for more than 16 aggregator ports per group must use the address-based “custom” algorithm.
- For Summit X670 stacking, all nodes in the stack must be Summit X670s if more than 8 aggregator ports per group are configured. For existing configuration with a LAG with more than 8 ports, any new non-X670 node added to the stack causes an EMS error.
- With Distributed ARP mode on, the maximum number of aggregator ports on BDx is limited to 16.

Service Verification Tool

This feature provides a test tool for operators to verify a new service instance prior to turning the service over to their customers.

CLI Commands

- `run esvt traffic-test {vlan} <vlan_name> loopback-port <loopback-port> peer-switch-ip <ipaddress> packet-size <packet_size> rate <rate> [Kbps|Mbps | Gbps] duration <time> [seconds | minutes | hours]`
- `stop esvt traffic-test {{vlan} <vlan_name>}`
- `show esvt traffic-test {{vlan} <vlan_name>}`
- `clear esvt traffic-test {{vlan} <vlan_name>}`

Limitations

- The service verification tool can be used to verify L2 services only and test network must not cross L2 boundaries.
- One loopback port needs to be assigned for each L2 service test.
- All L2/L3 protocols for the service VLAN should be disabled unless specifically included in the test tool instructions.
- Only modules that support egress ACLs are supported in stacking (BlackDiamond 8800 series switches).
- Remote switch must be an Extreme Networks switch, or a switch that does not generate ICMP errors, and packets are L3 routed within the same VLAN.
- Only one test per VLAN can be actively running.



Supported Platforms

- Summit X460 switches
- Summit X480 switches
- Summit X670 switches
- Summit X650 switches
- SummitStacks with X460, X480, X670, and X650 switches
- E4G-200 and E4G-400 cell site routers
- BlackDiamond 8800 series switches

OpenFlow

The OpenFlow feature enables an external OpenFlow Controller to manipulate data flows within an Extreme Networks switch using a standard protocol to dynamically configure a flow table abstraction. Flow table entries consist of a set of packet matching criteria (L2, L3, and L4 packet headers), a set of actions associated with a flow (flood, modify, forward, divert to controller, etc.), and a set of per flow packet and byte counters. Flow table entries are implemented using hardware ACLs.

Feature highlights:

- Supports line-rate implementation of the OpenFlow flow table by instantiating the flow table in hardware lookup tables.
- Provides the ability for multiple OpenFlow controllers to be configured, with automatic failover to the alternative controller if connectivity is lost with the currently active one.
- Provides the ability for particular ports and VLANs to be configured for OpenFlow control. A particular port can support OpenFlow-managed and non-OpenFlow managed VLANs. Currently, OpenFlow only supports a single VLAN.



NOTE

OpenFlow, XNV, and IDM are all features that enable an external agent to control resources on a switch. Due to their interaction models and resource requirements, these features are mutually exclusive. The ExtremeXOS 15.3 OpenFlow implementation prevents these services from being simultaneously configured on the same port.



CLI Comamnds

- [enable | disable] openflow
- show openflow ports [all | <port_list>] {{vlan} <vlan_name>}
- configure openflow controller [primary | secondary] [in-band [port <portNumber> | discovery] | out-of-band [active [ipaddress <ipaddress> | hostname <hostName>] {<tcpPort>} | passive <tcpPort>]] {tsl} {vr <vrName>} {rate-limit <rate> {burst-size <burstSize>}}
- unconfigure openflow controller [primary | secondary]
- show openflow controller {primary | secondary}
- show openflow
- debug openflow show flows [vendor-table | exos-tree]

Limitations

- Only supports a single VLAN.
- Supported platforms do not implement both packet and byte counters simultaneously on dynamic ACL entries. Only packet counters are supported.
- IN_PORT and FLOOD forwarding actions are not implemented in hardware, and are instead implemented in the user-space forwarding plane, limiting the throughput that can be sustained for flows using these actions.
- Flows are implemented using ACL hardware. Platform hardware has limitations on the simultaneous combinations of flow match conditions that can be supported. These limitations are described in the platforms' ACL release notes. When receiving a flow match combination that cannot be supported with the platform's ACL hardware, the switch generates an OpenFlow error message to the controller.
- All flow table entries are implemented as wide-key dynamic ACLs, limiting the potential scalability of the flow tables.
- Default emergency flows are not automatically installed, and there is no CLI command to specify them.
- NORMAL forwarding action is not supported.

Supported Platforms

- Summit X440 switches
- Summit X460 switches
- Summit X480 switches
- Summit X670 switches



Generic Routing Encapsulation (GRE) Tunnel Support

This feature allows you to create a Generic Routing Encapsulation (GRE)-based IPv4 tunnel, and route IPv4 traffic over it. This feature supports:

- IPv4-based GRE tunneling support
- Forwarding based on static routes

Limitations

- IPv4 only (both the routed traffic, and tunnel protocol).
- Unicast forwarding only, no Multicast.
- Single IP address can be configured on a GRE tunnel.
- Duplicate Address Detection (DAD) is not supported on GRE tunnels.
- No routing protocol support (RIP, OSPF, etc. etc.), only static routes.
- Maximum of 255 system wide tunnels (this includes any combination of GRE/6in4/6to4).
- On a chassis or SummitStack system, all blades/nodes need to support GRE before the feature can be enabled.
- The GRE capable hardware does not support VRs, so you cannot create tunnels in any other VR than VR-Default. This is the same behavior as for the 6in4/6to4 tunnels.



NOTE

When the hardware matches the tunnel source and destination addresses it does not look at the incoming VR, therefore it's recommended to not create any additional VRs on the switch when using tunnels. If you must, make sure that tunnel traffic on other VRs does not match the tunnels configured on the switch.

Synchronous Ethernet (SyncE) to Derive Timing for Precision Time Protocol (PTP)

The Precise Time Protocol (PTP) synchronizes the network by transferring the master clock information in the form of timestamps in the PTP messages (Sync/FollowUp/DelayReq/DelayResp). In the slave clock, the clock offset is computed through the reception of PTP messages that carry master clock as timestamps.

In practice, a network could employ multiple synchronization methods in the same network. Synchronous Ethernet (SyncE) transfers the frequency of the reference clock through Ethernet's physical layer. The frequency recovered from SyncE is highly accurate when compared to the frequency recovered through PTP messages. However, SyncE does not carry the Time-of-Day (TOD) or the Phase information of the clock as PTP does. Networks that employ SyncE and PTP for synchronization can leverage the accuracy of time transfer through PTP by using SyncE. Such Hybrid networks use SyncE for frequency transfer and PTP for Phase/Time-of-Day transfer.



CLI Commands

- `configure network-clock ptp time-source [network-frequency | ptp-frequency]`
- `show network-clock ptp configuration`

Multi-switch Link Aggregation Groups (MLAG)-Link Aggregation Control Protocol (LACP)

This feature introduces Link Aggregation Control Protocol (LACP) support over Multi-switch Link Aggregation Groups (MLAG) ports with the following options:

- The MLAG peer having the highest IP address for the ISC control VLAN is considered the MLAG LACP master. The switch MAC of the MLAG LACP master is used as the System Identifier by all the MLAG peer switches in the LACPDUs transmitted over the MLAG ports. This is the default option.
- You can configure a common unicast MAC address to be used on all the MLAG peer switches. This MAC address is used as the System Identifier by all the MLAG peer switches in the LACPDUs transmitted over the MLAG ports. This configuration (like any other configuration item is not checkpointed to the MLAG peers) and you have to make sure that the same MAC address is configured on all the MLAG switches. You have to ensure that this address does not conflict with the switch MAC of the server node that teams with the MLAG peer switches.

CLI Commands

```
configure {mlag peer} <peer_name> lacp-mac [auto | <lacp_mac_address>]
```



Multi-session Mirroring

Mirroring is a function on existing Extreme Networks switches which allows copies of packets to be replicated to additional ports without affecting the normal switching functionality. This feature has been revised to support:

- Up to 16 named mirror instances can be created. The system creates a default mirror ("DefaultMirror") at start of day. This default instance supports legacy CLI operation. You can define up to 15 additional instances.
- You can activate up to four mirror instances, and a mirror with both ingress and egress filters represents two instances. This includes the default instance. If the default instance is not configured, it is not activated.
- Disabling mirroring no longer removes the source filter information. Therefore, you can disable, and then re-enable mirroring if necessary.
- Show commands have been enhanced to allow the following:
 - Legacy show command behaves as before, showing information only about the "default" mirroring instance.
 - Show commands now show a summary list of the configured and enabled mirror instances.
 - Show commands allow you to request detailed information about individual filters.
 - Show commands show detailed information about all configured filters.

Limitations

- SNMP support is limited to what was previously supported in ExtremeXOS.
- No XML support.
- The total number of allowed sources configured in the system is 128. Specifically, only FP filters are a limited resource. Port-based mirroring has no hardware configuration limitations.

ExtremeXOS Network Virtualization (XNV) Dynamic VLAN

This change enhances the ExtremeXOS Network Virtualization (XNV) feature to include dynamic VLAN support:

- Ability to enable/disable the XNV dynamic VLAN feature on a per-port basis.
- Ability to add a VLAN tag and a VR as attributes to both Local and Network Virtual Port Profiles (VPPs).
- Specify the VLAN tag as an attribute to the VM using:
 - For VMs managed through RL, the VM to VLAN tag mapping can be learned from RADIUS or can be specified as part of the mapping entry in the VM MAP file.
 - For VMs managed through CLI, the VM to VLAN tag can be configured through CLI commands.



- If the specified VLAN does not exist, ExtremeXOS dynamically creates the VLAN when the VM is detected and deletes the VLAN when the last VM that uses the VLAN is deleted.
- For VMs sending tagged traffic, if no VLAN configuration exists for the VM, XNV creates the VLAN (assuming the VLAN does not already exist) with the received packet's tag and adds the port to the VLAN as tagged.

Limitations

- As part of VLAN definition, you can only specify the VLAN tag and optionally a VR. ExtremeXOS internally generates names for dynamically created VLANs.
- You can configure a maximum of one tag for the VM as part of either a VM or VPP configuration.
- All restrictions applicable for MVRP-created VLANs are applicable for VLANs created by this feature.
- Dynamically created VLAN are not saved across reboots. However dynamic VLAN information is checkpointed to standby nodes.
- Uplink ports are always added to dynamic VLANs as tagged.
- Since there is time lag between VM detection and programming of the VLAN in hardware, traffic for the first few milliseconds may be flooded on the internal (or default) VLAN or may be dropped.
- This feature is not compatible with NetLogin. This feature cannot be enabled on NetLogin enabled ports.
- Since this feature creates an internal VMAN for VM detection and adds the enabled port as untagged to the internal VMAN, you cannot add the port as untagged to other VMANs.
- Before enabling the dynamic VLAN, you have to add the port to a “default” or “base” VLAN. VMs sending untagged traffic that have no VLAN configuration are classified to this VLAN.



New CLI commands

- configure vlan dynamic-vlan uplink-ports [add {ports} <port_list> | delete {ports} [<port_list> | all]]
- show vlan dynamic-vlan
- [enable|disable] vm-tracking dynamic-vlan ports <port_list>
- configure vm-tracking vpp <vpp_name> vlan-tag <tag> {vr <vr-name>}
- unconfigure vm-tracking vpp <vpp_name> vlan-tag
- create vm-tracking local-vm mac-address <mac> {name <name> | ip-address <ip_address> | vpp <vpp_name> | vlan-tag <tag> {vr <vr-name>}}
- configure vm-tracking local-vm mac-address <mac> [name <name> | ip-address <ipaddress> | vpp <vpp_name> | vlan-tag <tag> {vr <vr-name>}]
- unconfigure vm-tracking local-vm mac-address <mac> [name | ip-address | vpp | vlan-tag]

OpenStack

OpenStack is an open source cloud operating system that manages pools of compute, storage and networking resources. OpenStack has three key projects to offer as-a-service capabilities: compute-as-a-service, storage-as-a-service, and network-as-a-service. This feature introduces a plug-in to the network-as-a-service application (Quantum) of OpenStack to provide network-as-a-service capabilities on ExtremeXOS-powered Extreme Network switches. This plug-in provides:

- Quantum plugin version 1.1 (Essex) and version 2.0 (Folsom) API support to enable tenants to define network of compute, storage and network services.
- Multitenant isolation with VLAN and VMAN. Multitenant isolation to the compute host via integration with vSwitch from OVS plug-in.
- Support for Tenant networks beyond 4K limits using VMAN.
- Support for VM motion across L2 boundaries.
- Support for L3 connectivity and forwarding to and from public/external network via SNAT and floating IPs with Quantum L3 router.
- Support for change management. All changes to tenant networks are logged and tracked for auditing. This is unique to the ExtremeXOS Quantum plug-in.
- Transaction Management: Since mapping of virtual to physical networks is a multi-step transaction, each of the transactions is tracked and errors are reported to the north-bound clients. If an error occurs during transaction execution, the entire transaction is rolled back.



- A topology-aware scheduler that chooses compute hosts for new virtual machines (VMs) in close proximity to the other VMs of the tenant. This proximity algorithm minimizes the east-west traffic load, and also reduces switch table sizes. Without this optimization, VMs are placed randomly on any host attached to any switch causing all intermediate switches to learn the VMs, thus causing MAC table exhaustion. The topology aware scheduler also enables tenant networks to scale over the 4K limit by carefully assigning VLAN IDs on a per pod basis. This feature is unique to the Extreme Quantum plug-in.

Use Cases

The previous features enable cloud service providers to build the following use cases:

- Cloud-in-a-box: A fully functioning cloud-in-a-box with components: switches, routers, load balancer, firewall, compute and storage. This is economical for enterprises to buy or for cloud resellers to offer as private, hosted cloud service in a data center.
- Private clouds in an enterprise: Enterprises are migrating to cloud services to offer the same services that public clouds offer to their employees. Some of the incentives for this move include: Bring your own device (BYOD), IT central policy enforcements, DevOps, etc.
- Hosted private clouds: As Enterprises try to reduce costs, some enterprises choose to outsource their IT and buy services from cloud service providers.
- Web hosting services: Many small and large web-hosting companies offer low cost web services to consumers. Also, e-commerce companies such as eBay, Zynga, etc. need cloud data centers to scale up and down with demand.

Layer 2 Multicast Scaling

This feature provides an option to use the L2 table for IP multicast forwarding database entries to increase scale. It provides an option to use both L2 and IPMC tables for IP multicast forwarding database entries.

Limitations

- The “mixed-mode” configuration option is not allowed on Summit X150, X250e, X350, X450e, X450a, series switches and BlackDiamond 8800 “e2-series” and 8500-G48T-e.
- When the “mixed-mode” configuration option is engaged on BlackDiamond 8800 series switches, newly inserted slots, which do not support “mixed-mode” fail initialization. On SummitStack, this same condition generates the following message every 30 seconds:


```
<HAL.Ipv6Mc.Error> Stack slot %d is incompatible with the multicast forwarding lookup configuration. Either remove this node from the stack or change the multicast forwarding lookup configuration.
```
- When using the “mac-vlan” configuration option:
 - PIMv4/V6, MVR features cannot be used.
 - IGMPv3 should not be used in conjunction with this mode.



- Private VLAN multicast should not be used.
- Issues with IP multicast address to MAC address mapping:

All IPv4 multicast frames use multicast MAC addresses starting with 01:00:5e:xx:xx:xx. The lower order 23 bits of the IP multicast address are used in the MAC address derivation. As only 23 bits of MAC addresses are available for mapping layer 3 IP multicast addresses to layer 2 MAC addresses, this mapping results in 32:1 address ambiguity. For example, 225.129.1.1 maps to the same MAC address 01:00:5e:01:01:01.

CLI Commands

```
configure forwarding ipmc lookup-key [group-vlan | source-group-vlan | mac-vlan | mixed-mode]
```

In the `show igmp snooping` command, the “Forwarding Lookup-Key” in the output is removed.

255-Character Port Description String

This feature creates a new and separate “description-string” field for each port, which can be up to 255 characters with the same restrictions that exist with the existing “display-string” field. This new field can be set and retrieved via the CLI (`show port` command), SNMP, and XML interfaces.

Limitations

- The following characters cannot be used: ‘ ‘ ‘, “<”, “>”, “:”, “<space>”, “&”
- SNMP set of the ifAlias element sets both the display-string and the description-string.

CLI Commands

- `[config | unconfig] port <port_list> description-string <string>`
- `config snmp ifmib ifalias size [default | extended]`

Protocol Independent Multicast (PIM) Register Filtering

This feature allows to you filter the register message based on the policy file configured on the First-Hop Router (FHR) and/or Rendezvous Point (RP) in the Protocol Independent Multicast-Sparse Mode (PIM-SM) domain. There is a register policy mechanism to filter out specific PIM register messages which have encapsulated specific (S, G) packets. The filtering allows a network administrator to detect/deny malicious multicast packets to flow into a multicast shared tree and then create a service blackout. This is supported for both PIM IPV4 and PIM IPV6 mode.

CLI Commands

```
configure pim {ipv4 | ipv6} register-policy {rp} [<rp_policy_name> | none]
```



Flow Redirects Increased from 32 to 256

The number of flow redirects is increased from 32 to 256 for both IPv4 and IPv6.

Command to Locate a Switch Using Front Panel LEDs

This new command causes the front panel LEDs to display a unique pattern to allow you to visually locate a switch. The command is not stored and does not survive a reboot.

CLI Commands

- `enable led locator {timeout [<seconds> | none]} {pattern <pattern>} {slot [<slot> | all]}`
- `disable led locator {slot [<slot> | all]}`

Supported Platforms

- All Summit series switches

New Hardware Supported in ExtremeXOS 15.3

- BlackDiamond X8 10G48T Module

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 the following URL:

http://www.extremenetworks.com/libraries/services/HW_SW_Compatibility_Matrix.pdf



Upgrading to ExtremeXOS

See “Software Upgrade and Boot Options” in the *ExtremeXOS Concepts Guide* for instructions on upgrading ExtremeXOS software. 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.
- For the BlackDiamond 8800 series switches, a hitless upgrade to ExtremeXOS 15.2 from an earlier release is not supported and should not be attempted. Use the normal software upgrade process for these switches.
- Hitless upgrade from ExtremeXOS 12.0 and earlier to ExtremeXOS 12.1 and later is not supported on the BlackDiamond 12800 switch.
- SummitX software is required for E4G cell site routers.

Downloading Supported MIBs

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

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

ExtremeXOS Command Line Support

The following is true for all Summit X150 and X350 series switches:

- Summit X150 and X350 series switches do not support L3 functionality; this platform does not support CLI commands for L3 functionality.
- Summit X150 and X350 series switches do not support stacking; all CLI commands for stacking are not supported on this platform.
- Summit X150 and X350 series switches do not support IP forwarding; however, CLI commands that configure IP addresses function in order to access the management functionality of the switch are supported.
- Upgrade or trial licensing is not available on the Summit X150 and X350 series switches.



Tested Third-Party Products

This section lists the third-party products tested for ExtremeXOS 15.2.

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 the following location:

http://www.extremenetworks.com/services/service_notification_form.aspx



2 Limits

This chapter summarizes the supported limits in ExtremeXOS 15.3.

Supported Limits

Table 1 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 Concepts Guide*.



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 1 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 1 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 1: Supported Limits

Metric	Product	Limit
Access lists (meters) —maximum number of meters.	BlackDiamond 8000 series	
	e-series, group of 24 ports	512
	a-series, group of 24 ports	1,024
	c-series	2,048 ingress, 256 egress
	BlackDiamond 8900 series	
	8900-10G24X-c, group of 12 ports	1,024 ingress, 256 egress
	8900 xl-series, 8900-G96T-c	4,096 ingress, 512 egress
	8900-40G6X-xm	512 ingress 512 egress
	BlackDiamond X8 series	512 ingress, 512 egress
	E4G-200	1,024 ingress 256 egress
	E4G-400	2,048 ingress 256 egress
	Summit X150, X250e, X350, X450e group of 24 ports, Summit X440, per group of 24 ports	512
	Summit X450a, per group of 24 ports	1,024
	Summit X460, per group of 24 ports	2,048 ingress, 256 egress
	Summit X480	4,096 ingress, 512 egress
	Summit 650, group of 12 ports	512 ingress, 256 egress
	Summit X670 with VIM4-40G4x Summit X480 with VIM3-40G4X Summit 650, group of 12 ports with VIM3-40G-4x	512 ingress 512 egress
Access lists (policies) —suggested maximum number of lines in a single policy file.	All platforms	300,000



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Access lists (policies) —maximum number of rules in a single policy file. ^a	BlackDiamond 8000 series	
	a-series, group of 24 ports	2,048
	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	61,440 (up to)
	8900-40G6X-xm	2,048 ingress, 1,024 egress
	BlackDiamond X8 series	2,048 ingress, 1,024 egress
	E4G-200	2,048 ingress, 512 egress
	E4G-400	4,096 ingress, 512 egress
	Summit X150, X250e, X350, X440, X450e group of 24 ports	1,024 ingress
	Summit X440	1,024 ingress
	Summit X450a, group of 24 ports	2,048 ingress
	Summit 460	4,096 ingress, 512 egress
	Summit X480	(up to) 61,440 ingress, 1,024 egress
	Summit X650, group of 12 ports	2,048 ingress, 512 egress
	VIM3-40G4x	2,048 ingress, 1,024 egress
	Summit X670	2,048 ingress
	VIM4-40G4x	1,024 egress
	Summit X480	2048 ingress
	VIM3-40G4X	1024 egress



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Access lists (slices) —number of ACL slices.	BlackDiamond 8000 series	
	a- and c-series, group of 48 ports	16
	e-series, group of 24 ports	8
	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 series	10 ingress, 4 egress
	E4G-200	8 ingress, 4 egress
	E4G-400	16 ingress, 4 egress
	Summit X150, X250e, X350, X450e, group of 48 ports	8 ingress
	Summit X450a, group of 24 ports	16 ingress
	Summit X440	4 ingress
	Summit 460	16 ingress, 4 egress
	Summit X480	17 ^b ingress, 4 egress
	Summit X650, group of 12 ports	12 ingress, 4 egress
	VIM3-40G4x	10 ingress, 4 egress
	Summit X670	10 ingress, 4 egress
	VIM4-40G4x	
	Summit X480	
	VIM3-40G4X	10 ingress 4 egress
AAA (local) —maximum number of admin and local user accounts.	All platforms	16
BFD sessions—maximum number of BFD sessions	All platforms (default timers)	512
	All platforms (minimal timers)	10 ^c
BGP (aggregates) —maximum number of BGP aggregates.	All platforms with Core license or higher	256
BGP (networks) —maximum number of BGP networks.	All platforms with Core license or higher	1,024
	BlackDiamond X8 series	1,024



Table 1: Supported Limits (Continued)

Metric	Product	Limit
BGP (peers) —maximum number of BGP peers.	BlackDiamond X8 series	512
	BlackDiamond 8000 series	512
	BlackDiamond xl-series	512
	Summit X450a, X460, X650, X670	128*
	Summit X480	512
	* With default keepalive and hold timers.	
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 (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 X450a, X460, X650, X670	25,000
	Summit X480	524,256 (up to) ^b
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 X450a, X460, X650, X670	25,000
	Summit X480	1,000,000
BGP ECMP —maximum number of equalcost multipath for BGP and BGPv6.	All platforms except Summit X440	2, 4, or 8
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 X450a, X460, X650, X670	25,000
	Summit X480	524,256 (up to) ^b



Table 1: Supported Limits (Continued)

Metric	Product	Limit
BGPv6 (unicast address-family routes) —maximum number of unicast address family routes.	BlackDiamond 8900 xl-series	20,000
	BlackDiamond 8800 a-, c-series	6,000
	BlackDiamond 8000 e-series	240
	Summit X450e,X250e	240
	Summit X450a, X460, X650	6,000
	Summit X480	20,000
	Summit X670	8,000
	BlackDiamond X8 series	8,000
BGPv6 (non-unique routes) —maximum number of non-unique BGP routes	BlackDiamond 8900 xl-series	24,000
	BlackDiamond 8800 a-, c-series	18,000
	BlackDiamond 8000 e-series	720
	Summit X450e,X250e	720
	Summit X450a, X460, X650	18,000
	Summit X480, X670	24,000
	BlackDiamond X8 series	24,000
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 pseudo wires —maximum number of CES TDM pseudo wires 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
	BlackDiamond 8000 series	32
	BlackDiamond X8 series	32
CFM —maximum number of CFM up end points.	Summit series	32
	BlackDiamond 8000 series	32
	BlackDiamond X8 series	32
	Summit series X460, E4G200, E4G400 (Non-load shared Ports)	256
	Summit series X460, E4G200, E4G400 (Load shared ports)	32
	Summit series	32
CFM —maximum number of CFM down end points.	All other platforms	32
CFM —maximum number of CFM remote end points per up/down end point.	All platforms	2,000
CFM —maximum number of dot1ag ports.	All platforms	128



Table 1: Supported Limits (Continued)

Metric	Product	Limit
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 X450a, X650, X670	2,048
	Summit X480	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 X450a, X480, X650, X670	10
	with 50 DACLS with 500 DACLS	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, E4G-200, E4G-400	32
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
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, E4G-200, E4G-400	32
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, E4G-200, E4G-400	32



Table 1: Supported Limits (Continued)

Metric	Product	Limit
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
Forwarding rate —maximum L2/L3 software forwarding rate.	BlackDiamond 8000 series	10,000 pps
	BlackDiamond X8 series	20,000 pps
	Summit series	10,000 pps
FDB (blackhole entries) —maximum number of unicast blackhole FDB entries.	BlackDiamond 8800 a-series	16,000
	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 series	128,000
	E4G-200, E4G-400	32,000
	Summit X150, X250e, X350, X450e	8,000
	Summit X440, X450a	16,000
	Summit X480	524,288 (up to) ^b
	Summit X460	32,000
	Summit X650	
	VIM3-40G4x	32,000
	Summit X670	
	VIM4-40G4x	128,000



Table 1: Supported Limits (Continued)

Metric	Product	Limit
FDB (blackhole entries) —maximum number of multi-cast blackhole FDB entries.	BlackDiamond 8000 series	1,024
	BlackDiamond X8 series	1,024
	Summit series	1,024
FDB (maximum L2 entries) —maximum number of MAC addresses.	BlackDiamond 8800 a-series	16,384
	BlackDiamond 8000 c-series	32,768
	BlackDiamond 8000 e-series	8,192
	BlackDiamond 8000 (system), except 8900 xl-series	128,000
	BlackDiamond 8900 xl-series	524,488 (up to) ^b
	BlackDiamond X8 series	128,000
	E4G-200, E4G-400	32,000
	Summit X150, X350, X250e, X450e	8,192
	Summit X440	16,000
	Summit X450a	16,384
	Summit X480	524,488 (up to) ^b
	Summit X460, 650	32,768
	SummitStack (except X480)	128,000
	Summit X670	128,000
FDB (Maximum L2 entries)—maximum number of multi-cast FDB entries.	BlackDiamond X8	1024
	BlackDiamond 8800	
	All Summit series switches	
FIP Snooping VLANs	BlackDiamond X8	768
	BlackDiamond 8800 (8900-40G6X-c only)	
	Summit X670	
	Summit X650 series	
FIP Snooping Virtual Links (FPMA mode) per port group	BlackDiamond X8	1,908
	BlackDiamond 8800 (8900-40G6X-c only)	
	Summit X670	
	Summit X650 series	
FIP Snooping FCFs (with perimeter port) per port group	BlackDiamond X8	238
	BlackDiamond 8800 (8900-40G6X-c only)	
	Summit X670	
	Summit X650 series	
FIP Snooping FCFs (with Enode-to-FCF port)	BlackDiamond X8	212
	BlackDiamond 8800 (8900-40G6X-c only)	
	Summit X670	
	Summit X650 series	
Identity management —maximum number of Blacklist entries.	All platforms	512



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Identity management —maximum number of Whitelist entries.	All platforms	512
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	500



Table 1: Supported Limits (Continued)

Metric	Product	Limit
IGMP sender —maximum number of IGMP senders per switch (IP multi-cast compression disabled). ^h NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 a-series	1,024
	BlackDiamond 8800 c-series	2,048 ^d
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900-10G24X-c modules	2,048 ^d
	BlackDiamond 8900-G96T-c modules	4,096 ^d
	BlackDiamond 8900-40G6X-xm	3,000 ^e
	BlackDiamond 8900 xl-series	4,096 ^d
	BlackDiamond X8 series	4,096 ^f
	E4G-200	2,048
	E4G-400	2,048
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	64
	Summit X450a	1,024
	Summit X480	4,096
	Summit X460	2,048
	Summit X650	2,048
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e
IGMP sender —maximum number of IGMP senders per switch (IP multi-cast compression enabled). ^h NOTE: Assumes source-group-vlan mode. For additional limits, see: • “Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode.” on page 49 • “Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode.” on page 49	BlackDiamond 8800 a-series	2,000 ^e
	BlackDiamond 8800 c-series	6,000 ^e
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 c-series	6,000 ^e
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond 8900-40G6X-xm	3,000 ^e
	BlackDiamond X8 series	6,000 ^{e f}
	E4G-200	3,000 ^e
	E4G-400	6,000 ^e
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	192 ^e
	Summit X450a	2,000 ^e
	Summit X460	6,000 ^e
	Summit X480	12,000 ^b
	Summit X650	6,000 ^e
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e



Table 1: Supported Limits (Continued)

Metric	Product	Limit
IGMP snooping per VLAN filters —maximum number of VLANs supported in per-VLAN IGMP snooping mode.	BlackDiamond 8800 a-series	1,000
	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 series	1,000
	E4G-200, E4G-400	1,000
	Summit X150, X250e, X350, X440 X450e	448
	Summit X450a, X460, X650, X670	1,000
	Summit X480	4,000
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. ⁱ	BlackDiamond 8800 c-series	2,000
	BlackDiamond 8900 c-series	2,000
	BlackDiamond X8 series	2,000
	Summit series (except Summit X460, X480, X650, and X670)	1,000
IGMPv2 subscriber —maximum number of IGMPv2 subscribers per switch. ⁱ	Summit X460, X480, X650, X670	2,000
	BlackDiamond 8800 c-series	20,000
	BlackDiamond 8900 c-series	20,000
	BlackDiamond X8 series	20,000
	Summit series (except Summit X480, X650, and X670)	10,000
IGMPv3 maximum source per group —maximum number of source addresses per group.	Summit X460, X480, X650, X670	20,000
	All platforms	250
IGMPv3 subscriber —maximum number of IGMPv3 subscribers per port. ⁱ	BlackDiamond 8800 a-, 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)	1,000
	Summit X460	2,000



Table 1: Supported Limits (Continued)

Metric	Product	Limit
IGMPv3 subscriber —maximum number of IGMPv3 subscribers per switch. ⁱ	BlackDiamond 8800 a-, 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)	10,000
	Summit X460	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).	All 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	8,000
IPv4 ARP entries in hardware with distributed mode on —maximum number of IP ARP entries in hardware with distributed mode on	Per BlackDiamond 8900-10G8X-xl, up to 260,000 per system	32,500 ^b
	Per BlackDiamond 8900-G48X-xl or 8900-G48T-xl, up to 130,000 per system	16,250 ^b
	Per BlackDiamond 8000 c-series, up to 18,000 per system	8,000
	BlackDiamond 8900-40G6X-xm, up to 22,000 per system	8,000
	BlackDiamond X8 series, up to 28,000 per system	12,000
	All other platforms	N/A
IPv4 ARP entries in hardware with minimum LPM routes —maximum recommended number of IPv4 ARP entries in hardware, with minimum LPM routes present. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series switches, assumes number of IP route reserved entries is 100 or less.	BlackDiamond 8800 a-, c-, xm-series	8,000
	BlackDiamond 8000 e-series	1,000 ^e
	BlackDiamond 8900 xl-series	16,000
	BlackDiamond X8 series	16,000
	E4G-200	8,000
	E4G-400	16,000
	Summit X440	412
	Summit X250e, X450e	1,000 ^e
	Summit X450a, X650, X670	8,000
	Summit X460, X480	16,000



Table 1: Supported Limits (Continued)

Metric	Product	Limit
IPv4 ARP entries in hardware with maximum LPM routes —maximum recommended number of IPv4 ARP entries in hardware, with maximum LPM routes present. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series, assumes number of IP route reserved entries is “maximum.”	BlackDiamond 8800 a-series	2,000 ^e
	BlackDiamond 8800 c-, xm-series	6,000 ^e
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 xl-series	12,000 ^e
	BlackDiamond X8 series	12,000 ^e
	E4G-200	6,000 ^e
	E4G-400	12,000 ^e
	Summit X440	380
	Summit X250e, X450e	500 ^e
	Summit X450a	2,000 ^e
	Summit X460, X480	12,000 ^e
	Summit X650, X670	6,000 ^e
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 a-series	14,000 ^e
	BlackDiamond 8800 c-series	18,000 ^e
	BlackDiamond 8000 e-series	1,000 ^e
	BlackDiamond 8900 xl-series	40,000 ^b
	BlackDiamond 8900-40G6X-xm	22,000 ^e
	BlackDiamond X8 series	28,000 ^e
	E4G-200	18,000 ^e
	E4G-400	20,000 ^e
	Summit X440	448
	Summit X250e, X450e	1,000 ^e
	Summit X450a	14,000 ^e
	Summit X460	20,000 ^e
	Summit X480	40,000 ^b
	Summit X650	18,000 ^e
	Summit X670	22,000 ^e
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 X250e, X450a, X450e, X460, X650, X670, E4G-400, E4G-200	25,000
	SummitStack or standalone	
	Summit X480 SummitStack or standalone	524,256 (up to) ^b



Table 1: Supported Limits (Continued)

Metric	Product	Limit
IPv4 routes (LPM entries in hardware) — number of IPv4 routes in hardware.	BlackDiamond 8800 a-, c-series	12,000
	BlackDiamond 8000 e-series	480
	BlackDiamond 8900 xl-series	524,256 (up to) ^{bg}
	BlackDiamond 8900-40G6X-xm	16,000 ^e
	BlackDiamond X8 series	16,000 ^e
	E4G-200, E4G-400	12,000
	Summit X440	32
	Summit X250e, X450e	480
	Summit X450a, X460, X650	12,000
	Summit X480	524,256 (up to) ^{bg}
	Summit X670	16,000 ^g
IPv6 addresses on an interface — maximum number of IPv6 addresses on an interface.	All platforms	255
IPv6 addresses on a switch —maximum number of IPv6 addresses on a switch	BlackDiamond 8000 series	512
	BlackDiamond X8 series	512
	E4G-200, E4G-400	512
	Summit X440	254
	Summit X460, X480, X650, X670	512
IPv6 host entries in hardware —maximum number of IPv6 neighbor entries in hardware.	BlackDiamond 8800 a-series	1,000 ^e
	BlackDiamond 8800 c-, xm-series	3,000 ^e
	BlackDiamond 8000 e-series	250 ^e
	BlackDiamond 8900-10G24X-c modules	2,000 ^e
	BlackDiamond 8900-G96T-c modules	4,000 ^e
	BlackDiamond 8900 xl-series	8,192 (up to) ^b
	BlackDiamond X8 series	3,000 ^e
	E4G-200	2,000 ^e
	E4G-400	3,000 ^e
	Summit X440	192
	Summit X250e, X450e	250 ^e
	Summit X450a	1,000 ^e
	Summit X460, X670	3,000 ^e
	Summit X650	2,000 ^e
	Summit X480	8,192 (up to) ^b



Table 1: Supported Limits (Continued)

Metric	Product	Limit
IPv6 routes (LPM entries in hardware) —maximum number of IPv6 routes in hardware.	BlackDiamond 8800 a-, 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 X250e, X450e	240
	Summit X450a, X460, X650	6,000
	Summit X670	8,000
	Summit X480	245,760 (up to) ^b
IPv6 routes with a mask greater than 64 bits in hardware —maximum number of such IPv6 LPM routes in hardware.	BlackDiamond 8000 a-, c-, e-, xm-series	256
	BlackDiamond 8000 xl-series	245,760 (up to) ^b
	BlackDiamond X8 series	256
	E4G-200, E4G-400	256
	Summit X250e, X440, X450e, X450a, X460, X650, X670	256
	Summit X480	245,760 (up to) ^b
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 X250e, X450a, X450e, X460, X650, X670, 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).	BlackDiamond X8 series	512
	All BlackDiamond 8000 series and Summit family switches with Edge license or higher	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 configurable gateways used by equal cost multipath OSPF, BGP, IS-IS, or static routes. Routing protocols OSPF, BGP, and IS-IS are limited to 8 ECMP gateways per destination.	All platforms	2, 4, 8, 16, or 32



Table 1: Supported Limits (Continued)

Metric	Product	Limit
IP route sharing (total destinations) —maximum number of unique destinations used by multipath OSPF, OSPFv3, BGP, IS-IS, or static routes. 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. For other platforms, each limit is based on up to 8 gateways per destination for routing protocols, or up to 32 gateways per destination for static routes.	BlackDiamond 8800 a-series, c-series	12,256
	BlackDiamond 8000 e-series	480
	BlackDiamond 8900 xl-series	524,256 (up to) ^b
	BlackDiamond 8900-40G6X-xm	16,352
	BlackDiamond X8 series	16,000
	Summit X250e, X450e	480
	Summit X450a, X460, X650, E4G-200, E4G-400	12,256
	Summit X480	524,256 (up to) ^b
	Summit X670	16,352
	E4G-200, E4G-400	12,256
IP route sharing (total combinations of gateway sets) —maximum number of combinations of sets of adjacent gateways used by multipath OSPF, BGP, IS-IS, or static routes.	BlackDiamond 8800 a-, c-, xl-, and xm-series	
	default maximum gateways of 4	510
	if maximum gateways is 2	1,022
	if maximum gateways is 8	254
	if maximum gateways is 16	126
	if maximum gateways is 32	62
	BlackDiamond 8000 e-series	
	default maximum gateways of 4	30
	if maximum gateways is 2	62
	if maximum gateways is 8	14
	if maximum gateways is 16	6
	if maximum gateways is 32	2
	BlackDiamond X8 series	
	default maximum gateways of 4	510
	if maximum gateways is 2	1,022
	if maximum gateways is 8	254
	if maximum gateways is 16	126
	if maximum gateways is 32	62
	Summit X460, X480, X650, X670, E4G-200, E4G-400	
	default maximum gateways of 4	510
	if maximum gateways is 2	1,022
	if maximum gateways is 8	254
	if maximum gateways is 16	126
	if maximum gateways is 32	62
IP multinetting (secondary IP addresses) —maximum number of secondary IP addresses per VLAN.	All platforms	64
IS-IS adjacencies —maximum number of supported IS-IS adjacencies.	BlackDiamond 8000 series	128
	BlackDiamond X8 series	128
	BlackDiamond 8900 xl-series	255
	Summit X450a, X460, X480, X650, X670	128



Table 1: Supported Limits (Continued)

Metric	Product	Limit
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 X450a	5,000
	Summit X480	30,000
	Summit X460, X650, X670	20,000
IS-IS IPv4 L1 routes in an L1 router —recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router.	BlackDiamond 8000 series	25,000
	BlackDiamond X8 series	25,000
	BlackDiamond 8900 xl-series	120,000
	Summit X450a	5,000
	Summit X480	50,000
	Summit X460, X650, X670	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 X450a	5,000
	Summit X480	50,000
	Summit X460, X650, X670	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 X450a	20,000
	Summit X460, X480, X650, X670	20,000
IS-IS IPv6 L1 routes in an L1 router —recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond 8900 xl-series	40,000
	Summit X450a	5,000
	Summit X480	25,000
	Summit X460, X650, X670	10,000



Table 1: Supported Limits (Continued)

Metric	Product	Limit
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 X450a	5,000
	Summit X480	25,000
	Summit X460, X650, X670	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 X450a	3,000
	Summit X480	15,000
	Summit X460, X650, X670	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 X450a	5,000
	Summit X480	40,000
	Summit X460, X650, X670	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 X450a	5,000
	Summit X480	40,000
	Summit X460, X650, X670	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 X450a	3,000
	Summit X460, X480, X650, X670	20,000
Jumbo frames —maximum size supported for jumbo frames, including the CRC.	All platforms	9,216



Table 1: Supported Limits (Continued)

Metric	Product	Limit
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, X650	8,000
	Summit X670	15,000
	Summit X150 ,X250, X350, X450e	2,000
	Summit X450a, X440	4,000
Layer-2 IPMC forwarding caches — (IGMP/MLD/PIM snooping) in mixed-mode. NOTE: IPv6 and IPv4 L2 IPMC scaling is the same for this mode.	BlackDiamond 8800 e-series switches	N/A
	BlackDiamond 8800 xl- and c-series switches	8,000
	BlackDiamond 8800 xm-series switches	15,000
	BlackDiamond X8, Summit X670	15,000
	E4G-200 and E4G-400 cell site routers, Summit X460, X480, X650	8,000
	Summit X150, X250, X350, X450, X450a, X450e	N/A
	Summit X440	4,000
Layer-3 IPMC forwarding caches —(PIM, MVR, PVLAN) in mixed-mode. ^e 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 series switches	6,000
	E4G-200 cell site routers, Summit X670	3,000
	E4G-400 cell site routers, Summit X460, X480, X650	6,000
	Summit X150, X250, X350, X450e, X450a	N/A
	Summit X440	192



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Load sharing —maximum number of load-sharing groups. NOTE: The actual number of load-sharing groups that can be configured is limited by the number of physical ports present in the switch or SummitStack.	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 using address-based custom algorithm	
	With distributed IP ARP mode off (default)	128
	With distributed IP ARP mode on	64
	SummitStack with X670 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
	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
Load sharing —maximum number of ports per load-sharing group.	BlackDiamond X8 series	64
	Summit X670 (non-stacked)	32
	SummitStack of all X670s	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-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



Table 1: Supported Limits (Continued)

Metric	Product	Limit
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 BlackDiamond X8 series Summit series	128 128 128
Mirroring, one-to-many (monitor port) —maximum number of one-to-many monitor ports.	All platforms	16
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.	BlackDiamond 8800 series BlackDiamond X8 series Summit series	1 1 1
MPLS LDP enabled interfaces —maximum number of MPLS LDP configured interfaces per switch.	Summit X460 Summit X480 Summit X480-40G VIM Summit X670 Summit X670V-48t BlackDiamond 8900 xl-series BlackDiamond 8900-40G6x-xm BlackDiamond X8 series E4G-200 E4G-400	32 64 64 32 64 64 64 64 32 32



Table 1: Supported Limits (Continued)

Metric	Product	Limit
MPLS LDP peers —maximum number of MPLS LDP peers per switch.	Summit X460	32
	Summit X480, Summit X480-40G VIM	64
	Summit X670	32
	Summit X670V-48t	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	50
	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	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	2,048
MPLS LDP transit LSPs —maximum number of MPLS transit LSPs per switch.	BlackDiamond 8900 xl-series	4,000
	BlackDiamond 8900-40G6X-xm	3,791
	BlackDiamond X8 series	4,000
	E4G-200	3,535
	E4G-400	4,000
	Summit X460, Summit X480	4,000
	Summit X670	3,725
	Summit X670V-48t	4,000
	Summit x480-40G VIM:	3,725



Table 1: 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	7,821
	BlackDiamond 8900-40G6X-xm	3,791
	BlackDiamond X8 series	7,821
	E4G-200	3,535
	E4G-400	7,525
	Summit X460, X480	7,821
	Summit X670	3,725
	Summit X670V-48t	7,821
	Summit x480-40G VIM	3,725
MPLS static LSPs—maximum number of static LSPs.	All platforms	100
MSDP active peers —maximum number of active MSDP peers.	BlackDiamond 8000 series	32
	BlackDiamond X8 series	64
	BlackDiamond 8900 series	64
	Summit X460, X480, X650, X670	16
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, X650, X670	8,000
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, X650, X670	4
Multi-cast listener discovery (MLD) IPv6 multi-cast data sender—maximum number of IPv6 multi-cast streams supported on a switch ^{h e} NOTE: Assumes source-group-vlan mode. For additional limits, see: “Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode.” on page 49 “Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode.” on page 49	BlackDiamond 8800 a-series BlackDiamond 8800 c-series BlackDiamond 8800 e-series BlackDiamond 8900 c-series BlackDiamond 8900-40G6X-xm BlackDiamond 8900 xl-series BlackDiamond X8 series E4G-200 E4G-400 Summit X150, X250e, X350, X450e Summit X440 Summit X450a Summit X460 Summit X480 Summit X650 Summit X670	750 1,000 250 1,000 1,000 3,000 3,000 1,500 3,000 250 90 750 3,000 3,000 1,500 1,500



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Multi-cast listener discovery (MLD) snooping per-VLAN filters—maximum number of VLANs supported in per-VLAN MLD snooping mode.	BlackDiamond a-series	500
	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 X150, X250e, X350, X450e	250
	Summit X450a	500
	Summit X480	2,000
	Summit X440	250
	Summit X650, X670, E4G-200	500
Multi-cast listener discovery (MLD)v1 subscribers—maximum number of MLDv1 subscribers per port ⁱ	BlackDiamond 8800 c-series	500
	BlackDiamond xl-series	1,500
	BlackDiamond X8 Series	1,500
	Summit X450a, X450e, X440, SummitStack	750
	Summit X460, X480, X650, X670	1,500
Multi-cast listener discovery (MLD)v1 subscribers—maximum number of MLDv1 subscribers per switch ⁱ	BlackDiamond 8800 series	10,000
	BlackDiamond X8 series	10,000
	Summit X450a, X450e, X440, SummitStack	5,000
	Summit X460, X480, X650, X670	10,000
Multi-cast listener discovery (MLD)v2 subscribers—maximum number of MLDv2 subscribers per port ⁱ	BlackDiamond 8800 c-series	500
	BlackDiamond xl series	2,500
	BlackDiamond X8 series	2,000
	Summit X450a, X450e, X440, SummitStack	1,000
	Summit X460, X480, X650, X670	2,000
Multi-cast listener discovery (MLD)v2 subscribers—maximum number of MLDv2 subscribers per switch ⁱ	BlackDiamond 8800 series	10,000
	BlackDiamond xl series	10,000
	Summit X450a, X450e, X440, SummitStack	5,000
	Summit X460, X480, x650, X670	1,0000
Multi-cast listener discovery (MLD)v2 maximum source per group—maximum number of source addresses per group	All platforms	200



Table 1: 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 a-series	1,024
	BlackDiamond 8800 c-series	2,048 ^d
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^d
	8900-G96T-c modules	4,096 ^d
	8900 xl-series	4,096 ^d
	8900-40G6X-xm	3,000 ^e
	BlackDiamond X8 series	4,096
	E4G-200	
	E4G-400	2,048
	Summit X150, X250, X350, X450e	500 ^e
	Summit X440	64
	Summit X450a	1,024
	Summit X480	4,096
	Summit X460	2,048
	Summit X650	2,048
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. ^e ” on page 49	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e
	BlackDiamond 8800 a-series	2,000 ^e
	BlackDiamond 8800 c-series	6,000 ^e
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 c-series	6,000 ^e
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 series	6,000 ^e
	8900-40G6X-xm module	3,000 ^e
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	192 ^e
	Summit X450a	2,000 ^e
	Summit X460	6,000 ^e
	Summit X480	12,000 ^b
	Summit X650	6,000 ^e
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e
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



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Network login —maximum number of dynamic VLANs.	All platforms)	2,000
Network login VLAN VSAs —maximum number of VLANs a client can be authenticated on at any given time.	All platforms	10
OSPF adjacencies —maximum number of supported OSPF adjacencies.	BlackDiamond 8000 series	128
	BlackDiamond 8900 xl-series	255
	BlackDiamond X8 Series	255
	Summit X250e, X460, X650, X670	128
	Summit X440	128
	Summit X480	255
	E4G-400, E4G-200	128
OSPF areas —as an ABR, how many OSPF areas are supported within the same switch.	All platforms	8
OSPF ECMP —maximum number of equal cost multipath OSPF and OSPFv3.	All platforms, except Summit X440	2, 4, or 8
OSPF external routes —recommended maximum number of external routes contained in an OSPF LSDB.	BlackDiamond 8000 series	20,000
	BlackDiamond 8900 xl-series	130,000
	BlackDiamond X8 series	20,000
	Summit X250e, X450a, X460, X650, X670	5,000
	Summit X480	130,000
OSPF inter- or intra-area routes —recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB with one ABR in OSPF domain.	BlackDiamond 8000 series	7,000
	BlackDiamond 8900 xl-series	7,000
	BlackDiamond X8 series	7,000
	Summit X250e, X450a, X460, X650, X670	2,000
	Summit X480	7,000
OSPF routers in a single area —recommended maximum number of routers in a single OSPF area.	BlackDiamond 8000 series	100
	BlackDiamond 8900 xl-series	200
	BlackDiamond X8 series	100
	Summit X250e, X450a, X460, X650, X670	50
	Summit X480	200
OSPF subnets on a single router —recommended maximum number of OSPF routed subnets on a switch.	All platforms with Core license or higher	400
OSPF virtual links —maximum number of supported OSPF virtual links.	All platforms with Core license or higher	32
OSPFv2 links —maximum number of links in the router LSA.	All platforms	419
OSPFv3 active interfaces —maximum number of OSPFv3 active interfaces.	All platforms with Advanced Edge license	4
OSPFv3 areas —as an ABR, the maximum number of supported OSPFv3 areas.	All platforms with Core license or higher	16



Table 1: Supported Limits (Continued)

Metric	Product	Limit
OSPFv3 external routes —recommended maximum number of external routes.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond 8900 xl-series	60,000
	Summit X450a, X460, X650, X670	10,000
	Summit X480	60,000
OSPFv3 interfaces —maximum number of OSPFv3 interfaces.	BlackDiamond 8000 series	256
	BlackDiamond X8 series	256
	BlackDiamond 8900 xl-series	384
	Summit X450a, X460, X650, X670	128
	Summit X480	384
OSPFv3 inter- or intra-area routes —recommended maximum number of inter- or intra-area routes.	BlackDiamond 8000 series	6,000
	BlackDiamond X8 series	6,000
	BlackDiamond 8900 xl-series	6,000
	Summit X450a, X460, X650, X670	3,000
	Summit X480	6,000
OSPFv3 neighbors —maximum number of OSPFv3 neighbors.	BlackDiamond 8000 series	64
	BlackDiamond X8 series	64
	BlackDiamond 8900 xl-series	128
	Summit X450a, X460, X650, X670	64
	Summit X480	128
OSPFv3 virtual links —maximum number of OSPFv3 virtual links supported.	All platforms with Core license or higher	16



Table 1: Supported Limits (Continued)

Metric	Product	Limit
PIM IPv4 snooping —maximum number of (S,G) entries programmed in the hardware (IP multi-cast compression disabled). NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 ^d
	BlackDiamond 8000 e-series	500 ^d
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^d
	8900-G96T-c modules	4,096 ^d
	8900 xl-series	4,096 ^d
	8900-40G6X-xm	3,000 ^e
	BlackDiamond X8 series	4,096
	E4G-200	2,048
	E4G-400	2,048
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	64
	Summit X450a	1,024
	Summit X460	2,048
	Summit X480	4,096
	Summit X650	2,048
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e
PIM IPv4 snooping —maximum number of (S,G) entries programmed in the hardware (IP multi-cast compression enabled). NOTE: Assumes source-group-vlan mode. For additional limits, see: “Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode.” on page 49 “Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode.” on page 49	BlackDiamond 8800 a-series	2,000 ^e
	BlackDiamond 8800 c-series	6,000 ^e
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 c-series	6,000 ^e
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 series	6,000 ^e
	E4G-200	3,000 ^e
	E4G-400	6,000 ^e
	8900-40G6X-xm	3,000 ^e
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	192 ^e
	Summit X450a	2,000 ^e
	Summit X480	12,000 ^b
	Summit X460	6,000 ^e
	Summit X650	6,000 ^e
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e



Table 1: Supported Limits (Continued)

Metric	Product	Limit
PIM IPv4—maximum routes —maximum number of (S,G) entries installed in the hardware (IP multi-cast compression disabled). NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 ^d
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^d
	8900-G96T-c modules	4,096 ^d
	8900 xl-series	4,096 ^d
	8900-40G6X-xm	3,000 ^e
	BlackDiamond X8 series	4,094
	E4G-200	2,048
	E4G-400	2,048
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	64 ^e
	Summit X450a	1,024
	Summit X480	4,096
	Summit X460	2,048
	Summit X650	2,048
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e
PIM IPv4—maximum routes —maximum number of (S,G) entries installed in the hardware (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. ^e ” on page 49	BlackDiamond 8800 a-series	2,000 ^e
	BlackDiamond 8800 c-series	6,000 ^e
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 c-series	6,000 ^e
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 series	6,000 ^f
	E4G-200	3,000 ^e
	E4G-400	6,000 ^e
	8900-40G6X-xm modules	3,000 ^e
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	192
	Summit X450a	2,000 ^e
	Summit X480	12,000 ^b
	Summit X460	6,000 ^e
	Summit X650	6,000 ^e
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e



Table 1: 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 ^d
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^d
	8900-G96T-c modules	4,096 ^d
	8900 xl-series	15,000
	8900-40G6X-xm	3,000 ^e
	BlackDiamond X8 series	4,094
	E4G-200	2,048
	E4G-400	2,048
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	64
	Summit X450a	1,024
	Summit X480	4,096
	Summit X460	2,048
	Summit X650	2,048
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e
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. ^e ” on page 49	BlackDiamond 8800 a-series	2,000 ^e
	BlackDiamond 8800 c-series	6,000 ^e
	BlackDiamond 8000 e-series	500 ^e
	BlackDiamond 8900 c-series	6,000 ^e
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 series	6,000 ^e
	E4G-200	3,000 ^e
	E4G-400	6,000 ^e
	8900-40G6X-xm	3,000 ^e
	Summit X150, X250e, X350, X450e	500 ^e
	Summit X440	192 ^e
	Summit X450a	2,000 ^e
	Summit X480	12,000 ^b
	Summit X460	6,000 ^e
	Summit X650	6,000 ^e
	VIM3-40G4x	3,000 ^e
	Summit X670	
	VIM4-40G4x	3,000 ^e



Table 1: 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 a-series	750
	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 X150, X250e, X350, X450e	250
	Summit X440	90
	Summit X450a	750
	Summit X460	3,000
	Summit X480	3,000
	Summit X650	1,500
	Summit X670	1,500
PIM IPv4 (maximum interfaces) —maximum number of PIM active interfaces.	All platforms	256
PIM IPv4 (maximum interfaces) —maximum number of PIM snooping enabled interfaces.	All platforms	256
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	256
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
PIM IPv6 Limits—maximum number of secondary address per interface	All platforms	70
PIM IPv6 Limits—static rendezvous points	All platforms	32



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Policy-based routing (PBR) redundancy —maximum number of flow-redirects.	All platforms	256 ^j
Policy-based routing (PBR) redundancy —maximum number of next hops per each flow-direct.	All platforms	32 ^j
Private VLANs —maximum number of subscribers. Assumes a minimum of one port per network and subscriber VLAN.	BlackDiamond 8800 a-, 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.	All platforms	512
Private VLANs —maximum number of private VLANs in an L2-only environment.	BlackDiamond 8800 a-, c-, e-series BlackDiamond 8900 series BlackDiamond X8 series E4G-200 E4G-400 Summit X440 Summit X250e, X450a, X450e Summit X480, X650 Summit X670 Summit X460	384 2,046 2,046 597 1,280 254 384 2,046 597 820
PTP/1588v2 Clock Ports	E4G Platforms	32 for boundary clock 1 for ordinary clock
PTP/1588v2 Clock Instances	E4G Platforms	2 combinations: - Transparent clock + ordinary clock - Transparent clock + boundary clock
PTP/1588v2 Unicast Static Slaves	E4G Platforms	40 entries per clock port
PTP/1588v2 Unicast Static Masters	E4G Platforms	10 entries per clock type



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Route policies —suggested maximum number of lines in a route policy file.	All platforms	10,000
RIP-learned routes —maximum number of RIP routes supported without aggregation.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond 8900 xl-series	10,000
	Summit X250e, X450a	3,000
	X440	3,000
	Summit X460	10,000
RIP interfaces on a single router —recommended maximum number of RIP routed interfaces on a switch.	Summit X480, X650, X670	10,000
	BlackDiamond 8000 series	256
	BlackDiamond X8 series	256
	BlackDiamond 8900 xl-series	384
	Summit X250e, X450a, X440	128
	Summit X460	256
RIPng learned routes —maximum number of RIPng routes.	Summit X480	384
	Summit X650, X670	256
	BlackDiamond 8000 series	3,000
	BlackDiamond X8 series	3,000
	BlackDiamond 8900 xl-series	5,000
	Summit X250e, X450a	1,500
RSVP-TE interfaces —maximum number of interfaces.	Summit X480	5,000
	Summit X460, X650, X670	3,000
RSVP-TE ingress LSPs —maximum number of ingress LSPs.	All platforms	32
RSVP-TE egress LSPs —maximum number of egress LSPs.	All platforms	2,000
RSVP-TE transit LSPs —maximum number of transit LSPs.	All platforms	2,000
RSVP-TE paths —maximum number of paths.	All platforms	2,000
RSVP-TE profiles —maximum number of profiles.	All platforms	1,000
RSVP-TE EROs —maximum number of EROs per path.	All platforms	1,000
Spanning Tree (maximum STPDs) —maximum number of Spanning Tree Domains on port mode EMISTP.	All platforms	64



Table 1: Supported Limits (Continued)

Metric	Product	Limit
Spanning Tree PVST —maximum number of port mode PVST domains. Note: Maximum of 7 active ports per PVST domain when 128 PVST domains are configured.	All platforms	128
Spanning Tree —maximum number of multiple spanning tree instances (MSTI) domains.	All platforms	64
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) Summit X460	500 600
Spanning Tree —maximum number of VLANs on all MSTP instances.	All platforms (except Summit X460) Summit X460	1,000 1,024
Spanning Tree (802.1d domains) —maximum number of 802.1d domains per port.	All platforms	1
Spanning Tree (number of ports) —maximum number of ports including all Spanning Tree domains.	All platforms	2,048
Spanning Tree (maximum VLANs) —maximum number of STP protected VLANs (dot1d and dot1w).	BlackDiamond 8900 series Summit X460 All other platforms	1,024 600 560
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 a-, c-, e-, xl-series BlackDiamond X8 series Summit X150, X350, X250e, X450a, X450e, X460, X480, X650, X670	1,024 1,024 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
Virtual routers —maximum number of user-created virtual routers that can be created on a switch. Note: Virtual routers are not supported on Summit X150, X250e, X350, X440, X450a, and X450e series switches.	BlackDiamond 8000 c-, xl-, xm-series BlackDiamond X8 series E4G-200, E4G-400 Summit X460, X480, X650, X670	63 63 63 63
Virtual router forwarding (VRFs) —maximum number of VRFs that can be created on a switch.	BlackDiamond 8000 c-, xl-, xm-series BlackDiamond X8 series Summit X460, X480, X650, X670	190 190 190



Table 1: Supported Limits (Continued)

Metric	Product	Limit
VRF forwarding instances —number of non-VPN VRFs that can be created on a switch.	BlackDiamond 8000 c-, xl-, xm-series Summit X460, X480, X650, X670	190 190
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) Summit X440	1,000 256
VLANs —includes all VLANs.	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.	BlackDiamond X8 series All BlackDiamond 8000 series and Summit family switches with Edge license or higher Summit X440	512 512 254
VLANs (maximum active port-based) —(Maximum active ports per VLAN when 4,094 VLANs are configured with default license)	Summit X670, X650, X480,X460, E4G-400 Summit X440 E4G-200 Summit X450e, X350, X250e, X150 Summit X450a	32 32 16 12 2 1
VLANs (maximum active protocol-sensitive filters) —number of simultaneously active protocol filters in the switch.	All platforms	15
VLAN translation —maximum number of translation VLANs. Assumes a minimum of one port per translation and member VLAN.	BlackDiamond 8000 a-, c-, e-, xl-series with eight modules of 48 ports 8900-G96T-c modules BlackDiamond X8 series Summit X450a and X450e, group of 24 ports with two-port option cards without option cards Summit series	383 767 767 25 23 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.	All platforms	512



Table 1: Supported Limits (Continued)

Metric	Product	Limit
VLAN translation —maximum number of translation VLAN pairs in an L2-only environment.	BlackDiamond 8800 a-, c-, e-series	384
	BlackDiamond 8900 xl-series	2,046
	BlackDiamond X8 series	2,046
	Summit X460, E4G-400, E4G-200	2,000
	Summit X440	512
	Summit X250e, X450a, X450e	384
	Summit X480, X650, X670	2,046
VPLS: VCCV (pseudo wire Virtual Circuit Connectivity Verification) VPNs —maximum number of VCCV enabled VPLS VPNs.	All platforms	16
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 series	128,000
	E4G-200, E4G-400	32,000
	Summit X460	32,000
	Summit X480	512,000
	Summit X670, Summit X670V-48t	128,000
	Summit X480-40G VIM	121,000
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	1,023
	Summit X480-40G VIM	1,023
VPLS peers —maximum number of VPLS peers per VPLS instance.	Summit X480	64
	Summit X460	32
	BlackDiamond 8900 xl-series	64
	BlackDiamond 8900-40G6x-xm	64
	Summit X670	32
	Summit X670V-48t, Summit X480-40G VIM	64
	BlackDiamond X8 series	64
	E4G-200, E4G-400	32



Table 1: Supported Limits (Continued)

Metric	Product	Limit
VPLS pseudo wires —maximum number of VPLS pseudo wires per switch.	BlackDiamond 8900 xl-series	7,800
	BlackDiamond 8900-40G6X-xm	4,000
	BlackDiamond X8 series	7,800
	E4G-200, E4G-400	1,000
	Summit X460	1,000
	Summit X480	7,800
	Summit X670	4,000
	Summit X670V-48t	7,800
	Summit X480-40G VIM	3,716
Virtual Private Wire Service (VPWS): VPNs —maximum number of virtual private networks per switch.	Summit X460	1,000
	Summit X480	4,000
	Summit X480-40G VIM	2,033
	Summit X670	2,047
	Summit X670V-48t	4,000
	BlackDiamond 8900 xl-series	4,000
	BlackDiamond 8900-40G6X-xm	2,047
	BlackDiamond X8 series	4,000
	E4G-200, E4G-400	1,000
VRRP (maximum instances) —maximum number of VRRP instances for a single switch.	BlackDiamond X8 series	255
	BlackDiamond 8800 c-series	255
	MSM-48c	
	BlackDiamond 8900 xl-series	255
	8900-MSM128	128
	All other platforms with Advanced Edge license or higher	
VRRP (maximum VRID) —maximum number of unique VRID numbers per switch.	All platforms with Advanced Edge license or higher	7
VRRP (maximum VRIDs per VLAN) —maximum number of VRIDs per VLAN.	All platforms with Advanced Edge license or higher	7
VRRP (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. Hello interval: 100 milliseconds Hello interval: 1 second	All platforms with Advanced Edge license or higher	
		2 4
VRRP (maximum iproute tracks) —maximum number of IP route tracks per VLAN.	All platforms with Advanced Edge license or higher	8
VRRP —maximum number of VLAN tracks per VLAN.	All platforms with Advanced Edge license or higher	8



Table 1: 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 BlackDiamond 8900 series with 100 DACLs with 500 DACLs Summit X450a, X480, X650, X670 with 100 DACLs with 500 DACLs	10 3 10 3 4 1
XNV authentication —maximum number of VMs that can be processed (combination of local and network VMs).	All platforms	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 Ingress Egress	2,048 512
XNV —maximum number of policies/dynamic ACLs that can be configured per VPP. ^k	All platforms Ingress Egress	8 4
XNV network VPPs —maximum number of XNV network VPPs. ^k	All platforms Ingress Egress	2,048 512

- a. The table shows the total available.
- b. Limit depends on setting configured for `configure forwarding external-tables`.
- c. When there are BFD sessions with minimal timer, sessions with default timer should not be used.
- d. Applies only if all enabled BlackDiamond 8000 I/O modules are BlackDiamond 8000 c-, xl-, or xm-series modules.
- e. 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, and Summit X460, X480, X650, and X670 switches compared to BlackDiamond 8800 a-series and 8000 e-series modules and Summit X250e, X450e, and X450a switches.
- f. 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.
- g. The limit depends on setting configured with `configure iproute reserved-entries`.
- h. 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.
- i. If IGMP and MLD are simultaneously configured on the switch, the number of effective subscribers supported would be appropriately lessened.
- j. Sum total of all PBR next hops on all flow redirects should not exceed 1024.
- k. The number of XNV authentications supported based on system ACL limitations.



3 Open Issues, Known Behaviors, and Resolved Issues

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

This chapter contains the following sections:

- Open Issues on page 69
- Known Behaviors on page 93
- Resolved Issues in ExtremeXOS 15.3 on page 97

Open Issues

The following are the open issues for supported features in ExtremeXOS 15.3.

Table 2: Open Issues, Platform-Specific and Feature PDs

PD Number	Description
General	
PD4-3312768711	Continuous pinging does not show ping statistics after canceling the ping by pressing CTRL + C.
PD4-3075940603	DHCPv6 control packets are CPU-forwarded, rather than forwarded via the hardware.
PD4-3317546201	Two issues in MAC-VLAN Mode: • Multicast cache entry for IGMP/MLD group is not deleted from hardware even after adding matching static FDB entry onto a different port in the same VLAN. • Multicast cache entry for IGMP/MLD group is not created in hardware after deleting matching FDB entry. Workaround: 1 Issue command: <code>clear igmp mld snooping vlan <vlannname> [vlannname on which the fdbentry is created or deleted]</code>. This command flushes the entire vlan. 2 Issue the command: <code>debug mcmgr clear <grpIp> <srcIp> <vlanName></code> [User has to know the srclp and grplp on which the cache is created].
PD4-3423282881	The virtual port filter is not working and it is missing from the output of the <code>show mirroring</code> command after deleting and adding the port to the VLAN, and then adding it back to the mirroring configuration. Workaround: Disable, and then enable mirroring.
PD4-3259170501	IPv6 multicast packets are being received by CPU when receiver VLAN is configured with PIM-SM and source VLAN is configured with PIM-DM. As a result, some traffic gets dropped.
PD4-3338939478	ESRPv6 + DAD: Duplicate IPv6 address is detected though there is no duplicate address while disabling/enabling ports between ESRP master and L2 switch.
PD4-3356030171	VRRP VRID limit should be increased to 255 based on the hardware stability.
PD4-3421465775	Process dot1ag ends unexpectedly with signal 11 when disabling ERPS ring after enabling CFM debug-data level log. Workaround: Do not enable debug.
PD4-3303209931	rtmgr process ends unexpectedly with signal 11 while rebooting neighboring switches.
PD4-3107133830	MSRP is not supported on LAG interfaces.
PD4-2991364782	In MVRP, dynamic VLAN propagation is not happening in MSTP (Multiple STP domains).

Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-2988947784	MVRP functionality is not working when integrated with LACP.
PD4-2988947661	MVRP feature is not working with static LAG.
PD4-3271203970	Debug message appears on console when enabling OpenFlow: "Nov 15 12:27:39 00001 lockfile INFO /var/run/openvswitch/.conf.db.-lock-: lock file does not exist, creating"
PD4-3251364953	Typo: "cliets" need to be changed to "clients" in CLI response when deleting CFM association.
PD4-3319542041	The total VLAN count in the <code>show vlan</code> command incorrectly includes the VMAN count. However, the command <code>show vman</code> shows the right number (it does not include VLAN count).
PD4-3057474297	Issuing command <code>create erps <></code> produces the following error message: "Unable to open sizing file system.size. System Size Data Structure updated with Default Values".
PD4-3107133811	"restart process gtp" and "terminate process gtp" sometimes cause the switch to reboot due to a kernel panic.
PD4-3289104822	PIMv6: Intermittently a router running in "dense" mode stops forwarding all received traffic when ports are restarted on the PMBR neighbor.
PD4-3298891721	PIMv6: Traffic drop occurs when all ports are restarted on a LHR/RP connected to a multi-access VLAN.
PD4-3312722306	AVB does not work with user VLAN. it is working only with default VLAN.
PD4-3312722295	AVB fails with different dot1w port encapsulation techniques like emistp and pvst-plus. It works with dot1d alone.
PD4-3298871993	BGP aggregation is not working after withdraw and advertise of routes.
PD4-3249665150	While MVR is enabled and a second VLAN is added as a MVR VLAN, traffic is not forwarded through its router ports.
PD4-3249665141	Traffic is not forwarded through primary/secondary ports when EAPS is disabled.
PD4-3289164206	Generic GTP error message. Need more specific message.
PD4-3289164196	Generic MRP error message. Need more specific message.
PD4-3346216771	BootPRelay - 6.1.1_7 Test script fails in <code>show dhcp-client state</code> command. Expected output is "BOOTP 172.16.250.1". Received no DHCP server IP in the output of the command, but current state shows "Received IP address configured on vlan".
PD4-3322786971	Advertisement interval shows "zero" in <code>show vrrp</code> command.
PD4-3218758136	ERPS: CFM association entry fails to get removed when disabling/enabling ERPS for the second time.
PD4-3320023671	The command <code>show iparp</code> fails to show the expected destination IP address after reboot.
PD4-3286100221	ACL_PERSISTANCE_RETRY - Verifying persistent Dynamic ACLs having higher priority compared to policy-based ACL, after save/ reboot, traffic failed after reboot.
PD4-3200169071	Within VR CLI context unable to configure the MLAG peer IP address without specifying the VR <vrname> parameters.
PD4-3259170494\	PIM-SM is configured on receiver and source VLAN. in this state when you change PIM mode from "sparse" to "dense" on the source VLAN, IPv6 multicast traffic to the receiver is stopped. The command <code>debug hal show ipv6mc</code> does not show the entry.
PD4-3328261697	L3 - Expected parameters failed to show up in the command <code>show iparp</code> in L3 module. Parameter details are : "Rejected Count" "Rejected Port" / "Dup IP Addr". This worked fine until ExtremeXOS v15.3.
PD4-3322820940	Traffic loss is occurring in OSPF-Graceful restart.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-3271246241	The command <code>show mpls l3vpn label received</code> displays no output.
PD4-3339092117	After enabling the trunk VLAN v7 at both FHR and RP, the output of the command <code>show pim cache detail</code> at RP does not update the source entry (S) flag for the source from FHR.
PD4-3213502873	VPNv4 routes received from peers are stored in the RIB only when there is a matching route target in the local VRF tables.
PD4-3296788112	Intra Area OSPF routes are not being learned properly.
PD4-3254256795	VRRP transition is not happening immediately when its receiving advertisements with priority zero.
PD4-3328850817	If a dynamic VLAN is added to OSPF, the OSPF configuration is retained even after the dynamic VLAN is removed.
PD4-3272987031	In the <i>ExtremeXOS Concepts Guide</i> , the configuration example given for ELSM needs to be corrected.
PD4-3282121907	On BlackDiamond X8 and BlackDiamond 8800 series switches, the ESRP state is a slave state when the maximum number of ESRP domain and track ping can be configured (256). However, the expected master state is seen on E4G-200 cell site routers. This issue occurs with ExtremeXOS 15.3 and 15.2.2.
PD4-3281007470	On BlackDiamond series switches, <code>show esrp domain1</code> command shows incorrect summary when one ESRP VLAN configured track ping and also track route.
PD4-3273653251	Packet loss is occurring after deleting/recreating the GRE tunnel.
PD4-3274384901	IPv6 address for ESRP domain cannot be configured and it shows 0.0.0.0 as VID, but when IPv4 address is also configured along with IPv6 then it shows IPv4 address as VID. This issue occurs in ExtremeXOS 15.3 and 15.2.
PD4-3322786980	VRRPv3 remains in backup state for 60 seconds.
PD4-3325786333	CLI execution fails and produces the following error: "enable sharing 6 grouping 6 10 algorithm address-based lacp" <Err:HAL.Port.Error> : Failed to disable static mac move drop on port 10.< > * (Engineering debug) X450e-24p.576 # - 21:14:38 < Illegal Line is 12/08/2012 21:14:37.32 <Err:HAL.Port.Error> : Failed to disable static mac move drop on port 10.
PD4-3189485339	Switch operation stops temporarily after issuing <code>delete vman test</code> command. This issue occur on ExtremeXOS 15.2.1 and 15.2.2, but not with 15.1.2.
PD4-3106307537	While start esvt test, existing FDB entry for the peer node is cleared from <code>show fdb table</code> automatically and port numbers connected to peer node are not showing in <code>show iparp table</code> . This floods traffic across all ports configured in service VLAN.
PD4-3172538561	VRRPv3: Lowest priority switch become the master with preempt mode "yes".
PD4-3076161855	Egress rate-limiting is allowed on MSRP enabled ports.
PD4-3238152635	NTP and DAD cannot co-exist. If you enable NTP and DAD on the same VLANs, save the configuration, and then reboot the switch, the NTP configuration for VLANs is removed when switch comes up.
PD4-3217352394	NTP packets from client do not egress until NTP is enabled on egress interface. This creates a problem when there are multiple gateways for the NTP server and the preferred gateway goes down. Need to enable NTP on all interfaces from which NTP server is reachable.
PD4-3217101680	NTP is not working when configured on loopback VLAN.
PD4-3223230501	The command <code>run msm fail-over</code> or <code>run fail-over</code> removes NTP configuration on a VLAN when NTP and DAD are enabled on the same VLAN.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-2885820591	Netlogin web is not working in ISP mode.
PD4-3000720361	Configuring NTP and then adding default route for reaching public NTP server causes re-enabling of NTP for synchronization. Until then, NTP won't get synchronized.
PD4-2900366121	On WindowsXP and earlier systems, Kerebros snooping does not work with user names that are 15 or more characters long. When a Kerberos client uses UDP for transport, some of the Kerberos packets exchanged between the client and the server may be fragmented, if the Kerberos user name is 15 or more characters long. Workaround: Ensure that the Kerberos client uses TCP for transport. For Microsoft WindowsXP and earlier clients, this patch from Microsoft ensures that the Windows Kerberos clients always use TCP for transport: http://support.microsoft.com/kb/244474
PD4-2483785534, PD4-2483785491	The option end-point should not be included in the configure vlan < vlan name > add ports < port no > tagged private-vlan command.
PD4-2255170647	After disabling and re-enabling a port using the <code>clear elsm port < port no > auto-restart</code> command, the ELSM state does not come up.
PD4-1402339707	The following Warning message should be deleted for BlackDiamond 8900 xl- and c-series cards when the switch is supporting 256 VRRP instances: WARNING: Number of VRs has reached the recommended maximum (128).
PD4-2330473390, PD4-2012884711	The following critical message may be logged followed by a system crash when making a configuration change to a private VLAN: System call ioctl failed failed: informCfgVlanAddPorts and 15
PD4-1688055111	A system crash occurs when the system is configured with 2,000 VPLS and 1,000 CFM instances while running the restart ports, or save and reboot commands.
PD4-1820554590	A switch sends an MLDv2 listener report even though MLDv1 is enabled on the switch.
PD4-1820554531	While running a TAHI conformance suite for MLDv2, a switch always sets the QQI value to 0 in MLDv2 general query messages. It should set QQI to querier query interval.
PD4-1842342875, PD4-1466022175	When an SNMP query is issued for non-existent IPv4 routes, the RtMgr process crashes with signal 11.
PD4-1842342815, PD4-1291631579	When working in network login, after a dot1x client logs out, the port is not moved to a MAC-based VLAN.
PD4-1842342767, PD4-1300978095	After installing a legacy CLI module, the CLI command <code>load script</code> returns the following error message: %% Unrecognized command: create vlan v\$x
PD4-1771325794	If the <code>configure ip-mtu</code> command is configured on VLANs that have only an IPv6 address, the <code>show configuration</code> command does not display the output for the <code>configure ip-mtu</code> command.
PD4-1540274936	An ISIS process crash with signal 6 may occur when disabling ISIS when ISISv6 and ISIS IP route compression are configured.
PD4-1556309411	With multiple NSSA areas where there is more than one link between areas, OSPF default routes are not installed in the routing table.
PD4-1549189647	In ISIS, when route summarization is configured with authentication, the authentication is not effective and all the routes are advertised, regardless of the type of authentication configured.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-1535268629	ISIS tx-only authentication also authenticates received LSPs. The received routes are not installed in the ISIS LSDB based on the authentication policy.
PD4-1620486143	When DHCP lease time is set to the maximum/infinity (4294967295), the DHCP client continuously sends renewal requests.
PD4-813961562	When a service VLAN is changed to include a dot1q tag on both sides in CFM VPLS, the RMEP entry is not learned on one side.
PD4-749060484	Errors are seen when a configuration having identifiers (SNMPv3 user name/EAPS domain name) with special characters are loaded through a script.
PD4-561358261	If you create a super VLAN and a sub-VLAN in different virtual routers you are able to bind the VLANs. Super VLANs and sub-VLANs should belong to the same virtual router.
PD4-460892051	Installing different versions of an ExtremeXOS image and an SSH image displays the following error message: Failed to install image- cannot read spec file" in the log "upgrade failed installation:got error from installer DLL"
PD3-132508261	When issuing the <code>enable jumbo-frame port all</code> command on a BlackDiamond 8800, the MTU size for the VLAN is not configured. Sending 5,000 byte traffic works correctly. However, if you disable jumbo-frames on the egress port the error message <code>Packet too big</code> is displayed.
PD3-104885349	When a shared link comes up, temporary traffic loss may occur until FDB entries are aged. Aging occurs by default every five minutes. Workaround: To reduce traffic loss, reduce the default age time.
PD3-132775269	Telnet sessions between two switches using two windows causes one session to hang if both sessions are edited but only one session is saved.
PD3-28378521	Enabling load sharing on a port that is being mirrored causes the mirroring to stop.
PD4-1678280326	File system commands such as <code>ls .*</code> exposes internal files.
PD4-2340987025	Bestpath is not computed when we have received routes from peer that is direct exported.
PD4-2365834125	When you attempt to delete VPLS peers, the following message is received: <code>Warn:FDB.VPLSPeerNotFound on deleting vpls peers.</code>
PD4-2377192305	Enabling and Disabling unicast-negotiation is not supported.
PD4-2383457620	PTP packets with TTL field set to one are discarded. The PTP master that sends packets with TTL as one will not be identified as master by the slave clock.
PD4-2402929258	printf message on console: <code>"Error: No valid Ethernet port(s) in the specified "taggedPorts" value """.</code>
PD4-2459996844	Multipath routes (ECMP) from EBGp peers are not displayed correctly for multi-cast routes.
PD4-2207141998	ERPS takes average of 186 ms to recover the L3 traffic when erps ring reverts from protection to idle state but G.8032 recommendation is to provide 50 ms recovery switching for Ethernet traffic in a ring topology.
PD4-2275772790	UDP CES packets are not forwarding to next hop when the CES TTL value is sets as 1.
PD4-2341681237	Some FDB entries were missing in network vlan file database while they are seen in subscriber VLAN.
PD4-2387510070	In a 7-node Summit Stack, the output of the command <code>show default</code> gets stuck after displaying some ports if the default VLAN has all the ports added.
PD4-2443869637	Looping is not prevented with VLAN aggregation using ERPSv2.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-2440608434	Default Sflow configuration uses configured Agent IP of 0.0.0.0 instead of Operational Agent IP of management IP (configured IP of the management port).
PD4-2355250081	After adding the trapreceiver <i>user</i> , it does not display after the <code>show snmpv3 user</code> command. Also the trapreceiver cannot be used as a RO or RW user even after given permissions to the required user.
PD4-2445640160	PSU state shows <code>Power_failed</code> for E4G-200 in Screenplay, although it shows as <code>Unsupported</code> in the CLI.
PD4-2501758416	Process route manager hits 99% CPU during link flap of LAG port in a PIM enabled VLAN.
PD4-2577130288	In E4G-400, when both XGM3S-2SF and XGM3-4SF are present, sometimes clock is not recovered via SyncE when the input clock source port and output clock ports are across the XGM modules.
PD4-2193216484	ERPS takes average of 94 ms to switch erps ring from idle to protection state (converge the traffic) when failover occurred in the ring but G.8032 recommendation is to provide 50ms protection and recovery switching for Ethernet traffic in a ring topology.
PD4-2415093851	If the CLI command <code>enable bgp export ospfv3 address-family ipv6-multicast</code> is given in <code>show conf bgp</code> it is shown incorrectly.
PD4-2495369831	Next-hop of the aggregate address is changed even though the routes filtered through aggregate policy has the same next-hop.
PD4-2479030541	Transmitted capabilities should not be shown when peering with neighbor that does not support capabilities.
PD4-2492841430	IPARP entries could not be learned as an identity on a shared port.
PD4-2473676475	When there are two VRRP instances with same VRID on two different VLANs, the switch is master for one VRRP instance and backup for another VRRP instance. Packets received with VRRP virtual mac on vlan-id for which switch is backup are consumed by the switch. The switch should not consume those packets.
PD4-2562996591	MIB value for OID <code>extremeInputPowerVoltage</code> wrongly shows incorrect or unknown for X460, E4G-400 and BlackDiamond 8806.
PD4-2769253950	The output of the command <code>show igmp</code> does not carry IGMP Snooping details. Workaround: Use <code>show igmp snooping</code> command instead.
PD4-2711224879	On DHCP servers when DHCP snooping is enabled, the server releases the IP to the client though it is not a trusted port.
PD4-2834808350	ESRP node doesn't L3 switch multi-cast traffic received on MLAG or non-MLAG ports when the ESRP node is slave for ingress VLAN and master for egress VLAN. Also, the ESRP slave doesn't switch L3 multi-cast traffic received on MLAG ports when it is slave for ingress and egress VLANs. As a result, traffic is being dropped when MLAG-PIMv4 is configured with ESRP.
PD4-2724480541	DHCP server does not work with PVLAN.
PD4-2860679997	512 BFD sessions are not stable (with default timers).
BlackDiamond 8800 Series Switches	
PD4-3108167929	Disabling IGMP snooping while PIMv6 is operating on IPv6 traffic causes the traffic to stop flowing across VLANs. Workaround: Do not disable IGMP snooping.
PD4-3405153751	When booting up, BlackDiamond 8900-10G8X-xl series switches go to failed state before coming to operational state and the following error message appears: "01/07/2013 04:29:43.68 <Error:HAL.Port.Error> MSM-A: Unable to configure DWDM channel to the transceiver for port 2:7 (-1) This issue is seeing only with v15_3_0_25 build. This issue is not seen with v15_2_2_7 build".



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-3291342581	While executing <code>run msm fail-over</code> command on BlackDiamond 8800 series switches, the following error messages appear: "11/23/2012 18:05:09.47 <Erro:Kern.Card.Error> Slot-9: exSmlpmcBitmapSetComplete: failed to set bitmap, group=39, unit=1, err=-7 11/23/2012 18:05:09.47 <Erro:Kern.Card.Error> Slot-9: exSmlpmcBitmapSetComplete: failed to set bitmap, group=74, unit=1, err=-7 01/16/2013 16:13:57.14 <Erro:HAL.SM.Error> MSM-A: aspenSmlpmcAddEgressPort: group does not exist 01/16/2013 16:13:57.14 <Erro:HAL.IPv6Mc.Error> MSM-A: SM failed to add"
PD4-3418269366	Process dcbgp ends unexpectedly with signal 11 while rebooting neighboring switches. "Process dcbgp pid 1737 died with signal 11 Code: 578e24 ac650004 sw a1,4(v1) 578e28 8c83001c lw v1,28(a0) 578e2c 8c850018 lw a1,24(a0) 578e30 <ac650000>sw a1,0(v1) 578e34 a0800031 sb zero,49(a0) 578e38 ac800018 sw zero,24(a0) 578e3c 03200008 jr t9 578e40 ac80001c sw zero,28(a0) 578e44 3c1c0012 lui gp,0x12 01/16/2013 12:42:10.86 <Crit:Kern.Alert> MSM-B: Process dcbgp pid 1737 died with signal 11 01/16/2013 12:42:11.18 <Crit:Kern.Alert"
PD4-3302111254	During longevity testing, OSPFv3 PID 5466 ends unexpectedly with signal 6 on BlackDiamond 8806 series switches.
PD4-3317420202	100FX (without phy) link is coming up connecting between BlackDiamond X8 series switches and Summit X440 stack. The following error message appears when you disable/enable the port, reboot and auto negotiate setting speed for BASET optics: "<Erro:Card.Error> Slot-6: aspendiags: Configuration error, attempting to set speed to 10 or 100 on SFP that does not support it on port 8" This error message does not occur on other all Summit platforms. This issue occurs on ExtremeXOS 15.2.2 and 15.3.
PD4-3335763779	Issue: 1 On BlackDiamond 8800 series switches with 1G48X-xl, Fx/Lx optics link goes down when a save and reboot and partner port (100FX without phy) is in the up state. Workaround: <pre>configure ports <ports list> auto on. configure ports <ports list> auto off speed 100 duplex full</pre> Issue: 2 On BlackDiamond 8800 series switches with 1G48X-xl, the following error message appears: "<Erro:HAL.Port.Error> MSM-A: Failed configuring speed for multi speed SFP on port 21" and link comes up at 100 Mbps speed when trying to set speed at 100 Mbps on Fx/Lx. This issue appears one time out of eight. This issue does not occur in ExtremeXOS 15.2.2.7.
PD4-3338131841	Process BGP consumes 99% of the CPU after restarting process BGP.
PD4-3337936364	PIM IPv6 Scaling: Packet loss occurs when BlackDiamond 8900-G48X-xl series switches have to forward traffic for 3000 (S;G)s. The total number of (S;G)s found in cache is 6000. No packet loss occurs when BlackDiamond 8900-G48X-xl series switches have to forward traffic for 3000 (S;G)s and the total number of (S;G)s found in cache is 3000.
PD4-3286134421	After restarting ports, the following errors occur in MSM-B: <Erro:HAL.SM.Error> MSM-B: aspenSmlpmcAddEgressPort: group does not exist 11/21/2012 16:59:53.70 <Erro:HAL.IPv4Mc.Error> MSM-B: SM failed to add 2:48 to IPMC 0, rv=-1 11/21/2012 16:59:53.70



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-3302111272	BlackDiamond 8806 rebooted with an following error on the CLI: rtMgrClientTransHandleSendPacket: failure ipmlSend ret=-112
PD4-3211693720	<Erro:HAL.IPV4ACL.Error> error message generated when clearing counters with ClearFlow enabled.
PD4-3211693711	ClearFlow "Last Value" counter is not resetting to zero after the defined sample period.
PD4-3291342588	DC-BGP signal 11 ends unexpectedly after executing disable/enable OSPF followed by MSM failover.
PD4-3331887655	Process dcbgp ends unexpectedly with signal 11, when disabling BGP with 512 peers (8 ports connected to IXIA, each with 64 VALN interfaces), and thousands of routes.
PD4-3103047401	Process VLAN ends unexpectedly with signal 11 on BlackDiamond 8800 series switches, while trying to send traffic through multiple tunnels (10 tunnels). Signal end occurs soon after initiating traffic from IXIA to multiple tunnels.
PD4-3121879961	BGP_GracefulRestart_NewStack: After an unplanned BGP process restart using the command <code>terminate process bgp graceful</code> , trace.bgp should be stored in the memory card if it exists. However, the trace .bgp is not stored in the memory card of the switch; instead the file is stored in the internal memory. This issue applies to ExtremeXOS 15.2.1 and later versions.
PD4-2836155071	On BlackDiamond 8800 and Summit X480 series switches, IPv6 automatic tunnels require hardware aging support when IPv6 uses external TCAM.
PD4-2820557781	Under certain conditions with VRRP, HAL process ends unexpectedly with signal 6: glibc detectedProcess hal pid 1277 ends unexpectedly with signal 6 and signal 11 (two different back trace).
PD4-2790183670	Graceful restart of BGP (in unplanned or both mode) does not work correctly with VPN VRF.
PD4-2840571161	License does not get applied if Backup MSM is not in synch but the CLI displays a messages saying that it was successfully applied. Workaround: Ensure that MSMs are in synch prior to performing a license upgrade.
PD4-2252055101	After disabling a 10G port on an 8900-10G24X-c module, entries are not flushed from the software. NOTE: This issue occurs on load-shared ports and can be reproduced by running the <code>disable port all</code> command.
PD4-2166404788	Some software forwarding of untagged VMAN traffic does not work when an 8900-MSM128 module is installed. That is, the software forwarding that results when limit-learning is configured does not work. A similar problem occurs with untagged VMAN traffic received from one of the following I/O modules regardless of the MSM installed: • 8900-G48T-xl • 8900-G48X-xl • 8900-G96Tc • 8900-10G8X-xl • 10G8Xc • 10G4Xc This problem is also seen in ExtremeXOS 12.3, 12.4, and 12.5 software releases.
PD4-2049460356	Only partial traffic is forwarded between a service VLAN and VPLS (LACP). Transmitting from a VLAN to VPLS, traffic loss is approximately 10%; from VPLS to a VLAN, traffic loss is approximately 98%.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-1633677741	On a BlackDiamond 8800 series switch, making link state changes during a large policy refresh can take more than 20 seconds and may cause duplicate packet forwarding in an MLAG configuration.
PD4-1546542587	ISIS process crashes with signal 6 while trying to change the metric-style to wide under scaled conditions.
PD4-1637972971	Beginning with ExtremeXOS 12.5, the mirroring feature stops working after downgrading and then upgrading the switch software.
PD4-1567438997	When an ExtremeXOS switch receives an OSPF user group LSA advertisement with a router ID field as 0.0.0.0, the switch does not process the advertisement and reboots OSPF.
PD4-1674379381	When installing new PSU controller firmware, log messages starting with <code><Crit:Kern.Critical></code> or <code><Erro:Kern.Error></code> may be generated by the backup MSM and can be ignored.
PD4-1530729359	When there are a large number of OSPF routes (> 100K) and the switch is restarted, OSPF session goes down and stays in EX_START and continues flapping to EXCHANGE and EX_START states.
PD4-750014887	If a failover occurs during a “refresh policy” the HAL process dies on a new master MSM. Workaround: Avoid performing a policy refresh if switching from one MSM to another.
BlackDiamond X8 Series Switches	
PD4-3417486171	Cannot manage BlackDiamond X8 series switches using XML. ScreenPlay does not start.
PD4-3397742865	On BlackDiamond X8 10G48X modules, link fails to come up after changing a port from a 1000BaseT SFP optic to a 10GBaseX SFP+ optic.
PD4-3381730844	The following logs appear when disabling tracking VLAN in the current master: “<Erro:Kern.IPv6Mc.Error> Slot-1: Unable to Add IPmc sender entry s,G,v=0000:0000:0000:0000:0000:0000:0000:0000,ff02:0000:0000:0000:0000:0001:ff00:0201,4093 IPMC 1 flags 92 unit 1, Entry exists” This issue does not occur with ExtremeXOS v15_3_0_20.
PD4-3282517920	CFM_MIB - snmp get of dot1agCfmMaNetName failed for one display parameter. Expected result is “Hex-STRING: 00 00 11 FF FF FF FF” but got “Hex-STRING: 22 11 00 FF FF FF FF”.
PD4-3252433587	I/O card becomes operational when all FM cards are in the “SHUTDOWN” state.
PD4-3288486281	Verifying the best route fails to show IGP routes in BlackDiamond X8 series switches with 5-node topology.
PD4-3202624621	MSDP SAs received from an eBGP neighbor are rejected.
PD4-3259170481	PIM-SM is configured on receiver VLAN; PIM-DM is configured on source VLAN. in this state when you change PIM mode from “dense” to “sparse” on the source VLAN, IPv6 multicast traffic is duplicated (doubled). Egress port is programmed for source and destination VLAN. it should be programmed for source VLAN only.
PD4-3282121167	During EAPS/STP fail-over, the following error messages appear (it seems for some IPv4/IPv6 multicast groups egress port does not get added in MM-B.): “<Erro:HAL.SM.Error> MM-B: esfmEntityAddEgressPort: entity does not exist 12/13/2012 16:45:19.26 <Erro:HAL.SM.Error> MM-B: esfmEntityAddIngressPort: entity does not exist entity=60 12/13/2012 16:45:19.26 <Erro:HAL.SM.Error> MM-B: esfmEntityAddIngressPort: entity same error messages are observed while doing disable/enable slot.”
PD4-3304371204	OSPFv3 neighbor between BlackDiamond X8 series switch and Summit X670 series switch stuck in EXCHANGE/DR state when you disable/enable OSPFv3 in NON-DR.
PD4-3241154791	On BlackDiamond X8 series switches, MSM modules stop working after enabling/disabling all ports.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-3280042531	Using ESPR priority-ports-track-mac algorithm, BlackDiamond X8 series switches running ExtremeXOS 15.2.2 have dual slave state. This is not seen in Summit switches.
PD4-3322786987	System performance is very slow when sending traffic to 256 VRRP interfaces.
PD4-3268227287	VRRP IP route tracking for IPv4 routes does not seem to work with 32 ECMP default static routes.
PD4-3172867544	After a process ends unexpectedly or a management module failover, the status stays in the "running" state and you have to unconfigure the switch.
PD4-3157840924	For BlackDiamond X8 series switches, the following error message appears "Error:Kern.Error & Erro:HAL.Port.ReadOpticDDMIDataFail" while executing command <code>show ports tr inf detail</code> on unsupported DDMI optics module. Part and serial number are detected properly after: <Erro:HAL.Card.Error> MM-A: aspenCardPortGetTransceiverInfoSfp: Unable to read DDMI info from Transceiver for port 1:18. <Erro:HAL.Port.ReadOpticDDMIDataFail> MM-A: Reading of DDMI data failed on the optical module in port 1:18. Workaround: Save and reboot, and then disable/enable slot in switch.
PD4-2787569240	OSPFv3 adjacency state changes after doing MM failover.
PD4-2968093558	If your configuration concurrently pushes the limits of both MSDP and (s,g) entries, you may experience undefined behavior.
PD4-2749874087	Do not attempt to upgrade a BlackDiamond X8 switch using the hitless failover process if newer BootROM or FPGA code has been installed as part of the installation of the newer ExtremeXOS image. Use the standard method of installing a core image show in the <i>Concepts Guide</i> , which requires rebooting the entire switch immediately after the installation.
PD4-2426499671	sFlow sampling is not working after running the disable slot and enable slot commands on a BDXA-10G48X module.
PD4-2486716171	The console window for the MM BootROM is limited to 80 columns × 25 rows. Using a larger size results in displayed text overwriting the screen. Larger console sizes work appropriately in ExtremeXOS.
PD4-2464885298	When we schedule a reboot time or cancel a reboot time, configuration allows changes on MM-A and B independently when it is actually a global setting.
PD4-2560208473	Uninstalling SSH XMOD fails on MM-A with <code>Device or resource busy</code> error.
Summit Series Switches	
PD4-3249665236	Need to automatically adjust for ingress and egress time stamp latency. This problem causes the latency between switches to seem longer than it actually is, which can keep gTP from working (specifically from becoming asCapable) on certain links, particularly BASE-T links. Workaround: Increase the 802.1AS parameter <code>neighborPropDelayThresh</code> on both switches connected by the link (on Extreme switches: <code>configure network-clock gtp ports <port_list> peer-delay neighbor-thresh <neighbor_thresh_time></code>).
PD4-3395770688	In some cases, on Summit X440 switches, OpenFlow sends empty packet_ins, eventually causing the switch to stop working and fail. Workaround: Reboot the switch and re-run the packet_in scenario.
PD4-3314306502	On Summit X450a/e switches with XGM2-2bt modules, normal/extended diagnostics fail.
PD4-3347131581	On Summit X440 switches, OpenFlow needs support to respond with ALL_TABLES_FULL.
PD4-3346145041	On Summit X440 switches, OpenFlow action to output special port controller does not work when used with a multiple action.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-3247019651	ERPS process ends unexpectedly with signal 11 after issuing the command <code>delete erps submetro</code> .
PD4-3337790808	Summit X670 series switches produce error messages for route manager: " <code><Error:RtMgr.Error> Gateway Gw vrid=2 Nh=2.5.1.2:iNh=3.5.3.2 sa=U flg=0x6 olfldx=0xf4270 inLabel=0x0 Type=IP outIf=v146_2 NHLFEIns=0x0 exfibShim=0x0 refcnt=1 (hdl=0x6fe658)</code> still present in hash table after removal"
PD4-3320438261	MVRP-created dynamic VLANs [SYS_VLAN] do not get deleted after disabling MVRP.
PD4-3199448140	The command <code>show protocol</code> does not show the difference in system-created/user-created protocol filter.
PD4-3199448131	Help option does not provide hex value format and range for Etypes to configure protocol filter.
PD4-3290151517	CFM_MIB - <code>snmp get on dot1agCfmMepXconCcmLastFailure</code> failed in NWI-E450A/x670 switch by displaying parameters in show logs output as "Hex-STRING: 00 --". Expected output is "Hex.*01 80 C2 00 00 30".
PD4-3288760818	For Summit X440 series switches, the command <code>configure accesslist with enable load sharing</code> failed with "Error: ACL install operation failed - slice hardware full for vlan".
PD4-3347099575	Disable OpenFlow VLAN with Active Rules.
PD4-3335987372	OpenFlow process ends unexpectedly with signal 6 after disabling/enabling OpenFlow.
PD4-3342351731	Change OpenFlow configuration to be VLAN-based, rather than port-based.
PD4-3321289423	Hung ACL while sending down two rules.
PD4-3340786971	Dynamically created VLANs get deleted on STP topology change.
PD4-3274999981	Configuring MVRP on a switch running EAPS and STP may cause an STP loop on the STP domain. Workaround: Configure <code>configure mvrp tag 3 ports 12 transmit off</code> for all EAPS-protected VLANs on MVRP-enabled ports.
PD4-3334236627	On Summit X440-24t series switches, the management port is flapping after rebooting the peer.
PD4-3174040301	When a mirroring instance has an ingress port filter and the VLAN to which this port belongs is added as a VLAN filter to another mirroring instance, traffic is mirrored to the VLAN filter only.
PD4-3338669767	On Summit X670 series switches, when a FlowMod with action <code>SetVlanPcp</code> is sent, only output works.
PD4-3314915463	AVB traffic fails on user-VR.
PD4-3330660401	When ESRP and ELRP are enabled with all the ports added to the master VLAN, the master switch stops working after receiving a <code>show vlan MasterVlan</code> command. The state in ESRP master switch switches between master and slave for some domains.
PD4-3336698507	On Summit X480 series switches, pings to neighboring router loopback addresses fail when no source is mentioned in the ping command.
PD4-3280042721	LAG configuration is lost on Summit X670 switches after rebooting. LAG is configured using combo [2:11] and non-combo port [2:46]. Before enabling LAG, these ports are configured with 1G speed.
PD4-3337647097	Forty percent traffic loss occurs for Summit X670 switches when in an IP route table, an IBGP route has been installed with a next hop that is not directly connected to the switch.
PD4-3346334931	The 10/100/1000 BASE-T optics link is flapping between unconfigured Summit X480-48x and X650-24x switches after several saves and reboots.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-3334236498	Ping is not successful on combo port of Summit X670v-48t switches and no optics specific.
PD4-3013252121	<Crit:VSM.ParmInv> messages occur when MLAG port is enabled/disabled.
PD4-3325786101	rtmgr process ends unexpectedly with signal 11 and the CLI <code>delete vlan ibgpv12</code> execution on Summit X450 series switches.
PD4-3292518613	On Summit X670v-48t, issuing command <code>config stacking-support stack-ports all selection alternate</code> , causes switch to crash: (kernel oops) - CPU 0 Unable to handle kernel paging request at virtual address 0000000c, epc == c4009008, ra == c402bcac Oops[#1]: This issue occurs with ExtremeXOS 15_2_2_7.
PD4-3298181451	idMgr signal 11 ends unexpectedly when creating 2 ldap domains and configuring domains for each.
PD4-3253270187	During automated testing, ripng process ends unexpectedly with signal 6: Listing: 60 isem->private ^ FUTEX_PRIVATE_FLAG); 61 62 /* Disable asynchronous cancellation. */ 63 __pthread_disable_asynccancel (oldtype); 64 65 if (err != 0 && err != -EWOULDBLOCK) 66{ 67 __set_errno (-err); 68 err = -1; 69 break;
PD4-3296877221	BGPbestpath_NewStack - RTMGR ends unexpectedly with signal 6. #0 0x2aae2630 in __pthread_rwlock_wrlock (rwlock=0x54a150) at pthread_rwlock_wrlock.c:75
PD4-3138196841	On Summit X460 switches, process dcbgp pid 1556 ends unexpectedly with signal 11 after issuing command <code>disable bgp</code> .
PD4-3187826591	For Summit X670-48t series switches, diagnostics version 5.10 has < 0.5% false failure in loopback phy and snake interface tests.
PD4-3159542451	For Summit X670v-48x series switches, ports with a SFP+_SR (SOURCEPHOTONICS) optic inserted with link up will flap if any other ports (two or more) are disabled/enabled. This issue does not occur with ExtremeXOS v15_1_2_12.
PD4-2488384169	In Summit X440 series switches, the maximum number of L3 interface supported is 256 only. Switch allows user to create more than 256 L3 VLANs and displays no error/warning message.
PD4-3012216202	Under certain conditions, in Summit X440 series switches, "kernel oops" error occurs during VRRP clean-up session. This issue occurs in the following versions ExtremeXOS 15.1.3.3, 15.1.3.1, 15.2.0.23, and 15.1.2.12.
PD4-2879629222	Under certain conditions, large number of packet drops occur when traffic is failed over to other active member ports in the LAG.
PD4-2821074451	In Summit X670V series switches, packet looping occurs with unknown L2 traffic type when 32-member ports are configured using custom address based LACP.
PD4-2994375866	In Summit X460 series switches, enabling load sharing is not allowed when ports in front panel ports (21 and 22) and XGM3-2sfpPort (49). The following error message appears: "Error: System cannot support Load Sharing among port media types with different maximum speed". This started in ExtremeXOS version v15.1.2.12.
PD4-3037333500	In Summit X670-48x series switches, only SFP+_SR (SOURCEPHOTONICS) ports are not coming up after restarting the port (25) and then save and reboot. This issue does not occur with Summit 650-24x and BlackDiamond 8800 series switches. Workaround: Disable, and then enable the far-end port.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-3064998712	In Summit X480 series switches with VIM3-40G4X and SummitStack, the following error message occurs after issuing <code>save</code> and <code>reboot</code> with custom LACP load sharing: “<Erro:HAL.Port.CfgTrunkFail> Failed to configure load sharing group 27 on slot 1 unit 0: Entry not found” Traffic does not appear to be affected. This problem does not occur with version ExtremeXOS 15.1.2.12.
PD4-2967046316	In Summit X460 series switches, IGMP join-and-leave delay is very long compare to other platforms.
PD4-2964780000	LAG does not forward traffic after restarting LAG ports and flushing FDB table.
PD4-2152137121	BFD sessions with minimal timers are not stable on Summit X460 and X670 switches.
PD4-2330727937, PD4-2120554424	Ports with BASE-T SFPs do not become active after disabling those ports and enabling them one by one.
PD4-2330727870, PD4-1600530241	A HAL crash with signal 11 occurs when connecting a ReachNxt device to a load-shared port. Workaround: Do not connect a ReachNxt device to a load-shared port.
PD4-1545964372	On a Summit X480 switch, the log message <code>Setting hwclock time to system time, and broadcasting time</code> is frequently displayed.
PD4-2094910917	An LDAP bind fails after a system reboot.
PD4-2026467027	When a Summit X670V-48x switch with a 40G port in partition mode connected to a 10G port using MTP breakout cables, the 10G side shows the link is up.
PD4-1950680298	On a Summit X670V-48x switch, inserting a copper GBIC (1000BASE-T) in a 10G capable port changes the speed and duplex setting of the port.
PD4-2092425173	When 500 virtual machines have 500 ingress or egress policies, running the <code>clear fdb</code> command results in removing some of the installed ACL entries from the hardware after receiving the following error message: <code>ACL refresh failed - updated policy has not taken effect.</code>
PD4-2093228351	On a Summit X670 series switch, FDB identities are not flushed from the identity management table.
PD4-2095792132	When clearing an FDB entry or updating a policy entry on VM-tracking, MAC addresses with a wide key policy installed results in the policy being uninstalled. When this occurs, the following error message is displayed: <code>ACL filter install failed on vlan *, port 1:16, rule "VM1:16_00:00:11:00:00:1a_E601" index 601, Invalid parameter (user-defined field (UDF)) "</code>
PD4-1824530443	The following behavior differences were observed between a Summit X480 and a Summit X450a switch: - Configured a general query interval of 125 seconds but a Summit X480 switch sends general queries every 130 seconds. This issue is not seen on a Summit X450a switch. - A Summit X480 is sending both MLDv1 and MLDv2 reports even though only MLDv2 is enabled. This issue is not seen on a Summit X450a switch.
PD4-1436226210	With default ethertype (0x88a8) configured, Summit X460 and X480 switches do not display an error message when adding a port as tagged to a VMAN when the port is already part of a tagged VLAN. When the port is already part of an EAPS control VLAN, EAPS goes to a failed state.
PD4-1721719301	The process HAL crashes with signal 11 after configuring a primary and secondary pseudo wire and sending unicast traffic through the pseudo wire.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-1676631313	When disabling sharing on a load-shared port that is part of multiple VLANs, VLAN statistics shows a "-" for some VLANs. Workaround: Unconfigure and reconfigure VLAN statistics.
PD4-1659644826	VLAN statistics monitoring is unconfigured for a specific VLAN if the VLAN name is changed.
PD4-1648023331	A Summit family switch configured with MLAG does not reply to the first ARP request received on an ISC port.
PD4-1664831900	VLAN statistics are included in the output of the <code>show configuration</code> command for load-sharing member ports even after the <code>unconfigure ports monitor vlan</code> command is issued. Workaround: Disable sharing, remove the VLAN statistics configuration, and enable sharing.
PD4-1676631313	When disabling sharing on a load-shared port that is part of multiple VLANs, VLAN statistics shows a "-" for some VLANs. Workaround: Unconfigure and reconfigure VLAN statistics.
PD4-1589959110	When adding ECMP routes using OSPF, a route flap occurs.
PD4-1603951551	On a Summit X460 stack, the following kernel warning is seen in the log during a failover. <Warn:Kern.Card.Warning> Slot-1: select_mux:line 250:I2C I/O operation failed
PD4-1590249340	Clearing FDB entries when a Summit X480 switch learns 512,000 MAC addresses from MLAG ports disrupts MLAG peer TCP sessions.
PD4-749682632	You cannot run the <code>configure port auto on</code> command on XGM2-2bt ports.
PD4-427423116	When a dot1x client is authenticated at multiple VLANs, the output of the <code>show netlogin port</code> command shows the client is sometimes authenticated at the local server and other times at a RADIUS server. NOTE: This occurs when dot1x and MAC authentication are enabled on the port.
PD4-1142692318	On Summit X480 switches, L3 multi-cast traffic sent from a service VLAN/VMAN to VPLS is not received at the VPLS peer. Workaround: When VPLS enabled VLANs exist on a port, all VLANs on that port must have IGMP snooping disabled. Also, the IP address cannot be configured on any other VLANs (including non-VPLS VLANs). Remove ports from the default VLAN.
PD4-448681226	The <code>show 12stats</code> command does not count ARP packets to the CPU, even though the packet goes to the CPU.
PD4-489142320	One Gigabit ports set to <code>auto on</code> flap twice during a switch reboot.
PD4-489359602	Conflicting Link Fault Signal (LFS) alarms are shown when disabling local ports.
PD4-274249122	If a Summit switch populated with an XGM2-2bt module is rebooted, a false link up is seen on 10G links connected to the XGM2-2bt ports approximately 30 to 50 seconds before the switch has fully booted.
PD3-43606168	If sFlow does not have a collector configured using the <code>configure sflow collector</code> command, the <code>show log</code> command generates the following messages: 08/23/2005 12:28:09.55 <Noti:sflow.debug.AddCntSmplFail> : Could not add the counter sample for port 0:1020, as receiver is not configured. 08/23/2005 12:07:49.55 <Noti:sflow.debug.AddCntSmplFail> : Previous message repeated 61 additional times in the last 1200 second(s).



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD3-40266236, PD3-40233121	Traffic on load share ports configured as redundant ports incorrectly moves to other ports in the load share group during link transition.
PD3-202013281	Learning is disabled by default on remote mirroring VLANs. Running the <code>enable learning</code> command on those VLANs may cause a loss of remote mirrored traffic.
PD3-202013298	The valid value range for tags on remote-mirroring VLANs is 1 to 4,094. Use these values for configuring the remote tag in the <code>enable mirroring</code> command.
PD4-2424777736	The following error message displays on Summit Stack on the new master node formed with E4G-400 and x460 during stack failover: <pre><Error:HAL.Port.Error> Slot-1:aspenCardPortLinkscanHandler(): port mapping failed for unit 1 port 1</pre>
SummitStack	
PD4-3298731387	Issue 1: In Summit X650-24x switch stacks, media type does not appear in command <code>show port configuration</code> output without any configuration. Issue 2: In Summit X670v/x670-48x switch stacks, media type does not appear in command <code>show port configuration</code> output after run failover without any configuration.
PD4-3422951171	Process <code>acl pid 1551</code> ends unexpectedly with signal 11: <pre>"#0 0x0043d9e8 in InstanceCompFunc0_1 (n1=<value optimized out>, n2=<value optimized out>) at acl_pol.c:3588 3588 if (inst1->vlanIfInstance != inst2->vlanIfInstance)"</pre>
PD4-3285096691	FIP snooping does not work after stack failover.
PD4-3303365072	On all four channels the optic is shown as "Q+LR4" when the stack is built/booted up for the first time. 40G QSFP+ SR4/LR4 link in the port by default 1x40G mode in Summit X650-24x stack. But it is expected to display only in channel 1. This issue does not occur in chassis and standalone switch. Workaround: Reboot the stack topology or Summit X650 slot.
PD4-3219427851	Stack port status is shown as "No Neighbor" when the stack link is up and properly configured in 320G stacking.
PD4-3310044391	traffic drop and following error messages are seen while removing/adding a port from/to LAG: <pre>"11/30/2012 11:27:09.46 <Warn:HAL.FDB.Warning> Slot-3: pibRemoveCPUFilter(1:25, 5) - VLAN 1011 not found in filter VLAN flood vector." "11/30/2012 11:27:09.47 <Warn:HAL.FDB.Warning> Slot-3: pibRemoveCPUFilter(1:26, 5) - VLAN 1011 not found in filter VLAN flood vector." "11/30/2012 11:27:28.41 <Warn:HAL.FDB.Warning> Slot-3: pibAddPortCPUFilter(2:18, 4, 00:e0:2b:00:00:04, 1010, 0) - MAC/STP/Flood mismatch for CPU filter." "11/30/2012 11:27:28.41 <Warn:HAL.FDB.Warning> Slot-3: pibAddPortCPUFilter(2:18, 5,"</pre>
PD4-3094883803	Kernel error occurs on Summit X440 series switches slot when enabling jumbo frames in stack.
PD4-3248963580	VRRPv3 IPv6 master election is not happening based on the link local address when priorities are same.
PD4-3271412898	Some IPv4/IPv6 multicast streams received on a Summit X670 stack are not sent out on egress ports. The egress port is 40G LAG and ingress port is a 10G LAG. mcast cache, HAL and hardware are programmed correctly, but traffic does not egress on 40G LAG ports.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-3334236528	On 10/100/1000 BASET optics link is coming up between Summit X440-8t and X460-24x V80-Stack. This is not observed in standalone Summit X460-24x switch. Also, BASET link is coming up in Summit X460-48x stack. Workaround: Remove stacking configuration and enabling stack.
PD4-3049363121	Inactive aging TTL set to value 45:12:13.
PD4-3295591451	Process ipSecurity pid 1588 signal 11 ends unexpectedly with cleanup session in IPAdresecurity.
PD4-3256376161	Process DCBGP ends unexpectedly with signal 6.
PD4-2894216251	"<Erro:pim.vsm.RtxTimerExpNoRspns> Slot-1: Ingress VLAN query timer expired for S 10.158.111.12, G 225.5.0.94 for ISC 1, max retries 5 over, no response received" observed continuously on console. If you disable BFD, these messages do not occur.
PD4-3068413902	In Summit X480 stacks under certain conditions, ExtremeXOS ends unexpectedly with "Process rtmgr pid 1502 died with signal 11".
PD4-3031740060	When rebooting any slot (master/backup/standby) in E4G-400 cell site router stacking, SyncE configuration gets deleted and clock is not forwarded through that ports in that particular slot.
PD4-3031740051	SyncE over stacking is not working after run failover when the clock source and output in are in different slots.
PD4-3036895393	Ports (1:21 & 1:22) in Summit X670v stack (AlterNative) ports do not go into down state even when the connected ports (through Tri Speed BASET link) in the other ports (6:21 & 6:22) DUT is disabled and re-enabled. Workaround: Ports go to down state after you save, and then reboot the switches in the stack.
PD4-2970543225	When Vim3-40G4X is configured in 4x10G mode in Summit X650-24X series switches, you cannot form alternate stacking using ports number 23 and 24. When the same card is configured in 1x40G mode, you can enable stacking.
PD4-2899595434	On Summit X450a series stacks with the L3 egress broadcast. Traffic is not successfully with mirrored in stack in enhanced mode. This issue occurs in the following version ExtremeXOS 12.6.1.1, 12.6.2.10, 12.6.3.1, 12.7, 15.1.1, 15.1.2, and 15.2.1.
PD4-2724480637	CFM_MIB : verification failed in SNMP walk while returning the correct values on dot1agCfmMepTable after MSM failover. This is failed in New Stack (v320).
PD4-2600920021	The counters of the CLI <show port stack-ports utilization> do not increment properly when different types of traffic streams are sent.
PD4-2131586058	A Summit X480 stack does not come up after a failover if the stack is configured with 50 identity roles.
PD4-2234586031	When using a SummitStack, known L2 traffic sent to port 64 is not forwarded in 4x10G mode.
PD4-1678164933	After upgrading to ExtremeXOS 12.5.1.4, the following error is shown in the log: <Erro:Kern.Ipv4Mc.Error> Slot-1: Unable to Del IPmc vlan 924 for 1:15 s,G=a9e6f05,e1010028 IPMC 186, unit 0 Entry not found.
PD4-1645865216	The following error message is seen when operating a stack where the backup switch does not have the same feature pack licenses as the master switch. Error: Backup execution timed out, command execution aborted! Workaround: When using VRs, be sure that while forming the stack using mixed platforms, the master node and backup node have the same licenses feature pack.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-928567091	Running the <code>synchronize</code> command on a Summit X650 in a SummitStack causes the system to time out and the stack to not synchronize for an extended period of time. This also results in the master node no longer being accessible.
PD3-181304741	After inserting a XENPAK in a stack (XGM2-2xn, XGM-2xn) and performing an <code>snmpwalk</code> on the <code>entityMib entPhysicalDescr</code> variable, XGM- is always shown, not the complete module description.
PD3-209191768	After running the <code>disable port all</code> command on a SummitStack, some port LEDs may sometimes light green even though ports are not up.
PD3-204744742	IPv6 neighbor-discovery in a management VLAN in a SummitStack resolves to the node address of the stack master, instead of the stack MAC address.
PD3-136493921	If a switch is added as a master-capable node of a stack whose master node has a license level that is not equal to the level of the switch, the switch will fail. The complete condition can be seen using the <code>show slot detail</code> command. In this state, the switch does not have AAA services available. You will only be able to log into the switch using the failsafe account that was last assigned to it. You must log into the switch to upgrade the license. If the switch is not using the failsafe account configured on the stack, you can use the <code>synchronize stacking {node-address <node-address> slot <slot-number>}</code> command to copy the failsafe account information from the master switch to the failed switch NVRAM. You can also use the <code>configure stacking license-level</code> command to configure a license level restriction on the entire stack and then reboot the stack. Once the stack is restarted, there is no longer a license mismatch, enabling you to log into the switch and upgrade the license. From the master switch, run the <code>unconfigure stacking license-level</code> command to get the stack to operate at the desired license and then reboot the stack.
E4G-200 and 400 Cell Site Routers	
PD4-3394733284	Ping fails through VPN after creating 125th VPN VRF.
PD4-3386857057	Hybrid Clock—E4G-400 cell routers do not synchronize to GM in hybrid mode sometimes.
PD4-3386231103	Process SNMP ends unexpectedly with signal 5 when rebooting the device. Device configuration file has four SNMPv3 target addresses added.
PD4-3234389983	Wintegra error occurs when creating a pseudowire with UDP ports same as that of existing PWE.
PD4-3210756739	ACL error occurs during E4G-200 reboot. Does not occur on ExtremeXOS 15.2.2.6. "10/19/2012 13:40:24.48 <Error:HAL.IPV4ACL.Error> EXOS application attempting to install incompatible ACL: slice vlan y1731_A1C1, port * (rule "dot1agSysMAC_1000340", index -2)"
PD4-3299329980	L3 VPN – After dynamically changing the route-distinguisher value in PE-1 node, remote PE-2 fails to advertise the updated VPNv4 routes as IPv4 routes to its connected CE.
PD4-3346216381	L3 VPN SOO-PE router fails to transmit the VPNv4 as IPv4 routes CE as soon as the back door link is disabled.
PD4-3298763670	E4G200-12x cell site routers stop working after PTP-OC is configured and running over a long period (greater than two days).
PD4-3240872343	FDB errors occur when initiating ESVT between E4Gs that has VPLS/VPWS configurations running.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
BGP	
PD4-3300432331	BGP is not preserving routes when GR is enabled. Issue also occurs in BGPv6.
PD4-3296418447	Router does not allow re-enabling OSPF/OSPFv3 routes into BGP. The following error message appears: "Error: Cannot change export policy for protocol ospf while export is enabled" You need to reboot the switch after issuing the command <code>enable bgp export ospf ipv4-unicast</code> to make this configuration.
PD4-3336005381	BGP_GracefulRestart_NewStack: Parameter "Policy for NLRI Type ipv6-unicast" in the command <code>show bgp neighbor <ipaddress></code> is not getting the expected established state in a specific configuration of "BGP GR from both/planned/unplanned to aware-only". This works fine except in ExtremeXOS 15.3.0.
PD4-3218758011	In BGP when executing the command <code>enable bgp peer-group p1 capability ipv4-unicast</code> , the following error appears: "Error:missing close-brace". And then the CLI does not work.
PD4-3227329371	Aggregate policy is not functioning in BGP.
PD4-3346945750	BGP_QAD_NewStack - DCBGP process ends unexpectedly with signal 11 on BlackDiamond X8 and Summit X670 series switches running ExtremeXOS 15.3.0.20: #0 0x08259f40 in bgp_rtmgr_agt_restart_timer_pop (vrf_cb=0x836a108, reg_cb=0x836e938) at bgp_rtmgr_reach.c:1071 1071 EMS_BGP_RtMgr_NextHopUnrch(vrf_cb->vr_id,
PD4-3138621068	BGP_QAD_NewStack: Issuing the command <code>enable bgp export ospf address-family ipv4-multicast</code> produces the error: "Error: Cannot change export policy for protocol ospf while export is enable". This issue occurs in ExtremeXOS 15.2.1 and 15.2.2, but not in 15.1.2.
PD4-2897496481 PD4-2972980627 PD4-2918099300 PD4-2925483199	In BGP, various processes end unexpectedly and switches fail occasionally under certain remote switch rebooting and/or disabling peers conditions.
PD4-2797497741	For some filtering situations there is no indication that incoming routes have been rejected by a policy filter.
PD4-2744063235	If an RD (or RT) of "23:1" is set, the standard (RFC 4364) is vague as to whether that should be a type 0 or type 2. An AS of "23" is a 2-byte AS, so it could be encoded as a type 0. However, it could also be a type 2.
PD4-2965165797	BGPv6_ceasenotification - Verifying Notification message ADMINISTRATIVELY SHUTDOWN is not showing the expected result for LastError Admin peer shutdown.
PD4-2933087364	BGP neighborship between IXIA to DUT fails when BGP router-id in DUT is configured between 248.X.X.X to 255.X.X.X range.
PD4-2861990096	Advertised routes counter in BGP displays wrong values after a DUT reboot with an out-bound route-policy configured.
PD4-2969922397	BGP_GracefulRestart_NewStack : Checking End-of-RIB Message with restart flag unset failed.
PD4-1599215746	Ping fails for remote loopback addresses.
PD4-1689076501	Rebooting a BlackDiamond 20800 series switch causes the system to crash in process HAL signal 11.
PD4-1891612811	BGP routes are not made inactive if IP forwarding is disabled.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-2092705405	A switch reports the following error when running the CLI command <code>configure bgp neighbor x.x.x.x no-route-reflector-client</code> if the peer is not an RR client: Error: RR Client Peer remote AS xxxxx does not match local AS xxxxx
PD4-2151105474	On a Summit X480 stack, a watchdog timer reboots when a backup slot is unconfigured or powered down when the stack has 524,000 BGP routes.
CLEAR-Flow	
PD4-278443631	CLEAR-Flow commands display on platforms that do not support this capability, including the Summit X150, X250, X350, and X450e series switches, as well as BlackDiamond 8800 non-c-series switches.
Connectivity Fault Management (CFM)	
PD4-2492531418	When doing the restart process <code>do1ag</code> the following message appears: "*** acl: transHandleSend sendPacket peerId = 31 ret = -1"
PD4-2562050001	ERPS does not move to <i>idle</i> state when we disable sharing because CFM does not learn some MEPs do not show any issues when sharing is enabled.
PD4-2415095641	Unicast CCM, CFM Ping and trace route do not work on BD10k and BD12k platforms from 15.1 onwards. Only the multi-cast CCM will work
EAPS	
PD4-2343129201, PD4-2038216594	If a port is included in an EAPS protected VLAN, it cannot be configured for redundancy.
PD4-1673032272, PD4-1676753651	EAPSV2 segment health-check packets received on a ring port may be dropped if the EAPS node on a Summit family or BlackDiamond 8800 series switch has a different EAPS shared port on any other ring ports.
PD4-749215481	Disabling the EAPS master primary port when there are no other ports configured on a protected VLAN will cause a disruption of L2/L3 multi-cast traffic. Workaround: Enable loopback on all EAPS protected VLANs.
PD4-471892924	Restarting the EAPS process on a controller generates the following error messages on a console, but does not impact switch performance. BD-8806.80 # restart process eaps Step 1: terminating process eaps gracefully ... Step 2: starting process eaps ... Restarted process eaps successfully BD-8806.81 # ERROR:VmgrProtocolIfRegister protoId:0 numIf:1 ERROR:VmgrProtocolIfRegister protoId:0 numIf:3 ERROR:VmgrProtocolIfRegister protoId:0 numIf:1
ERPS	
PD4-3349333182	ESRPv6 + DAD: Duplicate IPv6 address is detected though there is no duplicate address while disabling/enabling ports between ESRP master and L2 switch.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
IP Routing Protocols	
PD4-2905790321	The ip-option record-route is processed even while it is disabled in the intermediate switches.
PD4-2965217751 PD4-2965218003 PD4-2921267090 PD4-2925325781 PD4-2799443103 PD4-3041932280	In OSPF/RTMGR, various processes end unexpectedly and switches fail occasionally under certain conditions.
PD4-2914006121	OSPF router TLV disappears in the <code>ospf lsdb detail lstype opaque</code> after disable/enable port two to three times or pulling out the cable two to three times.
PD4-2996986009	Unable to configure authentication for VRRPv2.
PD4-2886857541	PBR: Reject the configuration if the flow redirect next hop is same as vlan IP address
PD3-39411271	icmpInMsgs counter will display the incoming ICMP packets for VR-Default only.
PD3-128093864	MSDP Source-Active responses received from non-RPF peers are not processed.
PD3-192821161	For Summit X650, X450 a-series and e-series switches, and the BlackDiamond 8800 series of switches, the maximum number of supported IP interfaces is 512 (IPv4 and IPv6 combined). If there are more IP interfaces configured, the following log message is displayed: <Info:HAL.VLAN.Info> Maximum # of IP interfaces (512) already configured. Could not add IP address 0x0A010101 mask 0xFFFFFFFF00
PD3-202580681	Enabling IP route compression may cause temporary slow path forwarding for some of the L3 traffic.
PD4-718946965	Directed broadcast traffic is not being forwarded.
Mirroring	
PD4-2804878767	Dynamic ACL counter is not incremented properly. Packets are double counted. (Packets sent=10000, counted=20000)
PD3-79867211	If you create a load sharing group (trunk), then enable mirroring to a port, the software allows you to add the mirroring port to the load sharing group.
MPLS	
PD4-3240664297	The command <code>show iproute mpls</code> shows only seven routes from MPLS, when there are eight static LSPs configured.
PD4-1992679531	In VPLS, when configuring 7,190 pseudo-wires, FDB entries are not learned on 100 pseudo-wires.
PD4-1592270405	The run <code>msm-failover</code> command shows the following warning message in the log. <Warn:Kern.IPv4FIB.Warning> Slot-4: dest 0x0A9E6D7C / 30 nexthop 0x0A9E6D39: Unable to add route to unit 1, rc Entry exists. Shadow problem.
PD3-93218551	If either an egress or a transit LSP traverses the system, and an MPLS labelled packet containing a router alert label is received, that packet is not forwarded.
PD3-203917264	When an explicit route object (ERO) is changed for an LSP session that is up, the LSP that is already up is not torn down. LSP stays up based on the older values. The retry count continues to increment as LSP tries to come up with new values by querying routes every 30 seconds. This is done while the earlier LSP session is still active using the previously configured values. See the retry count in the command output for the <code>show mpls rsvp-te lsp <lsp_name> detail</code> command.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
Multi-cast	
PD4-339945634	When a load-sharing group is a member of a mirrored VLAN, packets ingressing on the member of the load-sharing group in the mirrored VLAN should be mirrored. On the Summit family switches and BlackDiamond 8800 modules, packets ingressing on member ports other than the master port of the load-sharing group in the VLAN are not mirrored. Workaround: Packets ingressing non-master ports in the load sharing group on the mirrored VLAN can be mirrored by adding virtual port mirroring filters for each of the non-master member ports.
PD3-78144711	The <code>show ipstats</code> command does not increment IGMPv3 statistics.
PD3-79383551	IGMPv3 Report Record type "5" does not work as expected when sent after a type "2" or a type "4" message.
Network Login	
PD4-763062511	Hitless upgrade is not supported for network login in ExtremeXOS 12.3.1.
PD4-752731351	You should not be able to enable network login if a VLAN is a VLAN-aggregation subVLAN. The system should generate a syntax error.
PD4-2930611052	Unable to enable netlogin in user-vr.
Network Services	
PD3-93829391	Configurations using a VR-Mgmt interface as a RADIUS client IP may not load at boot-up. However, using an interface in VR-Default will load correctly.
PD3-67727590	Creating two sets of VMAN ACLs with 4000 entries each and performing a VMAN ID translation on each ACL may generate the following error: <pre>03/15/2006 17:57:28.84 <Info:pm.config.openingFile> MSM-B: Loading policy RLL20k from file /config/RLL20k.pol ...03/15/2006 17:57:32.46 <Info:pm.config.loaded> MSM-B: Loaded Policy: RLL20k number of entries 4002Error in alloc txmi txmi 0x9f2 txmdi 0xffffffff Error in alloc txmi txmi 0x9f4 txmdi 0xffffffff Error in alloc txmi txmi 0x1102 txmdi 0xffffffff Error in alloc txmi txmi 0x9f6 txmdi 0xffffffff Error in alloc txmi txmi 0x9f8 txmdi 0xffffffff </pre>
OSPF	
PD4-3288760782	The command <code>show iproute</code> after exporting all ospf routes fails to show the required routes while verifying <code>enable rip export ospf intra area route</code> .
PD4-3288760735	The command <code>show ospf lsdb detail area with lstype as-external</code> failed to show up the required routes while verifying export static routes on ASBR.
PD4-3296418465	Router is not re-advertising OSPFv3 routes with new link-local address when link-local address is changed.
PD4-3206618829	OSPFV3 is not selecting highest area ID to forward traffic (to ASBR).
PD4-3135276301	OSPF is not properly installing next best hop for E1 routes.
PD4-2271653883	Using a specific configuration, OSPF neighbors may become a designated router even though the switch only has an Advanced Edge license.
PD4-2078865382, PD4-1493257018	The forwarding address in an external LSA should not be set for an interface that is configured as passive.
PD4-1641495299	When 5,000 routes are received via the OSPF neighbor and advertised to 253 neighboring OSPF routers, all 253 sessions go down. It then takes nearly 20 minutes for the sessions to come back up.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-1548969848	OSPF neighbors remain in the Exchange state after disabling and enabling an OSPF instance.
PD4-2148345931	Process ospfv3 pid 1527 died with signal 6 crash after issuing the command <code>unconfiguring ospfv3</code> then trying the <code>show configuration ospfv3</code> command.
PD4-2100715451	OSPFv3 neighbor-ship goes to INIT state once after withdrawing ospfv3-external routes when traffic is flowing.
PD4-2434344631	While flapping OSPFv3 Virtual link x670 crashed with the following message: "ospfv3: ospf6_lsProcess ospfv3 pid 1588 died with signal 6"
PD4-2430322044	Unable to delete an OSPFv3 area which was configured for virtual link.
PD4-2110841706	OSPFv3 neighbors keeps flapping once after disabling and enabling IP forwarding IPv6 with default timers.
QoS	
PD3-67431351	Configuring an ingress traffic queue and an egress traffic queue association to multiple ports in sequential order generates the following error: Egress queue already associated to this ingress queue Configuration failed on backup MSM, command execution aborted!
PD3-16578296	The member ports of a trunk will retain the QoS profile configuration of the trunk (based on the master port) after load sharing is disabled, or if a port is removed from the trunk.
RMON	
PD3-12950492	Issuing the <code>clear counter</code> command might cause a high number to be displayed in variables such as <code>etherHistoryOctets</code> , <code>etherHistoryPkts</code> , and <code>etherHistoryTable</code> .
ScreenPlay	
PD3-111344472	ScreenPlay allows you to configure DHCP but you cannot enable DHCP.
Security	
PD4-3248608611	After configuring vm-tracking server repository, configuring the refresh interval is not possible. You have to unconfigure the existing configured server and re-configure it.
PD4-3318405524	IdMgr queries for <code>userPrincipalName=user@domain*</code> may not work if the UPN suffix is different from the domain name.
PD4-3293652967	Radius authentication fails with move-fail-action authenticate configuration in BlackDiamond X8, Summit X670v, and all SummitStack. This issue is not observed in standalone Summit X670v and BlackDiamond X8 series switches. This issue occurs in ExtremeXOS 15.3.0, 15.2.2.7, 15.1.3.3, and 15.1.2.9.
PD4-3294445945	ACL signal 6 ends unexpectedly when there are stale entries in the command <code>show access-list dynamic</code> .
PD4-2955462615	IDM: After MSM failover Netlogin table does not show correct users on enabling netlogin dot1x. (Dot1x user is shown in Netlogin table in master and MAC user was shown in netlogin table in backup.)
PD4-2900851991	IDMPolicyOrder: Adding policies before creating child roles is not generating an error on adding child roles with policies that are the same as parent and because of this. On changing order: Error: Policy "policy6" added already in the role's hierarchy of role "role1". Because of this, re-ordering is not occurring.
PD4-2743463189 PD4-2743526221	IDM Multiple Windows Domain: For show identity-management entries, user does not show entries for Kerberos user (only dot1x and unknown users are shown). Show IDM entries domain does not show entries for all the domains in IDM entries table.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD3-205012219	The source IP lockdown dynamic deny ACL counter is not working properly and increments valid traffic from a trusted client.
PD3-186939931	Ingress mirroring is not working for DHCP snooping when snooping is enabled on BlackDiamond 12800 series switches. DHCP snooping works correctly when DHCP snooping is disabled.
PD3-75120608	The <code>unconfigure radius</code> and <code>unconfigure tacacs</code> commands do not reset the timeout value to the system default of 3 seconds.
SNMP	
PD4-3346123370	On Summit X440 series switches, snmpMaster process ends unexpectedly with signal 1.
PD4-2251137003	The following information will be added to the ExtremeXOS 15.0 version of the <i>ExtremeXOS 15.1 Concepts Guide</i>: The SNMP context <code>Name</code> should be set to the name of the virtual router for which the information is requested. If the context <code>Name</code> is not set, the switch retrieves the information for <code>VR-Default</code>. If the SNMP request is targeted for a protocol running per a virtual router, then the context <code>Name</code> should be set to the exact virtual router for which the information is requested. Refer to the “Adding Routing Protocols to a Virtual Router” section in the <i>ExtremeXOS Concepts Guide</i>. Following is a list of protocols that run on a virtual router: • BGP • OSPF • PIM • RIP • OSPFv3 • RIPng • MPLS • ISIS
PD4-1388191921	When changing an SNMP master configuration using <code>SNMP set</code>, the changes are not immediately reflected in the <code>show configuration snmp</code> command output. Workaround: Run the <code>save configuration</code> command to see the changed configuration in the <code>show configuration snmp</code> output.
PD4-705730556	AES/3des users created using ExtremeXOS 12.3.1 software cannot be used for SNMP operations in ExtremeXOS 12.1 or earlier releases. This may cause the SNMP master to crash.
PD4-2562996591	The SNMP MIB OID <code>extremeInputPowerVoltage</code> response is not correct for BD8806 having 48VDC power supplies in place and the response on X460, E4G-400 is unknown.
Spanning Tree Protocol	
PD3-189927343	A temporary loop occurs when a root bridge is taken down by disabling all ports or powering down the switch.
UPM	
PD4-1664927541	UPM profiles for events <code>identity-detect</code> and <code>identity-undetected</code> are not executed when many unique kerberos users login simultaneously from two client PCs. This happens when 50 unique users login continuously from PC1, and another 50 unique users login continuously from PC2 at the same time.



Table 2: Open Issues, Platform-Specific and Feature PDs (Continued)

PD Number	Description
VLAN	
PD4-3353962932	PVLANS can be configured as EAPS control VLANs.
PD4-3353869081	In EAPS with PVLAN, when FDB is cleared, traffic b/w non-isolated ports stops in one direction. Workaround: Stop the traffic, clear FDB, and then re-start the traffic.
WAN PHY	
PD3-101226461	When show wan-phy commands are run on non WAN PHY ports, the ports display the headers. It should only display the error wan command is not supported on non-wanphy port 25.



Known Behaviors

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

Table 3: Known Behaviors, Platform-Specific and Feature PDs

PD Number	Description
General	
PD4-468366251	A network login client is not authenticated if the username is 32 characters. Only 31 character user names are supported, even if the user can create a 32-character username.
PD4-3315488109	After VLAN tag is modified on PMBR, traffic is not forwarded through dense circuit. Workaround: Enable MLD snooping on all interfaces.
PD4-3251586520	The IGMP groups are not deleted from the MVR VLAN when leave packets are sent in the edge VLAN. Workaround: Enable proxy when MVR is configured.
PD4-3352005341	Ping from ESRP master switch to host connected to ESRP slave (connected via host attach port) fails when the switch to which the host is connected is master for another ESRP domain.
PD4-3206377171	After a policy attached as a MVR static address range is modified, the ports for all MVR VLAN groups are deleted, while the addresses that are permitted by the policy are not deleted.
PD4-3320002080	Sending vr-value as vr1 with vlan-tag from vm-map file with no user vr configured in the switch results in no vlan-tag and vr-name creation for that VM, and then creating user vr vr1 and running repository synch does not create vm with vlan-tag and vr vr1 as configured in vm-map files. Workaround: Restart process vmt and do run vmt repository synch.
PD4-3154851418	When both dynamic VLAN uplink port is configured and ISC port is configured dynamic VLAN created, add the uplink port as tagged, but the flag is only *1lgV (should be with both "U" and "V" flag).
PD4-3269060537	The command <code>clear fdb</code> does not zero dropped counter in <code>show fdb</code> command.
PD4-3215358251	The help text (by pressing TAB) does not appear after issuing the command <code>configure vrrp vrrpvlan vrid 1 authentication none</code> .
PD4-3075918118	VM Statistics - Egress counter is not working.
PD4-3098599761	Changing from untagged to tagged traffic with the same VLAN tag does not add the ports as "tagged" and this results in traffic loss.
PD4-3080980181	When STP is configured and enabled the dynamic VLAN propagation suddenly vanishes.
PD4-3143967884	After authenticating two computers using the dot1x authentication method, a Remote Desktop session between two computers fails. Workaround: Enable computer-only authentication using the procedure from the Microsoft Windows knowledge base: http://support.microsoft.com/kb/929847 .
PD4-2039102228	There is no CLI command to verify which control and protected VLANs belongs to which ERPS domains. Also missing are flags reserved for ERPS in <code>show vlan</code> command.
PD4-2744727326	<code>clear bgp</code> or <code>show bgp</code> commands should fail in VRF context.
PD4-2770013534	Valid MLD messages with unspecified address (::), are treated as valid MLD packets. instead, these should be silently dropped.



Table 3: Known Behaviors, Platform-Specific and Feature PDs (Continued)

PD Number	Description
PD4-2770013437	MLD protocol Timer value is not inherited from the current querier's query message— instead it is always the configured value on this router.
PD4-2835155421	On a private VLAN [network VLAN] translated port, the packets egressing when captured, shows the subscriber VLAN tag instead of Network VLAN.
PD3-57182431	For the incoming traffic with alignment errors, the "RX Align" counter in the output of the <code>show ports <port number> rxerrors</code> command is not incremented. Instead, the "RX CRC" counter is incremented.
PD4-1663984367	Whenever a (*, G) join is received, all hardware entries installed on non upstream (*, G) interfaces are cleared. Therefore, every 60 seconds, the L2 switching is affected, traffic comes to the CPU, and entries are re-learned.
PD4-2110742669	When using SCP to transfer files to an Extreme switch, the transfer fails with an "incomplete command" error.
BlackDiamond Series Switches	
PD4-3266706442	The command <code>clear counters</code> followed by <code>show port congestion</code> results in an error if executed within about eight seconds of each other on BlackDiamond 8800 series switches.
PD4-3155474589	In BlackDiamond 8800 series switches, port flapping can occur when enabling mirroring instance with loopback port configured.
PD4-2938691821	For the BlackDiamond X8 series switches, differv value is not getting replace when dot1p examination is enabled on ingress port with diffserv replacement enabled for slow path I3 traffic.
PD4-2339271510, PD4-2180251961	For the BlackDiamond 8800 series switches, when running the <code>show tech</code> command on a backup MSM, an error message is displayed.
PD4-2854254274	For the BlackDiamond X8 and 8800 series switches, L3 traffic with IPv4 options or IPv6 ext. headers is not sent to the CPU when redirected by flow-redirection.
Summit Series Switches	
PD4-3349524091	On Summit X460/X440 series switches, slot 2 reboots when more than the maximum permitted FDB entries are learned.
PD4-2913932450	QSFP+LR (ColorChip) optics link is coming up when connected back-to-back and after a save and reboot of ExtremeXOS fails to detect the media type when a Q+LR4 in ports <port lists> on Summit X650 and X480-24x series switches. Workaround: Only the media type is getting properly set after unplug/plug the optics from ports. This issue is not seen with QSFP+ SR4 on Summit X650 and X480 series switches.
PD4-2760140871	For Summit X670 series switches, dynamically learned FDB entries on non-FIP snooping VLANs disappear after FIP snooping is disabled and re-enabled on the VLAN.
PD4-2835588361	For Summit X670 series switches, the <code>show conf bgp</code> command is not vr-aware. VR-default BGP config appears even though inside vrf-a context. It should show BGP config for vrf-a
PD4-2857038040	For Summit X650 series switches, differv value is not getting replace when dot1p examination is enabled on ingress port with diffserv replacement enabled for slow path I3 traffic.
PD4-2857038031	For Summit X650 series switches, dot1p value changes to zero when both dot1p and diffser examination is enabled on ingress port.
PD4-1637091230	With 4,000 VPWS sessions, traffic recovery takes approximately 8 minutes before a port flap occurs. Workaround: On a Summit X460, it is recommended that you only configure 1,000 VPWS instances.



Table 3: Known Behaviors, Platform-Specific and Feature PDs (Continued)

PD Number	Description
E4G-200 and E4G-400 Cell Site Routers	
PD4-3316332889	ESVT fails to run even with just 10 ERPS rings configured in an E4G-200 cell site router node. ACL slice errors occur due to a hardware limitation.
PD4-3312009417	TDM UDP PWE fails to work when created between PE routers. Transmitted TDM packets are not received at the destination.
ACL	
PD4-2649674514	Policies with "redirect-port-list" as the action modifier do not get installed on a G48Te2 nor a 10G2Xc. It is not installed on either Summit X450a nor x450e.
PD4-2761666711	Redirect port list changes its behavior when sending slow path traffic.
PD4-1933402713, PD4-1933225935	The ACL action "copy-cpu-and-drop" is not copying EAPS control packets to the CPU.
PD3-77983510	Summit X450a and Summit X450e series switches and BlackDiamond 8800 a-series and e-series modules provide more powerful ACL capabilities. Because of this, the amount and complexity of ACL rules will naturally impact the time needed to process and apply the ACL rules to the switch. This will also impact switch bootup time. Access Control List limitations fall into two areas: physical and virtual. Physical Limits—Summit X450a and Summit X450e series switches: The per-VLAN, wildcard (port any), and single-port access list installation limitations are 1,024 rules for the Summit X450e and 2048 rules for the Summit X450a. Physical Limits—BlackDiamond 8800 a-series and e-series modules: The per-VLAN, wildcard (port any), and single-port access list installation limitations are 1,024 rules for the e-series modules, and 2048 rules for the a-series modules. Extreme Networks recommends that you configure ACLs as per-VLAN, wildcard, or single-port. If either of the following is true, you will have to configure ACLs with multi-port lists: Your application requires that ports do not have a homogeneous ACL policy. When BlackDiamond 8800 original series modules are operational in the same chassis, it may be necessary to configure ACLs to specific port-lists instead of as wildcard or per-VLAN. This is because the original series modules have smaller physical limits. Virtual Limits—Summit X450a and Summit X450e series switches: When configuring a multi-port ACL, use the following guideline. The total ACL count (as calculated by ACL rules times ports applied to) should not exceed 48,000 total ACL rules. For example, applying a 1,000 rule policy file to a 48 port multi-port list is supported (1,000 rules * 48 ports in the list <= 48,000). Virtual Limits—BlackDiamond 8800 a-series and e-series modules: When configuring a multi-port ACL, use the following guideline. For any a-series or e-series blade in the system, its total ACL count (as calculated by ACL rules times ports applied to) should not exceed 48,000 total ACL rules. For example, applying a 1,000 rule policy file to a 48 port multi-port list on an a-series module on slot 1 and an e-series module in slot 2 is fine. Neither module exceeds the 48,000 total ACL rules. Excessive boot times and CPU resource starvation can be seen with larger total rule counts. If your application requires additional capacity, contact Extreme Networks.



Table 3: Known Behaviors, Platform-Specific and Feature PDs (Continued)

PD Number	Description
BGP	
PD4-2216087479	A confederation ID is used as an aggregator ID when in a confederation instead of an AS-number.
PD4-2125200453	A backup slot does not come up when rebooted with 1,000 non-unique routes on a Summit X480 stack.
MPLS	
PD4-521915271	The Internet Group Management Protocol (IGMP) group reports may occasionally change from Version 2 to Version 3.
PD4-581950231	Multi-cast traffic is not received even though the rendezvous point (RP) tree and source information is shown in the PIM cache table
PD4-475414505	In more complex topologies, detour Label Switched Path (LSP) connections are not set up.
PD4-475414370	The following warning message is seen numerous times after changing VLAN Virtual Private LAN Services (VPLS) mappings: <Warn:MPLS.LDP.InternalProb>
PD4-464587012	All unicast traffic routed by MPLS is stopped when penultimate hop popping (PHP) is enabled on all MPLS VLANs. VPLS traffic is not impacted.
PD3-203917264	If an LSP is already Up, and an ERO is added such that a subsequent path calculation will fail, the LSP will remain Up, and at the same time, continue to retry to calculate a new path with the new ERO. This situation is not clearly visible in the <code>show mpls rsvp-te lsp detail</code> output. The retry counters incrementing is really the only indication that this is happening. The fields showing the "Msg Src," "Msg Time," "Error code," and "Error Value" should be shown because the reason for the path calculation failure is shown in these fields.
PD3-93069318	Only VLANs configured as protocol <i>any</i> should be added to MPLS.
PD3-92653036	The <code>show mpls label</code> , <code>show mpls rsvp-te label</code> , and <code>show mpls rsvp-te lsp</code> command output currently does not display egress LSPs using advertised implicit NULL labels.
PD3-157687121	ExtremeXOS software uses Control Channel Type 2 to indicate router alert label mode. In MPLS Router Alert Label mode, VCCV packets are encapsulated in a label stack. However, the existing VCCV packets are sent like a stack without any PW label.
PD3-104731701	When a traceroute is performed by setting the MPLS TTL to the IP TTL, ExtremeXOS does not correctly send back an ICMP response. The result is "*" characters in the traceroute for the routers that timed out. If a route is available, ExtremeXOS should attempt to send back an ICMP response.
PD3-93630853	LDP should not advertise a label mapping for a direct VLAN that does not have IP forwarding enabled.
PD3-139423053	Running the <code>show mpls rsvp-te lsp summary</code> command on a system configured with 2,000 ingress LSPs takes an excessive amount of time to process.
PD3-111544904	When a router receives an explicit NULL label, it is incorrectly treated as an implicit NULL label, so rather than sending label 0, no label is sent.
PD3-184989177	When an <code>LDP advertise static</code> setting is set to <code>all</code> , all static routes are treated as egress routes and egress LSPs are created. That is, a label is generated and advertised for the static route. If the router at the end of the static route advertises a label matching that static route, the LSP that was previously an egress LSP becomes a transit LSP. An ingress LSP should also be created whenever a label is received, however, the ingress LSP is never created. Workaround: Do not use the <code>LDP advertise static all</code> configuration in situations where an ingress LSP for a static route is required.



Resolved Issues in ExtremeXOS 15.3

The following issues were resolved in ExtremeXOS 15.3. ExtremeXOS 15.3 includes all fixes up to and including ExtremeXOS 11.1.4.4, ExtremeXOS 11.2.3.3, ExtremeXOS 11.3.4.5, ExtremeXOS 11.4.4.7, ExtremeXOS 11.5.2.10, ExtremeXOS 11.6.5.3, ExtremeXOS 12.0.5, ExtremeXOS 12.1.6, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.4, ExtremeXOS 12.6.2, ExtremeXOS 12.7.1, ExtremeXOS 15.1.2, ExtremeXOS 15.2.1, and ExtremeXOS 15.2.2. For information about those fixes, see the release notes for the specific release.

Table 4: Resolved Issues, Platform-Specific and Feature PDs in ExtremeXOS 15.3

PD Number	Description
General	
PD4-3253491193	DHCPV6Relay: Config upload/download .xsf on executing <code>unconfig switch all</code> and loading <code>dhcpv6config</code> makes interface ID show an invalid tag value than the configured tag value for the VLAN.
PD4-3415999201	EMS messages/events supported.
PD4-3420578731	Static CFM group association with ERPS ring fails after the ring is disabled/enabled.
PD4-2797829304	The command <code>show conf</code> does not show when SSH2 has been enabled.
PD4-3007653061	Quitting the command <code>show config</code> causes memory leaks.
PD4-3405669115	In some circumstances, filtering on a port for "FDB.MACTracking" is not logging any informational messages.
PD4-3357584975	System MAC for VLAN gets overwritten with dynamic MAC in case of hash collision, causing L3 reachability issues.
PD4-3404733334	Switch logs "Info:RtMgr.Server.ProcGetRtMsgFail" message appears frequently (every 30 seconds) when it tries to export an unfeasible route. Informational messages should appear less frequently.
PD4-3138803991	Switches stop working after deleting virtual-router <code>vir_1</code> . This issue occurs in ExtremeXOS 15.2.1 and 15.2.2.
PD4-3400038115	Multicast cache entries associated with router port get deleted after timer expires, if MVR is enabled.
PD4-3300114424	Ingressing local multicast traffic to the super-VLAN via sub-VLAN ports is sent back to the same port.
PD4-3314306583	IGMP general query dropped due to multicast loop detection.
PD4-3118104426	OSPF process ends unexpectedly when moving a VLAN from one area to another area.
PD4-3199780495	Virtual link is not working in OSPF.
PD4-2202806910	OSPFv3 NSSA area configurations are not shown in the CLI command <code>show config ospfv3</code> and <code>show ospfv3 area</code> ; OSPFv3 area type always remains in the Normal area.
PD4-2797679070	BGP config for VRF is not cleaned up after a VRF is deleted.
PD4-3330304088	VRRP process signal 6 ends unexpectedly while adding VLAN and VRRP instances.
PD4-2853885980	Hot plugging management cable with IPv6 address results in "exvlan_ioctl_handler:3036:" error messages in the logs and remains in the "Tentative" state.
PD4-3122862720	ELRP process with signal 6 ends unexpectedly when STP tries to disable a trunk port.
PD4-3328741256	ExtremeXOS 12.6.2.10: Port stops learning on a VLAN after limit learning is enabled on the VLAN.



Table 4: Resolved Issues, Platform-Specific and Feature PDs in ExtremeXOS 15.3 (Continued)

PD Number	Description
PD4-3317958618	Packet loss up to two seconds occurs when one of the MLAG peer switch is rebooted.
PD4-3122863433	"hallsCardAlive-183 Slot 0" message appears at console when a trunk port is disabled by STP.
PD4-3170880275	MPLS packets are being forwarded in software instead of hardware on Summit X670, BlackDiamond X8, and BlackDiamond 8900-40G6X-xm when IP DAD is enabled, causing LDP sessions to drop/timeout.
PD4-3122863468	ARP requests are getting dropped due to the ARP validation when it is received from a DHCP-client.
PD4-3030089102	If TCP tracking for LAG port fails, and then becomes active, the LAG health-check remains down.
PD4-3227396996	NTP and DAD cannot co-exist. If you enable NTP and DAD on the same VLANs, and then reboot the switch by saving the configuration, the NTP configuration for VLANs is removed when switch comes back up.
PD4-3138242320	BFD-protected static route does not failover when the BFD session goes down.
PD4-3124148155	Refresh policy triggers ACL process to end unexpectedly while backing up MSM, and it goes out of sync.
PD4-3138604241	[10063] 100BASE (with phy) FX SFP optics link is not coming up after issuing command <code>config ports 9-11 auto off speed 100 duplex full</code> after saving and then rebooting Summit X440/BlackDiamond 8800 series switches and although partner link is up. This issue is seen between the Summit X440 and X460. This issue occurs in ExtremeXOS: v1522b0-br-SR1-4, 15.2.1.5, and 15.1.2.12.
PD4-3176665990	Combo port does not come up with auto-negotiation off after reboot when preferred medium is configured.
PD4-3204711497	On Summit X670v-48x series switches, Part Number and Serial Number are incorrect for 4th channel when configured in 4x10G mode. This issue also occurs on Summit X650, X480, and X670v.
PD4-3273707491	Rx jabber counter increments with jumbo frames.
PD4-2930580825	<code>vmtFileFetch: execv() failed /usr/bin/wget -v 0 -O /tmp/vmt/MANIFEST</code> when <code>run vm-tracking rep sync-now</code> command is executed and sometimes sync fails.
PD4-3136817360	Duplicated multi-cast packets appear while receiving J/P packets from LHR when a switch is acting as a assert-winner and a non-RP.
PD4-3184827915	IPMC forwarding configuration is corrupted after enabling license.
PD4-3178317084	Multi-cast packets are getting duplicated when source IP address is 0.0.0.0.
PD4-2841783640	<code>show bgp route address-family vpnv4 detail all</code> is showing label as part of the vpnv4 prefix.
PD4-3109325898	LAG member ports do not appear in <code>show ports tag <tag number></code> command output.
PD4-3161130688	Need a way to turn off temporary flooding in STP domains using the CLI.
PD4-3234332590	STP process ends unexpectedly when restricted role is enabled in dot1d mode.
PD4-3250650157	After a radius authentication failure, you are unable to access the switch using failsafe account.
PD4-3237306270	DHCP snooping does not work if VLAN translation and ip-security are configured together on a switch.
PD4-3237904817	Unable to upload the dhcp-binding manually to the server.



Table 4: Resolved Issues, Platform-Specific and Feature PDs in ExtremeXOS 15.3 (Continued)

PD Number	Description
PD4-3237904954	Error message occurs when uploading dhcp-binding to the server after changing the file name.
PD4-3199137619	Switch does not get model information (EVENT.DEVICE_MODEL) when an LLDP-enabled device is unplugged from the switch.
PD4-3128801005	After saving switch configuration with downloaded certificate, switch displays <code>Error Validating Certificate</code> error on thttpd restart.
PD4-3316990487	XML is not working with IPv6, making is you cannot manage the switch from web using an IPv6 address.
PD4-3267725700	ACL process ends unexpectedly when unconfiguring an access-list profile which has 32 characters.
PD4-3234332760	ACL refresh fails with error message - "unavailability of hardware resource or system error".
PD4-3160447091	After downloading and installing SSH module, <code>Image Selected</code> information from <code>show switch</code> command, changes back to the booted partition.
PD4-2678796534	If you use an ACL to mirror traffic to a port and then disable mirroring on that port, traffic continues to flow to that port.
PD4-1996237147	The priority flow control is not disabled when using the <code>disable</code> command with the <code>all</code> option.
PD4-3234332692	SNMP OID for ExtremeportQP TxBytes always return 0.
PD4-3054251501	Typo: In the error message when adding a duplicate log filter in EMS, "your are" needs to be changed to "you are".
PD4-3298782274	SNMP query of <code>extremeBootTime.0</code> returns a value in the local configured time zone instead of UTC.
PD4-3314659004	ISIS process ends unexpectedly while configuring interlevel-filter with a lengthy policy file name.
PD4-3314659087	OSPF process ends unexpectedly if a lengthy password is used.
PD4-3267770777	The number of blank lines printed under the RIP module of the <code>show configuration</code> command output increases with an increase in the number of Layer 3 VLANs.
PD4-3316937162	FDB process crashes while clearing neighbor-discovery cache
PD4-3078305908	"Failed to find_rtMgrClient_hash(0xFFFF000D)" error messages are logged, if netTools fails to register as route manager client on rare occasions after system reboot.
PD4-3273898398	NetTools process ends unexpectedly while displaying auto-provision for a virtual-router with a lengthy name.
PD4-3242760247	Manually configured ACL system application priority reduced by one for every reboot and more application seen in ACL zone.
PD4-3331822984	ACL process ends unexpectedly when creating flow-redirect with a large string for redirect-name.
PD4-3316990444	IPFIX configuration does not take the VR configuration into effect.
BlackDiamond 8800 Series Switches	
PD4-3134475729	Error message appears on the switch while disabling/enabling the slot with mirroring configuration.
PD4-2968093427	With BGP max peering, the <code>disable bgp neighbor all</code> command causes the switch to fail.
PD4-2582883670	Process DCBGP ends unexpectedly with signal 11 when trying to delete 100 inactive neighbors.



Table 4: Resolved Issues, Platform-Specific and Feature PDs in ExtremeXOS 15.3 (Continued)

PD Number	Description
PD4-2931845350	Issuing CLI to enable and apply export policy to VPN VRF on PE (to existing configuration) returns this error: "Error: vr red: Cannot change export policy for protocol bgp while export is enabled"
PD4-3297769048	Switch-fabric egress port is getting deleted when pseudo-wire adds and deletes arrive out of order.
PD4-3193434741	On BlackDiamond 8800 series switches, when the number of packets queued for transmit by the CPU exceeds 500, additional slow path forwarded packets are dropped.
PD4-2051483560	The <code>show iproute summary</code> command output occasionally shows the non-zero number of compressed routes even though the routes are no longer there.
PD4-3190863671	I/O card 10G4Xc goes to failed state after reporting "conduit asynchronous transmit error encountered," and then reboots.
PD4-3298782375	On BlackDiamond 8806 and 8810 series switches running ExtremeXOS 12.5 and later, I/O cards inserted in a dual-use slot (MSM or I/O slot) do not become operational.
PD4-3230196968	Executing "snmpwalk" operation on "pethMainPseTable" causes snmpMaster to end unexpectedly occasionally with signal 6.
BlackDiamond 10800 Switches	
PD4-3255799520	When an EAPS ring has multiples of 284 as protected VLANs, then ports are not blocked and can cause a traffic loop after rebooting the switch.
BlackDiamond X8 Series Switches	
PD4-3328224544	Counters for some VMs are not installed and the following error appears in log: "12/10/2012 11:10:04.33 <Error:HAL.IPV4ACL.Error> MM-B: Rcv checkpoint - dynRuleInst for vlan=0 port=70018 NULL (rule xnv_ing_dyn_rule_007bb5b8ae)".
PD4-3415694835	FAN speed reaches 5,000 RPM after hot-swapping FAN modules and does not return to normal range (3,000 RPM).
PD4-3404733158	IDMGR process ends unexpectedly while binding a user with a lengthy password.
PD4-3178905992	40G24X modules fail with conduit errors when system is stressed with ARP requests/replies.
PD4-3199780261	DCBGP process ends unexpectedly while configuring IPv6 address to the VLAN which was associated with BGP.
PD4-3297768981	<Error:HAL.MPLS.Error> MM-A: pibMplsPwUpdate pibMplsPwNhlfeIImWrite failed with error - 1., appears after <code>restart ports all</code> .
PD4-3144429723	MLAG-MLDv2: Record type for MLDv2 reports is modified after management module failover.
PD4-3282651251	The command <code>show ntp association</code> statistics does not display any output.
PD4-3092483367	ExtremeXOS image download is causing DOS-protect information when downloading over the gig connection.
PD4-3173100402	Spurious characters appear on console when enabling SFLOW.
PD4-3300086043	xmlc process ends unexpectedly while creating xml-notification with lengthy parameters.



Table 4: Resolved Issues, Platform-Specific and Feature PDs in ExtremeXOS 15.3 (Continued)

PD Number	Description
SummitStack	
PD4-3418269038	Process rtmgr ends unexpectedly with signal 6 on backup while rebooting neighboring switches. "01/16/2013 17:55:46.18 <Warn:RtMgr.Server.NtfyIpmIQFull> Slot-2: Notify IPML queue full. clnt=131138, ntfy peer-id=14. 01/16/2013 17:55:53.48 <Erro:RtMgr.Client.ReplyTimeOut> Slot-1: Client with ID=0x00020F09 Timed out waiting for (LKUPRPF). Process rtmgr pid 1530 died with signal 6 Code: 489eac 00000000 nop 489eb0 8c430094 lw v1,148(v0) 489eb4 8c640038 lw a0,56(v1) 489eb8 <1080001a>beq a0,zero,0x489f24 489ebc 00000000 nop 489ec0 8c830060 lw v1,96(a0) 489ec4 12430003 beq".
PD4-3334236490	On Summit V80 stacks, the following error message appears when booting up the stack without any configuration: "<Erro:HAL.Port.Error> Slot-1: Unable to get media type from slot 4 port 4 error -1".
PD4-2908405391	On SummitStacks, the following warning message appears while rebooting the stack: "<Warn:Kern.Card.Warning> Slot-1: pci 0000:02:00.1: warning: supported max payload size less than 256 bytes <Warn:Kern.Card.Warning> Slot-1: pci 0000:02:00.0: warning: supported max payload size less than 256 bytes". This issue does not occur in ExtremeXOS v15_2_0_21.
Summit Series Switches	
PD4-3206781851	Fan tray failure error messages occur on Summit family switches.
PD4-3327782647	ExtremeXOS should allow creation of new VLANs up until there is 20 MB left of free memory (instead of 30 MB, which is what is happening currently).
PD4-3110253401	The command <code>show port rxerror</code> displays an error if the value for Rx jabber is too large {integer value too large to represent while executing "format "%8u" \$xmlData(reply.message.show_ports_rxerrors.rxJabber)"}.
PD4-2517224847	Maximum CPU sample limit on Summit series switches is limited to 1000 pps and this needs to be documented.
PD4-3330304159	In Summit X440 and X460 PoE capable switches, log message and SNMP traps are not generated when power usage threshold are reached.
PD4-3170880999	For Summit series switches, <code>show switch</code> command output shows incorrect information when switch is below minimum temperature.
Summit X440 Series Switches	
PD4-3178370231	Summit X440 switches are getting stuck at 29.5 C temperature and logging a hot spot temperature on the console instead of switch temperature when it reaches maximum limit.
PD4-3012552953	In Summit X440-8t series switches, 10/100/1000 Base-T link when speed is set to 100 Mbps with "AutoNeg Off", the link is coming up at speed of 1 Gbps and full-duplex mode in X440 and other end X460 link is coming up at speed 100 Gbps. Cannot ping between switches.
PD4-3013094551	In Summit X440-8t series switches, 10/100/1000 Base-T link when speed is set to 100 Mbps with "AutoNeg Off" between Summit x460 and x440. Observed BASET link is not coming up after restart ports. x460-48x x440-8t BASET 9 <----->9 10<----->10 Port (9, 10) in Summit X440-8t is not coming to active state after a save and reboot of the Summit x440 in the other DUT (Summit x460) ports (9, 10) link is coming up.
PD4-3311903079	CRC Error/bad packets should increment the RxError counter only, but the Congestion counter is incorrectly incremented as well.



Table 4: Resolved Issues, Platform-Specific and Feature PDs in ExtremeXOS 15.3 (Continued)

PD Number	Description
Summit X450 Series Switches	
PD4-1673106807	For certain match conditions involving SIPv6 and DIPv6, packets may not hit an ACL in Summit X450a switches.
PD4-3253186902	Runtime diagnostics on Summit X450a switches fail with a XGM2-2Xf module with XFP installed.
Summit X460 Series Switches	
PD4-2700144775	Diffserv examination is not working at the GRE gateway at ingress from the LAN.
PD4-3092483478	QSFP+ direct-attach active optical cable does not link up when used with SummitStack-V80 stacking module in Summit X460 series switches.
PD4-3228083460	Links on VIM modules become active during bootup of switches before configuration is loaded. This occurs on Summit X460 series switches with XMG3-2sf and XMG3-4sf modules.
PD4-3212070797	On Summit X460 series switches, packets are not logged when icmp-type is used.
PD4-3187720889	Summit X460 series switches with XGM3SB-4sf modules show the wrong temperature value in the <code>show temperature</code> command output.
PD4-3092483593	The output for <code>show temperature</code> is not showing the correct maximum/normal temperature of the switch in stacking.
Summit X480 Series Switches	
PD4-2839090271	IPv6 routes are missing in kernel after a port is restarted.
PD4-2839681491	When SX Mini-GBIC SFP is inserted after inserting and removing on that same port, the 100BASE-FX SFP, the link does not become active.
Summit X650 Series Switches	
PD4-2919055907	In Summit X650-24x series switches, 1000 BASE-BX-D/U, ZX and SX copper port fails to negotiate flow control capability with peer ports. As a result, flow control is always set to "none" after save and reboot.
Summit X670 Series Switches	
PD4-2940562102	DCBGP ends unexpectedly with signal-11 when adding/deleting 100 eBGP neighbors for a long duration.
PD4-3192082965	Traffic is not getting forwarded once the PFC Rx-pause is disabled. dot1p Traffic is sent from one device to another, and PFC receives accordingly. When the PFC Rx-pause is disabled at receiving device side, dot1p traffic is blocked in that port.
PD4-3185656278	Ethernet tx pause (802.3x) is not working across a slot in a stacking setup: (2 node stack - 160G stack) Slot-1 x670v-48t<--10G-->ixia-1 Slot-2 x670v-48x<--1G--->ixia-2 VIM: VIM4-40G4X sending traffic from ixia-1 to ixia-2 (rate-limit 2M) and expecting the Rx Ethernet pause at ixia-2 It looks like RX_PAUSE_EN is not getting enabled in 40G High Gig port.
PD4-3178905832	On Summit X670v switches with L2 LACP load sharing, traffic does not get forwarded on one of the member port after a save and reboot. Issue occurs with the following versions: ExtremeXOS 15.1.1.1, 5.1.2.12, and 15.2.1.
PD4-2823233670	oui and vpn-index range errors are printed in decimal.
PD4-2823233880	Add "(default)" to "both" keyword when configuring route-target.
PD4-3170880245	For Summit X670v-48t series switches, maximum local IPv4 hosts and IPv6 hosts are not allowed in the LPM and L3 hash table. Also, update theoretical maximum in legend of <code>show iproute reserved-entries statistics</code> command.



Table 4: Resolved Issues, Platform-Specific and Feature PDs in ExtremeXOS 15.3 (Continued)

PD Number	Description
PD4-3116293491	QoS profile deletion is incorrectly allowed when the PFC rx-pause is enabled for the same QoS profile.
PD4-3116293726	Display shows both rx and tx are enabled when only one mode is enabled in PFC.
PD4-3185869422	Update temperature settings.
PD4-3334930168	In the <i>ExtremeXOS 15.X Concept Guide</i> , need to update the default stacking protocol for Summit X670 series switches.
E4G Cell Site Routers	
PD4-3346820155	Error "RMEP Creation Failed due to HAL problem" appears when creating 256 MEPs on E4G-200-12x cell site routers.
PD4-2989427990	Log message appears "Unable to add route to unit 0, rc Entry exists. Shadow problem" during reboot.
PD4-2667034651	If a PTP clock port is configured as "slave-only", its state should show as "PTP- Slave" in "show network-clock ptp boundary port" output. However, "Master" appears, instead of "Slave-only".
PD4-2441010796	You can create more than 256 CES services on E4G cell site routers. After creating 496 CES services in E1 mode, there is no CES traffic between nodes in CES after 256 services. Only 256 CES services are working fine. Need to restrict number of services supported on the E4G cell site routers.
PD4-2983613821	For E4G-200 cell site routers, get error log message for added ring ports on control VLAN.
PD4-3253500300	Frequent warning level log messages appear for the timer expiration from CFM process.
PD4-3314862043	For E4G-400 cell site routers, you are able to assign non-existing QoS profile to CES Pseudowire.
AAA	
PD4-3404733231	AAA process ends unexpectedly when radius and tacacs servers are configured with a large input string as the shared-secret.
EAPS	
PD4-3086010511	ExtremeXOS 12.4.4-patch1-4: EAPS shared-port controller stuck in preforwarding [F] even with the shared port down. This can produce a super loop. Issue occurs when there are subsecond link flaps between the partner and controller.
PD4-3314858144	MLAG and EAPS: When an ISC port is part of multiple EAPS domains, and through the ISC port is not a secondary port on the master domain, and only a secondary port in transit domain, the MLAG peer configuration fails with an error.
PD4-3237904730	EAPS fail time range needs to be modified for milliseconds.
L3 VPN	
PD4-2778106871	VPN site temporarily loses route when secondary power is turned on or off.
MPLS	
PD4-3055296411	EXOS_LDP fails test 5.3 due to an internal log message warning for MPLS.
Optics	
PD4-3192802811	When 10G SFP+ passive copper cable is used, Rx errors occur on that port. This issue occurs on Summit X670v and BlackDiamond X8 series switches.



Table 4: Resolved Issues, Platform-Specific and Feature PDs in ExtremeXOS 15.3 (Continued)

PD Number	Description
OSPF	
PD4-3423309981	With two adjacencies established, OSPFv3 process ends unexpectedly when link local address is replaced.
Security	
PD4-3206070707	Deploying identity management (IDM) on ExtremeXOS 15.2 with Ridgeline 3.1 does not work properly. If IDM is deployed on earlier versions of ExtremeXOS, and then upgraded to ExtremeXOS 15.2, IDM functions properly.
PD4-3271740768	While using Dot1x and MAC-based netlogin on the same port, the MAC re-authentication timer should stop when the client is authenticated with dot1x credentials.
PD4-3288465234	Log messages are not generated when SSH2 access is rejected by access-profile.
PD4-3157387244	Idmgr process ends unexpectedly after <code>configure identity-management delete ports all</code> command is issued.
PD4-3302318749	When solicited ARP violation is triggered on a LAG port, the following error message appears: <Error:FDB.ArpError> Unable to retrieve pif, for slot:port=0:0.
PD4-3178933920	The command <code>unconfig identity-management</code> returns the error "no such variable" when no identity management configuration exists.
PD4-3238981482	Implement back door username and passwords, update patent numbers, and update copyright dates.
SNMP	
PD4-2770013137	SET operation on downloadControl table is not functioning.
PD4-2983106991	EXTREME-CFGMGMT-MIB:extremeLastSaveCfgTable has lexicographical error.
PD4-3033106992	Lexicographical error on extremeDownloadImageTable while walking for a stack setup.
PD4-3332155202	snmpMaster process ends unexpectedly with signal 6.
PD4-3333709136	When an SNMP get operation is performed from an SNMP manager and through CLI script, snmpMaster process ends unexpectedly with signal 6.
PD4-3174219240	snmpmaster memory leak occurs while polling from Ridgeline.
PD4-3146090044	SNMPv3 request authentication fails after switch reboot.
PD4-3240646472	SNMPv3 polling causes memory leak in snmpMaster process.
PD4-3291095414	snmpSubagent process ends unexpectedly while sending the snmpset with 1.3.6.1.2.1.3.1.1.2.1 OID.
PD4-3174250688	SNMP walk does not work after restarting snmpSubagent process.
VLANs	
PD4-3269759161	Memory leak occurs after ending a telnet session with the <code>show vlan detail</code> command.
PD4-3158123625	In PVLAN configurations, when an IP address is configured for network VLANs, multicast traffic with a destination address 224.0.0.0/24 is sent back and forth between network VLANs. This affects switch's performance.
PD4-3316990558	Configuring translation VLAN as sub-VLAN causes mcmgr process to end unexpectedly.
PD4-3189051734	Need to update kernel error message when the switch receives the ARP-reply when there is no ARP entry for ARP-sender.

