Extreme®
Connect Beyond the Network

ExtremeXOS Release Notes

Software Version ExtremeXOS 16.2

121160-00

Published August 2016

Legal Notice

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

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

Trademarks

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

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

For additional information on Extreme Networks trademarks, please see:
www.extremenetworks.com/company/legal/trademarks

Support

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

For information, contact:
Extreme Networks, Inc.
145 Rio Robles
San Jose, California 95134
USA

Table of Contents

Preface

Conventions

This section discusses the conventions used in this guide.

Text Conventions

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

Table 1: Notice Icons

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

Table 2: Text Conventions

Convention	Description
`Screen displays`	This typeface indicates command syntax, or represents information as it appears on the screen.
The words **enter** and **type**	When you see the word "enter" in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says "type."
[Key] names	Key names are written with brackets, such as **[Return]** or **[Esc]**. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press **[Ctrl]**+**[Alt]**+**[Del]**
Words in italicized type	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.

Platform-Dependent Conventions

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

- ExtremeSwitching® switches
- Summit® switches
- SummitStack™

When a feature or feature implementation applies to specific platforms, the specific platform is noted in the heading for the section describing that implementation in the *ExtremeXOS 21.1 Command Reference Guide*. In many cases, although the command is available on all platforms, each platform uses specific keywords. These keywords specific to each platform are shown in the Syntax Description and discussed in the Usage Guidelines.

Terminology

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

Providing Feedback to Us

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

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

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

Getting Help

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

- Global Technical Assistance Center (GTAC) for Immediate Support
  - **Phone:** 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact
  - **Email:** support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- GTAC Knowledge — Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- The Hub — A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- Support Portal — Manage cases, downloads, service contracts, product licensing, and training and certifications.

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

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related Return Material Authorization (RMA) numbers

Related Publications

ExtremeXOS Publications

- *ACL Solutions Guide*
- *EMS Messages Catalog*
- *ExtremeXOS 21.1 Command Reference Guide*
- *ExtremeXOS 16.2 Command Reference Guide*
- *ExtremeXOS 21.1 Feature License Requirements*
- *ExtremeXOS 16.2 Feature License Requirements*
- *ExtremeXOS OpenFlow User Guide*
- *ExtremeXOS Quick Guide*
- *ExtremeXOS 21.1 User Guide*
- *ExtremeXOS 16.2 User Guide*
- *ExtremeXOS Legacy CLI Quick Reference Guide*
- *ExtremeXOS Release Notes*
- *Extreme Hardware/Software Compatibility and Recommendation Matrices*
- *Switch Configuration with Chalet*
- *Using AVB with Extreme Switches*

Hardware Documentation

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

Extreme Management Center™ Documentation

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

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

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

Open Source Declarations

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

1 Overview

These release notes document ExtremeXOS 16.2 which adds features, adds supported hardware, and resolves software deficiencies.

New and Corrected Features in ExtremeXOS 16.2

This section lists the new and corrected features supported in the 16.2 software:

Cisco Discovery Protocol (CDPv2)

Support for Cisco Discovery Protocol (CDPv1) was added in ExtremeXOS 15.4. This update to the feature adds support for Cisco Discovery Protocol (CDPv2). CDPv2 is a proprietary protocol designed by Cisco to help administrators collect information about nearby, and directly connected, devices. Support for listening, lifting, processing, and periodic transmitting of the CDPv1/v2 control packets on a per-port basis is implemented in this current release.

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

- SNMP is not supported.

Changed CLI Commands

Changes are underlined.

configure **cdp voip-vlan** [*vlan_name* | *vlan_id* | **dot1p** | **untagged** | **none**] **ports** [*port_list* | **all**]

configure cdp **trust-extend** [**untrusted** | **trusted**] **ports** [*port_list* | **all**]

configure cdp **cos-extend** *cos_value* **ports** [*port_list* | **all**]

show cdp ports {*port_list*} {**configuration**}

configure cdp **power-available** [**advertise** | **no-advertise**] **ports** [*port_list* | **all**]

The output of the following show commands is changed (shown in bold):

```
X460-48t.1 # show cdp
CDP Transmit time               : 60 seconds
CDP Hold time                   : 180 seconds
CDP Device ID                   : 00:04:96:8B:C2:CA
CDP Enabled ports               : 1-2, 7
Power Available TLV Enabled ports: 1-2,23
```

```
X460-48t.23 # show cdp ports
Neighbor Information
--------------------
Port  Device-Id            Hold time  Remote CDP  Port ID
                                      Version
----  ---------            ---------  ----------  --------------------
1     Eni-Extreme-x440-sw> 149        Version-1   Slot:  1, Port:  1
2     00:04:96:8B:9D:B0    160        Version-2   Slot:  1, Port:  2
7     00:04:96:8B:C1:ED    138        Version-2   Slot:  1, Port:  7
> indicates that the value was truncated to the column size in the output.
Use the "show cdp neighbor detail" command to see the complete value.
```

```
X460-48t.3 # show cdp neighbor
Device Id           Local       Hold   Capability  Platform       Port Id
                    Interface   Time
-------------------------------------------------------------------------------
Eni-Extreme-x440-sw> 1          150     T          X440-24t-10G   Slot:  1, P>
00:04:96:8B:9D:B0    2          171     T          X440-48t       Slot:  1, P>
00:04:96:8B:C1:ED    7          134     T          X460-48t       Slot:  1, P>
-------------------------------------------------------------------------------
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,
                  S - Switch, H - Host, I - IGMP, r - Repeater
> indicates that the value was truncated to the column size in the output.
Use the "detail" option to see the complete value.
```

```
X460-48t.7 # show cdp neighbor detail
--------------------------------------------------------------------------
Device ID               : Eni-Extreme-x440-switch-1
Port ID (outgoing port) : Slot:  1, Port:  1
Advertisement Version   : 2
IP Addresses            : 10.10.10.2
Platform                : X440-24t-10G
Interface               : 1
Holdtime                : 173

Version                 :
ExtremeXOS version 15.7.0.22 fixes_v1570b9 by kosharma
     on Tue Feb 24 11:53:33 IST 2015

Native VLAN             : 1
Duplex                  : Full
SysName                 : X440-24t-10G
Location                : Chennai
Power Request Id        : 24333
Power Management Id     : 2
```

```
Power Drawn             : 1500 mW
Power Consumed          : 3454 mW
```

```
X460-48t.11 # show cdp ports configuration
Local Port Information
----------------------
Port     Trust       COS   Voice-VLAN
----     ----------  ----  --------------------
1        Trusted     0     none
2        Untrusted   4     none
7        Untrusted   0     Default
```

Virtual Router Redundancy Protocol (VRRP) Fabric Routing

Virtual Router Redundancy Protocol (VRRP) has one master router that does L3 routing and one or more backup routers that perform L2 forwarding of packets toward the master router, as per VRRP RFC specification. With this method, L3 routing capability of backup routers goes unused. This also causes loss of bandwidth in the links that connect master and backup routers. This issue is present in any topology where host traffic is flowing using the backup routers. With multiple backup routers, traffic from hosts attached to some backup routers have to traverse multiple links to reach the master router. This causes loss of bandwidth in multiple links toward the master.

This feature allows backup routers to take part in L3 routing for the packets it receives with the destination address equal to VMAC. Backup routers enabled with this feature are called Fabric Routing Enabled Backup (FREB) routers. This feature allows:

- Load sharing of traffic between VRRP routers
- Bandwidth savings on the links connecting master and backup routers

This solution is applicable for all topologies, such as MLAG, EAPS, or STP.

Supported Platforms

- Summit X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

- Fabric Routing feature is not supported for VRRP VR for which Virtual IP is same as interface IP (owned IP).
- Traffic sent from host destined for VIP is L3 forwarded by FREB router if FREB router sits in between, even though both are in same subnet. VIP cannot be used to run protocols between host and VRRP router, which expects TTL value to not be decremented, for example BFD.
- PVLAN configuration will not be supported in this release.
- VLAN Aggregation configuration will not be supported in this release.

New CLI Commands

```
configure vrrp {vlan vlan_name vr vr_id | all} fabric-route-mode [on | off]
```

Virtual Router Redundancy Protocol (VRRP) Host Mobility

The Virtual Router Redundancy Protocol (VRRP) Host mobility feature solves the Asymmetric routing problem associated with VRRP where the path to return to an end host may be different and longer

than necessary. This feature uses host-routes to indicate where in the network an end host resides. Using other routing protocols such as OSPF, other routers then pick the shortest path back to the end host when multiple paths are available using Equal Cost Multi Path (ECMP) route entries.

Supported Platforms

- Summit X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

- Bound to FDB's ARP limitations
- Bound to Route Manager's entry limitations

Changed CLI Commands

Changes are underlined.

`configure vrrp {vlan} vlan_name vrid vridval host-mobility [{on | off} {exclude-ports [add | delete] port_list}]`

`configure iproute {ipv4} priority [static | blackhole | rip | bootp | icmp | ospf-intra | ospf-inter | ospf-as-external| ospf-extern1 | ospf-extern2 | ebgp | ibgp | mpls | isis | isis-level-1 | isis-level-2 | isis-level-1-external | isis-level-2-external | host-mobility] priority {vr vrname}`

`unconfigure iproute {ipv4} priority [static | blackhole | rip | bootp | icmp | ospf-intra | ospf-inter | ospf-as-external | ospf-extern1 | ospf-extern2 | ebgp | ibgp | mpls | isis | isis-level-1 | isis-level-2 | isis-level-1-external | isis-level-2-external | host-mobility | all ] {vr vrname}`

`configure iproute ipv6 priority [static | blackhole | ripng | icmp | ospfv3-intra | ospfv3-inter | ospfv3-as-external| ospfv3-extern1 | ospfv3-extern2 | isis | isis-level-1 | isis-level-2 | isis-level-1-external | isis-level-2-external | host-mobility] priority {vr vrname}`

`unconfigure iproute ipv6 priority [static | blackhole | ripng | icmp | ospfv3-intra | ospfv3-inter | ospfv3-as-external | ospfv3-extern1 | ospfv3-extern2 | isis | isis-level-1 | isis-level-2 | isis-level-1-external| isis-level-2-external | host-mobility | all ] {vr vrname}`

The existing `enable ospf export`, `disable ospf export`, and `configure ospf export` commands are expanded to allow a new route type of "host-mobility." Configuring host-mobility to be exported causes OSPF to redistribute host-mobility routes.

The existing `enable ospfv3 export` and `disable ospfv3` commands are expanded to allow a new route type of "host-mobility." Configuring host-mobility to be exported causes OSPFv3 to redistribute host-mobility routes.

The output of the following show commands is changed (shown in bold):

```
# show vrrp detail
VLAN:  vlan23   VRID:  1        VRRP:  Disabled State:  INIT
```

```
Virtual Router:  VR-Default
Priority:  100(backup)  Advertisement Interval:  1 sec
Version: v3-v2  Preempt:  Yes   Preempt Delay: 0 sec
Virtual IP Addresses:
Accept mode:  Off
Host-Mobility:  On
Host-Mobility Exclude-Ports: 1, 10
Checksum:  Include pseudo-header
Tracking mode:  ALL
Tracked Pings:  -
Tracked IP Routes:  -
Tracked VLANs:  -
Fabric Routing: Off
```

```
# show ospf
OSPF                : Disabled         MPLS LSP as Next-Hop: No
RouterId            : 0.0.0.0          RouterId Selection  : Automatic
ASBR                : No               ABR                 : No
ExtLSA              : 0                ExtLSAChecksum      : 0x0
OriginateNewLSA     : 0                ReceivedNewLSA      : 0
SpfHoldTime         : 3                Lsa Batch Interval  : 30s
CapabilityOpaqueLSA : Enabled
10M Cost            : 10               100M Cost           : 5
1000M Cost (1G)     : 4                10000M Cost (10G)   : 2
40000M Cost (40G)   : 2
100000M Cost (100G) : 1
Router Alert        : Disabled         Import Policy File  :
ASExternal LSALimit : Disabled         Timeout (Count)     : Disabled (0)
Originate Default   : Disabled
SNMP Traps          : Disabled
VXLAN Extensions    : Disabled
Redistribute:
Protocol              Status   cost  Type Tag        Policy
direct                Disabled 0     0    0          None
static                Disabled 0     0    0          None
rip                   Disabled 0     0    0          None
e-bgp                 Disabled 0     0    0          None
i-bgp                 Disabled 0     0    0          None
isis-level-1          Disabled 0     0    0          None
isis-level-2          Disabled 0     0    0          None
isis-level-1-external Disabled 0     0    0          None
isis-level-2-external Disabled 0     0    0          None
host-mobility         Enabled  0     2    0          None
```

```
# show ospfv3
OSPFv3              : Disabled           RouterId            : 0.0.0.0
RouterId Selection  : Automatic          ASBR                : No
ABR                 : No                 ExtLSAs             : 0
ExtLSAChecksum      : 0x0                OriginateNewLSAs    : 0
ReceivedNewLSAs     : 0                  SpfHoldTime         : 3s
Num of Areas        : 1                  LSA Batch Interval  : 0s
10M Cost            : 100                100M Cost           : 50
1000M Cost (1G)     : 40                 10000M Cost (10G)   : 20
40000M Cost (40G)   : 20                 100000M Cost (100G) : 10
Graceful Restart    : None               Grace Period        : 120s
Import Policy File  : none
Redistribute:
  Protocol              Status    Cost      Type  Tag    Policy
  direct                Disabled  20        2     ---    none
  e-bgp                 Disabled  20        2     ---    none
  i-bgp                 Disabled  20        2     ---    none
  ripng                 Disabled  20        2     ---    none
  static                Disabled  20        2     ---    none
  isis-level-1          Disabled  20        2     ---    none
  isis-level-2          Disabled  20        2     ---    none
  isis-level-1-external Disabled  20        2     ---    none
  isis-level-2-external Disabled  20        2     ---    none
  host-mobility         Enabled   0         2     ---    none
```

```
show iproute
Ori  Destination        Gateway         Mtr  Flags         VLAN       Duration
 d   192.168.24.0/24    192.168.24.44   1    -------um---- vlan24     0d:4h:20m:48s
*hm  192.168.23.1/32    192.168.23.1    1    UGHD---u---f- vlan23     0d:0h:16m:5s

      (mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2,
      (oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM,
```

```
       (r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (hm) Host-mobility, (un) UnKnown,
       (*) Preferred unicast route (@) Preferred multicast route,
       (#) Preferred unicast and multicast route.

Flags: (b) BFD protection requested, (B) BlackHole, (c) Compressed, (D) Dynamic,
       (f) Provided to FIB, (G) Gateway, (H) Host Route, (l) Calculated LDP LSP,
       (L) Matching LDP LSP, (m) Multicast, (p) BFD protection active, (P) LPM-routing,
       (R) Modified, (s) Static LSP, (S) Static, (t) Calculated RSVP-TE LSP,
       (T) Matching RSVP-TE LSP, (u) Unicast, (U) Up, (3) L3VPN Route.

MPLS Label: (S) Bottom of Label Stack
Mask distribution:
     1 routes at length 24

Route Origin distribution:
     1 routes from Direct

Total number of routes = 1
Total number of compressed routes = 0
```

```
# show iproute ipv6
Ori Destination                                 Mtr  Flags         Duration
    Gateway                                     Interface
*hm 2000::/128                                  1    UGHD---u---f- 0d:0h:0m:7s
    2000::2                                     vlan23
#d  2000::/64                                   1    U------um--f- 0d:20h:19m:46s
    2000::1                                     vlan23
#d  fe80::%vlan23/64                            1    U------um--f- 0d:20h:19m:46s
    fe80::204:96ff:fe51:f96d                    vlan23

Origin(Ori):(b) BlackHole, (be) EBGP, (bg) BGP, (bi) IBGP, (bo) BOOTP,
            (ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext,
            (e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2,
            (is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (ma) MPLSIntra,
            (mr) MPLSInter, (mo) MOSPF (o) OSPFv3, (o1) OSPFv3Ext1, (o2) OSPFv3Ext2,
            (oa) OSPFv3Intra, (oe) OSPFv3AsExt, (or) OSPFv3Inter, (pd) PIM-DM, (ps) PIM-
SM,
            (r) RIPng, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (hm) Host-mobility, (un)
UnKnown,
            (*) Preferred unicast route (@) Preferred multicast route,
            (#) Preferred unicast and multicast route.

Flags: (b) BFD protection requested, (B) BlackHole, (c) Compressed Route,
       (D) Dynamic, (f) Provided to FIB, (G) Gateway, (H) Host Route,
       (l) Calculated LDP LSP, (L) Matching LDP LSP, (m) Multicast,
       (p) BFD protection active, (P) LPM-routing, (R) Modified, (s) Static LSP,
       (S) Static, (t) Calculated RSVP-TE LSP, (T) Matching RSVP-TE LSP,
       (u) Unicast, (U) Up, (3) L3VPN Route.

Mask distribution:
     2 routes at length  64

Route Origin distribution:
     2 routes from Direct

Total number of routes = 3
Total number of compressed routes = 0
```

```
# show iproute priority
Direct                10
MPLS                  20
Blackhole             50

Static                1100
HostMobility          1150
ICMP                  1200
EBGP                  1700
IBGP                  1900
OSPFIntra             2200
OSPFInter             2300
Isis                  2350
IsisL1                2360
IsisL2                2370
RIP                   2400
OSPFAsExt             3100
OSPFExt1              3200
```

```
OSPFExt2                3300
IsisL1Ext               3400
IsisL2Ext               3500
Bootp                   5000
```

```
# show iproute ipv6 priority
Direct                  10
Blackhole               50

Static                  1100
HostMobility            1150
ICMP                    1200
EBGP                    1700
IBGP                    1900
OSPFv3Intra             2200
OSPFv3Inter             2300
Isis                    2350
IsisL1                  2360
IsisL2                  2370
RIPng                   2400
OSPFv3AsExt             3100
OSPFv3Ext1              3200
OSPFv3Ext2              3300
IsisL1Ext               3400
```

Internet Protocol Flow Information Export (IPFIX) Mirroring Enhancement

This feature enhances the mirroring capabilities in ExtremeXOS by adding IPFIX flow traffic support, in addition to the previously supported port and VLAN traffic. With the ability to mirror IPFIX flow traffic, you can leverage the combined capabilities of Internet Protocol Flow Information Export (IPFIX) and ExtremeAnalytics to provide additional information about flows. IPFIX can detect flows and collect flow statistics, but it cannot do deep packet payload inspections. ExtremeAnalytics, however, can do deep packet inspection beyond Layer 4, if it is provided with a copy of the packet payload. This feature mirrors the first 15 packets of any IPFIX flow to a port where ExtremeAnalytics is able to receive the packets for deep packet inspection.

Supported Platforms

- Summit X460, X460-G2 series switches
- BlackDiamond X8 series switches (40G12X-XL, 100G4X-XL, and 100G4X)

Changed CLI Commands

Changes are underlined.

configure mirror {*mirror_name* **|** *mirror_name_li***} add | delete [vlan** *name* **{ingress | port** *port* **{ingress}} | ip-fix | port** *port* **{vlan** *name* **{ingress} | ingress | egress | ingress-and-egress | anomaly}]**

The output of the following show command is changed (shown in bold):

```
# show mirror

DefaultMirror   (Disabled)
    Description:    Default Mirror Instance, created automatically
    Mirror to port: -

MyMirror   (Disabled)
    Description:
    Mirror to port: 2:1
    Source filters configured :
        Ports 2:2-3, all vlans, ingress and egress
        Port  2:5, ip-fix
```

Ethernet Ring Protection Switching (ERPS) Management Information Base (MIB) Support

Managed objects for Ethernet Ring Protection Switching (ERPS) Management Information Base (MIB) are defined in ExtremeXOS 16.2. ExtremeXOS 16.2 implements:

- extremeErpsProtectedVlanTable—contains the grouping of set of protected VLANs
- extremeErpsRingTable—each entry in extremeErpsRingTable has information about one ring in the switch
- extremeErpsStatsTable—contains statistics information for each of the rings present in the switch
- extremeErpsGlobalInfo—contains the information of ERPS configured globally in the switch
- extremeErpsNotification—contains two types of traps, extremeErpsStateChangeTrap and extremeErpsFailureTrap

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

Groups and tables are implemented as read-only.

ExtremeCFM Management Information Base (MIB)

This feature introduces the proprietary ExtremeCFM Management Information Base (MIB) that provides information about the Connectivity Fault Management (CFM) Group. This is an extension to IEEE8021-CFM-MIB.

The following objects are defined in the CFM Group MIB module:

- extremeCfmNotifications
- extremeCfmMibObjects
- extremeCfmMibConformance

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Link Aggregation Control Protocol (LACP) Fallback Option Feature

Preboot Execution Environment (PXE) is an industry standard client/server environment that allows workstations to boot from the server before their full operating system is up and running. PXE images are too small to take advantage of Link Aggregation Control Protocol (LACP) functionality, and therefore it is up to the administrator to statically configure the switch for correct connectivity. This also means that after the full operating system is up and running, the switch needs to be reconfigured for LACP. The LACP Fallback option automates this process.

The LACP Fallback option lets you select a single port that is automatically added to the aggregator if LACP data units (LACPDUs) do not appear on any of the member ports within the specified period of time. If LACPDUs are exchanged before this timeout expires, an aggregator is formed using traditional means. If LACPDUs are not received, an active port with the lowest priority value is automatically added to the aggregator (enters fallback state). If ports have the same priority value, the lowest port number on the lowest slot number is chosen.

The selected port stays in the fallback state until fallback is disabled or until LACPDUs are received on any of the member ports, at which point the old aggregator is removed and a new one is selected based on information propagated in the LACPDUs. The new fallback port may also be re-elected if the existing fallback port changes its state (for example, port priority change, link bounce, port disable/enable, etc.).

The LACP fallback option configuration consists of:

- Selecting a fallback port by setting its LACP port priority (optional)
- Configuring the fallback timeout (optional)
- Enabling fallback (mandatory)

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

When using LACP fallback with MLAG, fallback port is selected only on the LACP master.

New CLI Commands

`configure sharing port lacp fallback [enable | disable]`

Changed CLI Commands

The `show lacp lag group-id detail` command now shows fallback information (shown in bold):

```
# show lacp lag 17 detail

Lag   Actor    Actor  Partner           Partner  Partner Agg   Actor
      Sys-Pri  Key    MAC               Sys-Pri  Key     Count MAC
--------------------------------------------------------------------------------
17          0  0x03f9 00:00:00:00:00:00       0  0x0000      1 00:04:96:6d:55:13

Enabled          : Yes
LAG State        : Up
Unack count      : 0
Wait-for-count   : 0
Current timeout  : Long
Activity mode    : Active
Defaulted Action : Delete
Fallback         : Enabled
Fallback timeout : 40 seconds
Receive state    : Enabled
Transmit state   : Enabled
Minimum active   : 1
Selected count   : 1
Standby count    : 0
LAG Id flag      : Yes
  S.pri:0   , S.id:00:04:96:6d:55:13, K:0x03f9
  T.pri:0   , T.id:00:00:00:00:00:00, L:0x0000
```

```
Port list:

Member     Port      Rx           Sel          Mux            Actor     Partner
Port       Priority  State        Logic        State          Flags     Port
-------------------------------------------------------------------------------
17         10        Initialize   Unselected   Detached       A-G-----  0
18         5         Initialize   Fallback     Collect-Dist   A-GSCD--  1018
19         5         Idle         Unselected   Detached       --------  0
===============================================================================
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
             C-Collecting, D-Distributing, F-Defaulted, E-Expired
```

The **show lacp member-port** *port* **detail** command now shows fallback information (shown in bold):

```
# show lacp member-port 18 detail

Member     Port      Rx           Sel          Mux            Actor     Partner
Port       Priority  State        Logic        State          Flags     Port
-------------------------------------------------------------------------------
18         5         Initialize   Fallback     Collect-Dist   A-GSCD--  1018
Up            : Yes
Enabled       : Yes
Link State    : Up
Actor Churn   : False
Partner Churn : True
Ready_N       : Yes
Wait pending  : No
Ack pending   : No
LAG Id:
  S.pri:0   , S.id:00:04:96:6d:55:13, K:0x03f9, P.pri:65535, P.num:1018
  T.pri:0   , T.id:00:00:00:00:00:00, L:0x0000, Q.pri:65535, Q.num:1018
Stats:
  Rx - Accepted                                    : 0
  Rx - Dropped due to error in verifying PDU       : 0
  Rx - Dropped due to LACP not being up on this port : 0
  Rx - Dropped due to matching own MAC             : 0

  Tx - Sent successfully                           : 1162
  Tx - Transmit error                              : 0
===============================================================================
Actor Flags: A-Activity, T-Timeout, G-Aggregation, S-Synchronization
             C-Collecting, D-Distributing, F-Defaulted, E-Expired
```

OpenSSL Federal Information Processing Standards (FIPS) Object Module v2.0

This feature adds Federal Information Processing Standards (FIPS) compliance Object Module v2.0 (an open source library named openssl-fips-ecp-2.0.9).

OpenSSL is a software library used in applications to secure communications against eavesdropping or to ascertain the identity of the party at the other end. This feature does not validate the OpenSSL module itself, but instead implements a new software component called the OpenSSL FIPS Object Module.

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

New CLI Commands

configure security fips-mode [on | off]

```
show security fips-mode
```

Link Aggregation Group (LAG) Support for Audio Video Bridging (AVB)

This feature completes the capability to use Link Aggregation Group (LAG) ports with Audio Video Bridging (AVB) by adding support for LAG ports with Multiple Stream Reservation Protocol (MSRP).

This feature adds two modes for how MSRP calculates the available bandwidth of a LAG for use in making stream reservations:

- Single-port mode simply provides link redundancy and the LAG effective bandwidth is the same as the bandwidth of a single member port.
- Cumulative mode allows bandwidth aggregation and the LAG effective bandwidth is set to a configurable percent of aggregate bandwidth of the member ports in the LAG. This feature also adds generalized Precision Time Protocol (gPTP) configuration support at the LAG level. Only the LAG master port needs to be specified when making gPTP configurations. However, the protocol is still running on each member port at the physical port level.

Supported Platforms

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

Changed CLI Commands

show msrp ports {*port_list*} **detail**

For the preceding command, with LAG support, the port speed is replaced with “effective speed.”.For physical ports, the effective speed is equivalent to the port speed (shown in bold).

```
Port  Enabled   Oper  Effectv  Dplx  Jumbo  Jumbo  Cls  Bndry   State   Sr-Pvid
                          Speed                Size              App/Reg
----  -------  ------  ------  ----  -----  -----  ---  -----  -------  -------
*2g   Y        Up       150 M  Full  N       9216  A    N      QA/IN          2
                                                   B    N      QA/IN          2
*48   Y        Up      1000 M  Full  N       9216  A    N      QA/IN          2
                                                   B    N      QA/IN          2
```

With the **detail** option, and if the port is a LAG, additional information appears:

```
Load sharing ports:
                    Port     Port          BW Mode       Percentage
                             Speed
                    ----     ------     -----------     ----------
                    *2g      200 M      Cumulative       40%
```

Event Management System (EMS) IPv6 Syslog Server Support

This feature adds support for the Event Management System (EMS) to send log messages to Syslog servers having IPv6 addresses.

The Event Management System supports the logging of event occurrences to external Syslog server targets. Each Syslog server target is identified by its IP address, UDP port, VRID, and local use facility (for example: “local0” through “local7”). Previously, the IP address of a Syslog server target was limited to the IPv4 address family, but with this feature it can be of the IPv6 address family.

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Changed CLI Commands

The existing EMS ("log") commands relevant to Syslog server targets now support IPv6 server (and source, as applicable) addresses:

```
configure syslog add [ipaddress {udp-port udp_port}|ipPort] {vr vr_name}
[local0...local7]

configure syslog delete [all | ipaddress {udp-port udp_port}| ipPort] {vr
vr_name}{local0...local7}

configure log target syslog [all | ipaddress {udp-port udp_port} | ipPort] {vr
vr_name} {local} from source-ip-address

[enable|disable] log target [ . . . | syslog [[all | ipaddress {udp-port
udp_port} | ipPort] {vr vr_name} {local}]]

configure log target syslog [ipaddress {udp-port udp_port} | ipPort] {vr vr_name}
[local] severity severity {only}

configure syslog [ipaddress {udp-port udp_port} | ipPort] {vr vr_name} [local]
severity severity {only}

configure log target [ . . . | syslog [all | ipaddress {udp-port udp_port} |
ipPort] {vr vr_name} {local}] match {any | regex}

configure log target syslog [all | ipaddress {udp-port udp_port} | ipPort] {vr
vr_name} {local} format

unconfigure log target [ . . . | syslog [all | ipaddress {udp-port udp_port} |
ipPort] {vr vr_name} {local} | . . . ] format

show log configuration {target { . . . | syslog {ipaddress {udp-port udp_port} |
ipPort} {vr vr_name} {local}} | filter {filter-name}}
```

MAC Authentication Delay

Currently, when both dot1x and MAC authentication methods are enabled on a port, a new MAC address detection triggers ExtremeXOS to send a RADIUS request to authenticate the new client on that port using MAC-based authentication. This feature allows you delay/bypass the MAC authentication by configuring a MAC authentication delay period on a per port basis. The MAC authentication delay period's default value is 0 seconds for backward compatibility, with a permissible range of 0 to 120 seconds.

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches

- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Changed CLI Commands

Changes are underlined.

configure netlogin mac ports [*port_list* | all] timers [{reauth-period [*reauth_period*]} {reauthentication [on|off]} {delay [*delay*]}]

The output of the **show netlogin** command now includes the authentication delay period value (shown in bold):

```
NetLogin Authentication Mode : web-based DISABLED;  802.1x DISABLED;  mac-based DISABLED
NetLogin VLAN                : Not Configured
NetLogin move-fail-action    : Deny
NetLogin Client Aging Time   : 5 minutes
Dynamic VLAN Creation        : Disabled
Dynamic VLAN Uplink Ports    : None
Authentication Protocol Order: 802.1x, web-based, mac-based (default)
SNIPPED
-----------------------------------------------
          MAC Mode Global Configuration
-----------------------------------------------
Re-authentication period        : 0 (Re-authentication disabled)
Authentication Database         : Radius, Local-User database
Authentication Delay Period     : 0 (Default)
-----------------------------------------------
Number of Clients Authenticated  : 0
```

Configurable per Slot Link Aggregation Group (LAG) Member Port Distribution

Previously, ExtremeXOS switches would always distribute to all active members in a link aggregation group (LAG). This enhancement provides two options for specifying a subset of the active member ports as eligible for distribution on a per slot basis: “local slot distribution” and “distribution port lists”. The specific choice of configuration is described in the command line syntax as a “distribution-mode”. The choice of distribution mode is configurable per LAG. You may dynamically switch between distribution modes using the **configure sharing distribution-mode** command.

Local Slot Distribution

The “local-slot” distribution mode restricts distribution of unicast packets to the active LAG members on the same slot where the packet was received. If no active LAG members are present on the slot where the packet was received, all active LAG member ports are included in the distribution algorithm.

The “local-slot” distribution mode is useful for reducing the fabric bandwidth load of a switch. Reducing fabric bandwidth may be especially important for a SummitStack, which has significantly less fabric (inter-slot) bandwidth available in comparison to chassis switches. In many chassis or SummitStack hardware configurations, the “local-slot” distribution mode may reduce the switching latency of some flows distributed to a LAG.

Distribution Port Lists

The “port-lists” distribution mode configures one or more LAG member ports to be eligible for unicast LAG distribution on each slot in a switch. If a slot does not have a distribution port list configured or if

none of the configured member ports is active in the LAG, all active member ports are eligible for unicast distribution.

The use of the “port-lists” distribution mode should be taken into consideration when adding ports to a LAG with the **configure sharing** command. Any newly added port on a LAG is not available for unicast distribution unless it is also added to the distribution port list of at least one slot.

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches in stacks
- BlackDiamond X8 and 8000 series switches

Limitations

The distribution modes affect only the distribution of known unicast packets on a LAG. Non-unicast packets are distributed among all active members of a LAG.

Changed CLI Commands

Changes are underlined.

enable sharing *master_port* **grouping** *member_port_list* {**algorithm** [**address-based** {**L2** | **L3** | **L3_L4** | **custom**} | **port-based**]} {**distribution-mode** [**all** | **local-slot** | **portlists**]} {**lacp** | **health-check**}

configure sharing *master_port* **distribution-mode** [**all**| **local-slot** | **port-lists**]

configure sharing *master_port* **slot** *slot* **distributionlist** [*port_list* | **add** *port_list* | **delete** [*port_list*] | **all**]]

The **show sharing** and **show ports** *port_list* **sharing** commands now display the distribution mode for a LAG under the “Flag” column:

Distribution Mode Flags:

A - All: Distribute to all members

L - Local: Distribute to members local to ingress slot

P - Port Lists: Distribute to per-slot configurable subset of members

The **show sharing** and **show ports** *port_list* **sharing** commands now display the configured distribution mode and distribution port lists for LAGs:

show {**ports** *port_list*} **sharing** {**distribution configuration**}

```
Config Distribution Distribution
Master Mode Lists
=================================================================
1:1 Port Lists Slot 1: 1:1-10, 1:15
Slot 5: 1:11-22
1:25 Local Slot Slot 1: 1:25
Slot 5: 1:26
5:1 Port Lists
5:10 All Slot 1: 5:11
Slot 5: 5:10
```

Port Customer VLAN ID (CVID) on Port-Based or Customer Edge Port (CEP) VMAN Service

This feature introduces an optional port customer VLAN ID (CVID) parameter to the existing untagged and CEP VMAN port configuration options. When present, any untagged packet received on the port is double tagged with the configured port CVID and the SVID associated with the VMAN. If the port is untagged, packets received with a single CID still have the SVID added. If the port is CEP, only untagged and any specifically configured CVIDs are allowed. As double tagged ports are received from tagged VMAN ports and forwarded to untagged VMAN ports, the SVID associated with the VMAN is stripped. Additionally, the CVID associated with the configured port CVID is also stripped in the same operation. If the port is CEP and CEP egress filtering in enabled, only the specified port CVID and CVIDs are allowed to egress.

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches (except BD8K: G48Te2, G24Xc, G48Xc, G48Tc, 10G4Xc, 10G8Xc, S-G8Xc, S-10G1Xc, S-10G2Xc, and 8500-series)
- E4G-200 and E4G-400 cell site routers

Limitations

- Any limitations that currently exist with untagged VMAN ports also exist when the Port VLAN ID element is additionally applied.
- VPLS service VMANs are not allowed to have port-cvid configurations.

Changed CLI Commands

Changes are underlined.

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

configure vman *vman_name* **ports** [*port_list* | **all**] **add** [**cvid** *cvid_first* { - *cvid_last*} {**translate** *cvid_first_xlate* { - *cvid_last_xlate* } } |**port-cvid** *port_cvid*]

configure vman *vman_name* **ports** [*port_list* | **all**] **delete** [**cvid** *cvid_first* { - *cvid_last* } |**port-cvid** *port_cvid*]

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

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

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

Graceful Restart and Not-So-Stubby Area (NSSA) Supported for Open Shortest Path First (OSPFv3)

This feature upgrades Open Shortest Path First (OSPFv3) to support graceful restart and Not-So-Stubby Area (NSSA):

- Graceful OSPFv3 Restart—RFC 5187 describes a way for OSPFv3 control functions to restart without disrupting traffic forwarding. Without graceful restart, adjacent routers assume that information previously received from the restarting router is stale and should not be used to forward traffic to that router. However, in many cases, two conditions exist that allow the router restarting OSPFv3 to continue to forward traffic correctly. The first condition is that forwarding can continue while the control function is restarted. Most modern router system designs separate the forwarding function from the control function so that traffic can still be forwarded independent of the state of the OSPFv3 function. Routes learned through OSPFv3 remain in the routing table and packets continue to be forwarded. The second condition required for graceful restart is that the network remain stable during the restart period. If the network topology is not changing, the current routing table remains correct. Often, networks can remain stable during the time for restarting OSPFv3.
- NSSA—NSSA is an extension of OSPFv3 stub area. External routes originating from an Autonomous System Boundary Router (ASBR) connected to an NSSA can be advertised within the area and can be advertised to other areas as autonomous system (AS)-external link-state advertisements (LSAs).

Supported Platforms

- Summit X450-G2, X460, X460-G2, X480, X670, X670-G2, and X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

New CLI Commands

```
configure ospfv3 lsa-batch-interval seconds

configure ospfv3 area area-identifier nssa [nosummary | summary] stub-defaultcost
cost {translate}

configure ospfv3 restart [none | planned | unplanned | both]

configure ospfv3 restart grace-period seconds

configure ospfv3 [[vlan | tunnel] all | {vlan} vlan-name | {tunnel} tunnel-name |
area area-identifier] restart-helper [none | planned | unplanned | both]

enable ospfv3 [[vlan | tunnel] all | {vlan} vlan-name | {tunnel} tunnel-name |
area area-identifier] restart-helper-lsa-check

disable ospfv3 [[vlan | tunnel] all | {vlan} vlan-name | {tunnel} tunnel-name |
area area-identifier] restart-helper-lsa-check

enable ospfv3 virtual-link {routerid} router-identifier {area} area-identifier
restart-helper-lsa-check

disable ospfv3 virtual-link {routerid} router-identifier {area} area-identifier
restart-helper-lsa-check
```

Changed CLI Commands

Changes are underlined.

configure **ospfv3 area** *area_identifier* **add range** *ipv6netmask* [**advertise** | **noadvertise**] [**inter-prefix** | **nssa**]

configure **ospfv3 area** *area-identifier* **delete range** *ipv6Netmask* [**inter-prefix** | **nssa**]

configure ospfv3 **virtual-link** {**routerid**} *router-identifier* {**area**} **area-identifier restart-helper** [**none** | **planned** | **unplanned** | **both**]

The following show commands now display additional information (shown in bold):

```
show ospfv3

OSPFv3              : Enabled           RouterId            : 10.1.1.1
RouterId Selection  : Configured        ASBR                : No
ABR                 : No                ExtLSAs             : 0
ExtLSAChecksum      : 0x0               OriginateNewLSAs    : 3
ReceivedNewLSAs     : 0                 SpfHoldTime         : 10s
Num of Areas        : 1                 10M Cost            : 100
100M Cost           : 50                                    1000M Cost (1G)     :
40
10000M Cost (10G)   : 20                40000M Cost (40G)   : 20
100000M Cost (100G) : 10
Num of Areas        : 1                 LSA Batch Interval  : 30s
10M Cost            : 100               100M Cost           : 50
1000M Cost (1G)     : 40                10000M Cost (10G)   : 20
40000M Cost (40G)   : 20                100000M Cost (100G) : 10
Router Alert        : Disabled
ASExternal LSALimit : Disabled          Timeout (Count)     : Disabled (0)
Originate Default   : Disabled
Graceful Restart    : Both              Grace Period        : 120s
Restart Status      : None
Last Restart Exit Reason: None
Import Policy File  : none
Redistribute:
  Protocol              Status    Cost      Type  Tag    Policy
  direct                Disabled  20        2     ---    none
  e-bgp                 Disabled  20        2     ---    none
  i-bgp                 Disabled  20        2     ---    none
  ripng                 Disabled  20        2     ---    none
  static                Disabled  20        2     ---    none
  isis-level-1          Disabled  20        2     ---    none
  isis-level-2          Disabled  20        2     ---    none
  isis-level-1-external Disabled  20        2     ---    none
  isis-level-2-external Disabled  20        2     ---    none
```

```
show ospfv3 interfaces detail

Interface           : v100              Enabled             : ENABLED
Router              : ENABLED           AreaID              : 0.0.0.0
RouterID            : 10.1.1.2          Link Type           : point-to-point
Passive             : No                Cost                : 40/A
Priority            : 1                 Transit Delay       : 1s
Hello Interval      : 10s               Rtr Dead Time       : 40s
Retransmit Interval : 5s                Wait Timer          : 40s
Interface ID        : 19                Instance ID         : 0
State               : P2P               Number of state chg : 1
Hello due in        : 7s                Number of events    : 2
Total Num of Nbrs   : 1                 Nbrs in FULL State  : 1
Hellos Rxed         : 127733            Hellos Txed         : 127739
DB Description Rxed : 4                 DB Description Txed : 3
LSA Request Rxed    : 1                 LSA Request Txed    : 1
LSA Update Rxed     : 2121              LSA Update Txed     : 6156
LSA Ack Rxed        : 5962              LSA Ack Txed        : 2121
In Discards         : 0
DR RtId             : 0.0.0.0           BDR RtId            : 0.0.0.0
Restart Helper      : Both
```

```
Restart Helper Strict LSA Checking: Enabled
BFD Protection      : Off
```

```
show ospfv3 area detail

Area Identifier     : 1.0.0.0           Type                : NORM
Router ID           : 10.1.1.2          Num of Interfaces   : 1
Spf Runs            : 7                 Num ABRs            : 1
Num ASBRs           : 0                 Num DC-Bit LSAs     : 0
Num Indication LSAs : 0                 Num of DoNotAge LSAs: 0
Num LSAs            : 8                 LSA Chksum          : 0x4d0f7
Num ASBRs           : 1                 Num LSAs            : 2
Num Rtr LSAs        : 1                 Num Net LSAs        : 0
Num Inter-pref LSAs : 0                 Num Inter-rtr LSAs  : 0
Num Intra-pref LSAs : 1                 Num NSSA LSAs       : 0
LSA Chksum          : 0xbe09
Num of Nbrs         : 1                 Num of Virtual Nbrs : 1
Interfaces:
Interface Name                  Ospf State    DR ID           BDR ID
vlan101                         E    BDR      3.0.0.0         2.0.0.0
Inter-Area route Filter: none
External route Filter  : none
Configured Address Ranges:
Area: 0.0.0.1 Addr: 3100::/64 Type: 3 Advt: Yes
Addr: 3100::/64 Type: inter-prefix Advt: Yes
Addr: 3200::/64 Type: nssa         Advt: No
```

```
show ospfv3 area detail

Area Identifier     : 2.0.0.0           Type                : NSSA
Summary             : Yes               Default Metric      : 10
Translate           : Candidate (Elected)
Router ID           : 10.1.4.1          Num of Interfaces   : 1
Spf Runs            : 14                Num ABRs            : 1
Num ASBRs           : 2                 Num LSAs            : 10
Num Rtr LSAs        : 2                 Num Net LSAs        : 1
Num Inter-pref LSAs : 4                 Num Inter-rtr LSAs  : 0
Num Intra-pref LSAs : 1                 Num NSSA LSAs       : 2
LSA Chksum          : 0x3b142
Num of Nbrs         : 1                 Num of Virtual Nbrs : 0
Interfaces:
Interface Name                  Ospf State    DR ID           BDR ID
vlan400                         E    BDR      0.0.0.4         0.0.0.3
Inter-Area route Filter: none
External route Filter  : none
```

```
show ospfv3 lsdb area 0.0.0.2
                Router LSA for Area 0.0.0.2
Link State ID    ADV Router       Seq#        Age   Checksum  #Links
-------------------------------------------------------------------------
0.0.0.0          0.0.0.3          0x80000004  835   0x9b19    1
0.0.0.0          0.0.0.4          0x80000004  837   0x8431    1

                Network LSA for Area 0.0.0.2

Link State ID    ADV Router       Seq#        Age   Checksum
-------------------------------------------------------------
0.15.66.70       0.0.0.4          0x80000003  837   0x423c

                Inter Area Prefix LSA for Area 0.0.0.2

Link State ID    ADV Router       Seq#        Age   Checksum
-------------------------------------------------------------
0.0.0.2          0.0.0.3          0x80000003  829   0x734d
0.0.0.3          0.0.0.3          0x80000003  829   0x5521
0.0.0.4          0.0.0.3          0x80000003  829   0x543
0.0.0.5          0.0.0.3          0x80000003  808   0x4560

                NSSA LSA for Area 0.0.0.2

Link State ID    ADV Router       Seq#        Age   Checksum  MetricType
-------------------------------------------------------------------------
0.0.0.2          0.0.0.3          0x80000003  839   0x728f    type-1
0.0.0.8          0.0.0.4          0x80000003  898   0x5d7f    type-1

                Intra Area Prefix LSA for Area 0.0.0.2
```

```
Link State ID    ADV Router       Seq#        Age   Checksum  #Prefix  Reference
-------------------------------------------------------------------------------
0.1.0.0          0.0.0.4          0x80000005  838   0x6c9d    1        Network-LSA
```

```
show ospfv3 lsdb stats
Interface vlan100
------------------------
LSA Type           Count
------------------------
Link               2
Unknown            0

Interface v1
------------------------
LSA Type           Count
------------------------
Link               0
Unknown            0

Area ID 0.0.0.0
------------------------
LSA Type           Count
------------------------
Router             3
Network            1
Inter-Area-Prefix  7
Inter-Area-Router  1
NSSA               0
Intra-Area-Prefix  1
Unknown            0

Global
------------------------
LSA Type           Count
------------------------
AS External        1
Unknown            0
```

```
show ospfv3 lsdb stats lstype router

Area ID 0.0.0.0
------------------------
LSA Type           Count
------------------------
Router             3
Network            0
Inter-Area-Prefix  0
Inter-Area-Router  0
Intra-Area-Prefix  0
Unknown            0
```

Deleted CLI Commands

show **ospfv3 memory** {**detail** | *memoryType*}

Secure Shell (SSH) Server Upgrade

OpenSSH server listens for incoming connections. After authenticating, the server provides the client either shell access or access to the CLI, or performs a file transfer of configuration files. The server uses various services in ExtremeXOS including AAA for authentication, Policy Manager for access control, Session Manager for session reporting, and EMS for logging.

SSHServer is migrated from SSH toolkit to OpenSSH, where the SSH server is added as part of the exsshd process. ExtremeXOS 16.2 supports SSH protocol version 2 from OpenSSH. Although the SSH server is added to exsshd, the key generation is not performed by exsshd. This is done separately by another module from OpenSSH, ssh-keyGen, which is invoked from exsshd. The generated key is stored

in `/etc/ssh/ssh_host_dsa_key` and `/etc/ssh/ssh_host_dsa_key.pub`. The same format is used for any keys that are imported to OpenSSH.

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

- Keyboard interactive authentication is not supported.
- Host key algorithms are not configurable.

Resiliency Enhancement for IPv4 and IPV6 Static Routes

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

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

ExtremeXOS previously supported similar protection and resiliency using BFD on IPv4 static routes. However, BFD can only be used when the local and remote device both support BFD.

Supported Platforms

- Summit X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, and X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

New CLI Commands

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

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

enable iproute {ipv4 | ipv6} protection ping

disable iproute {ipv4 | ipv6} protection ping

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

Changed CLI Commands

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

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

ExtremeXOS Applications Environment

ExtremeXOS 16.2 introduces an environment that allows management applications, controllable through a web interface, to communicate directly with other switch management applications.

Applications are management software modules that manage, configure, or monitor specific functions within a switch. The applications leverage existing ExtremeXOS capabilities and protocols to simplify complex tasks. You may download applications to a switch independently from an ExtremeXOS release (see ezServiceablity (File Upload/Download) on page 29).

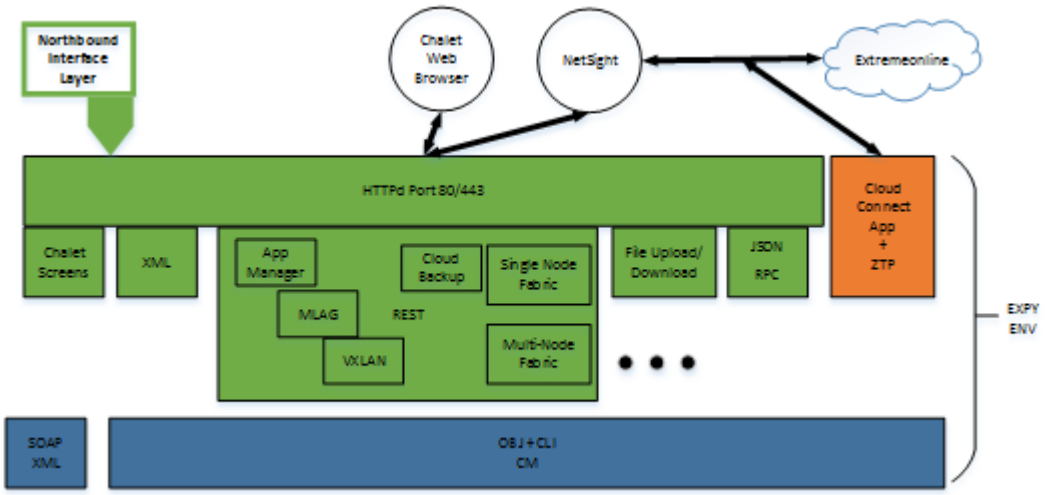

Figure 1: Application Environment Block Diagram

The HTTP interface is now a Python application based on CherryPy (3.7.0). This environment includes the following previously available interfaces:

- Web interface (Chalet)
- SOAP/XML interface

Additionally, the following new capabilities have been introduced with ExtremeXOS 16.2:

- Service applications.
- File upload/download (see ezServiceablity (File Upload/Download) on page 29)
- JSONRPC—provides a management automation interface (http://www.jsonrpc.org/specification). The JSONRPC implementation supports two methods:

    - CLI method—issues CLI commands to ExtremeXOS show commands and returns JSON data instead of formatted CLI data.
    - Python method—allows the remote system to send inline Python scripts to run on a switch. You can use inline Python scripting to perform complex tasks not available using the ExtremeXOS CLI.
- Configuration applications.
- Application manager—provides the ability to dynamically add management applications at run time. Applications may be developed independently from the ExtremeXOS release cycle.
- ezMLAG—works with Chalet web screens and peer switches. It can communicate with peer switches to perform the complex task of setting up and maintaining MLAG configurations.
- VXLAN—works with Chalet to manage VXLAN configuration coordination across multiple switches.

Supported Platforms

- Summit X430, X440, X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 series switches
- BlackDiamond X8 and 8000 series switches
- E4G-200 and E4G-400 cell site routers

ezServiceablity (File Upload/Download)

ezServiceablity is a web application that enables you to upload and download files to and from a switch instead of setting up a separate TFTP server. You can use this feature to push a new ExtremeXOS image to a switch directly when upgrading.

- The `app/file/<path>` URL provides the ability to send, retrieve, or delete files on a switch. The `<path>` parameter accepts the ExtremeXOS paths:
    - `/usr/local/cfg`
    - `/usr/local/tmp`
    - `/usr/local/ext`—Files located on a USB memory stick, if present.

    The allowed file extensions for `<path>` are: pol, cfg, xsf, py, pkt, and xml.
- The `app/file/cfg` URL is a shortcut for files in the `/usr/local/cfg` directory.

    For example, `http://<ip>/app/file/usr/local/cfg/myfile.py` is equivalent to `http://<ip>/app/file/cfg/myfile.py`. Upgrading a switch with a new ExtremeXOS image is covered using the app/upload interface. Use this interface in concert with the `app/filelist`, which provides the following capabilities:
    - Obtain the list of files on the switch.
    - Determine which file operations are supported for each file.

    This interface is useful for:
    - Sending policy, script, or config files to a switch directly from a web browser.
    - Retrieving files from a switch directly to a web browser, such as configuration files.
    - Retrieving/editing/returning files to a switch (provides a user-friendly way of editing files).
    - Deleting files on a switch.

Universal Port Management (UPM) on Summit X430 Series Switches

Support for Universal Port Management (UPM) is now extended to the Summit X430 series switches.

ExtremeXOS 16.2 Software Image Changes

The following information details changes to the ExtremeXOS 16.2 software image.

ExtremeXOS Images for BlackDiamond 8000 Series Switches

Due to additional functionality, the ExtremeXOS 16.2 and later software image is too large to download onto the BlackDiamond 8000 series switches. To resolve this issue, the diagnostics for the BlackDiamond 8900 I/O modules is now a separate image file (XMOD) in addition to the main software image file.

Table 3: BlackDiamond 8000 Series Switches Software Image Files

	Main Install image	Diagnostic image
Content	All BlackDiamond 8000 content (except BlackDiamond 8900 I/O module diagnostics)	BlackDiamond 8900 I/O module diagnostics
File Name	`bd8800-16.2.xx.yy.xos`	`bd8800-16.2.xx.yy-8900diags.xmod`
File Type	Standard ExtremeXOS image	XMOD image
Installation Notes	Other XMODs can be used with the BlackDiamond 8000 main ExtremeXOS image.	To update to a newer version of the diagnostics, you download and install the latest XMOD version. The diagnostics XMOD can be installed to the active or standby partition and diagnostics can be used immediately. There is no need to reboot or execute any other action to complete the installation.

The following scenarios produce an error or warning message:

- Attempting to run the diagnostic command on any BlackDiamond 8900 I/O module without the diagnostic image installed.
- Not having the diagnostic image installed (when system initializes).
- Installing the main BlackDiamond 8000 image without the diagnostics image present.

ExtremeXOS Images for Summit X480 Series Switches

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

Table 4: Summit X480 Series Switches Software Image Files

	Main Install image	Diagnostic image
Content	All Summit X480 content (except diagnostics)	Summit X480 diagnostics
File Name	`summitX480-16.2.xx.yy.xos`	`summitX480-16.2.xx.yy-diagnostics.xmod`
File Type	Standard ExtremeXOS image	XMOD image
Installation Notes	• Installing the main SummitX480 image over a previous release leaves the previous installation of the diagnostics image intact, as it is stored separately from the main ExtremeXOS image. You can continue to use the previously installed diagnostic version to run diagnostics. • Other Summit XMODs can be used with the Summit X480 main ExtremeXOS image.	To update to a newer version of the diagnostics, download and install the latest XMOD version. The diagnostics XMOD can be installed to the active or standby partition and diagnostics can be used immediately. There is no need to reboot or execute any other action to complete the installation.

The following scenarios produces an error or warning message:

- Attempting to run the diagnostic command without the diagnostic image installed.
- Not having the diagnostic image installed on a Summit X480 series switch or slot (when system initializes).
- Installing the main Summit X480 image without the diagnostics image present.
- Installing the general Summit image (summitX-16.2.xx.yy.xos, rather than the Summit X480-specific image) on a Summit X480 series switch.

Note

If Summit X480 series switches require rescue recovery, you can use the `summitX-16.2.xx.yy.xos` file image, and this image installs the diagnostics capability.

SSH Included in ExtremeXOS Base Image

SSH is now included in the ExtremeXOS base image starting with ExtremeXOS 16.2. A separate XMOD file is no longer required.

New Hardware Supported in ExtremeXOS 16.2

This section lists the new hardware supported in ExtremeXOS 16.2:

- BD 8800 MSM-96 module (also supported in ExtremeXOS 16.1.3)

The following transceivers are now supported in ExtremeXOS 16.2:

- 10335 40Gb ER4 QSFP+
- 10334 40Gb LM4 QSFP+
- 10329 40Gb MMF Bidirectional QSFP+
- 10325 10Gb Tunable DWDM SFP+
- 10GB-BX10-D 10Gb Bidirectional SFP+
- 10GB-BX10-U 10Gb Bidrectional SFP+

For a full list of supported platforms see *Extreme Hardware/Software Compatibility and Recommendation Matrices*.

For transceiver specifications, see *Extreme Networks Pluggable Transceivers Installation Guide*.

ExtremeXOS CLI Command Output Format Changes

The following information details format changes to the output of some ExtremeXOS 16.2 CLI commands.

VLAN Option Formatting in Commands

For commands with a **vlan_list** option, the input into this option must not contain spaces.

Example

The **enable stpd auto-bind** command VLAN ID input should be entered as:

```
enable stpd auto-bind vlan 10,20-30
```

Not:

```
enable stpd auto-bind vlan 10, 20-30
```

Output Change for Show FDB Command

The output of the **show fdb** command now accommodates longer VLAN names (32 characters long) and includes the VID, with the following additional formatting changes:

- **VID** and **Age** values are right justified.
- **Age** value leading zeros are removed.
- **Age** column is now six characters wide.

Old Output

```
Mac                     Vlan       Age  Flags      Port / Virtual Port List
------------------------------------------------------------------------------
00:00:5e:00:01:01       v123      0000  d mi       S   10
00:04:96:7e:13:7c       v123      0000  s m        S   10
00:04:96:7e:13:7c       v123      0000  s m        S   10
```

New Output

```
MAC               VLAN          VID     Age  Flags    Port/Virtual Port List
----------------- ------------- ---- ------  -------- ----------------------
00:00:00:00:00:01 32 characters  1234      0  spm      11
00:00:00:00:00:01 v123           123      46  d m    L  5
00:00:00:00:00:02 v123           123      46  d m    L  5
00:00:00:00:00:03 v123           123      46  d m    L  5
```

CLI Command Output Format of Ports Lists

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

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

Circuit Emulation Service (CES) No Longer Supported

Starting with ExtremeXOS 16.2, Circuit emulation service (CES) is no longer supported.

ExtremeXOS SSH Server Upgraded with OpenSSH v6.5

ExtremeXOS 16.1 and earlier versions generated DSA-2048 keys using `ssh keygen` provided by the SSH-Toolkit library. Starting with ExtremeXOS 16.2, ExtremeXOS generates more secure RSA-2048 keys due to switching to using the OpenSSH library, which does not support DSA-2048.

When upgrading to ExtremeXOS 16.2 and later, SSH keys generated by earlier ExtremeXOS versions (16.1 and earlier) are compatible and do *not* need to be re-generated.

Note
If a switch is downgraded from ExtremeXOS 16.2 or later to previous releases, with RSA key saved, the key becomes invalid.

OpenFlow No Longer Supported on SummitStack

For Extreme 16.2 and later, OpenFlow is not supported on SummitStack.

Extreme Hardware/Software Compatibility and Recommendation Matrices

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

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

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

Compatibility with Extreme Management Center (Formerly NetSight)

ExtremeXOS 16.2 is compatible with Extreme Management Center (formerly NetSight) version 6.3.0.182 and later.

Upgrading ExtremeXOS

For instructions about upgrading ExtremeXOS software, see "Software Upgrade and Boot Options" in the *ExtremeXOS 16.2 User Guide*.

Beginning with ExtremeXOS 16.2, an ExtremeXOS core image (.xos file) must be downloaded and installed on the alternate (non-active) partition. If you try to download to an active partition, the error message `Error: Image can only be installed to the non-active partition.` appears. An ExtremeXOS modular software package (.xmod file) can still be downloaded and installed on either the active or alternate partition.

Downloading Supported MIBs

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

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

Tested Third-Party Products

This section lists the third-party products tested for ExtremeXOS 16.2.

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

2 Limits

This chapter summarizes the supported limits in ExtremeXOS 16.2.

Table 5: Supported Limits on page 37 summarizes tested metrics for a variety of features, as measured in a per-system basis unless otherwise noted. These limits may change, but represent the current status. The contents of this table supersede any values mentioned in the ExtremeXOS books.

Note

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

The scaling and performance information shown in Table 5: Supported Limits on page 37 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 5: Supported Limits on page 37 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 5: Supported Limits

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

Table 5: Supported Limits (continued)

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

Table 5: Supported Limits (continued)

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

Table 5: Supported Limits (continued)

Metric	Product	Limit
AVB (audio video bridging)—maximum number of active streams. Note: * It is recommended that you do not use on more than 8 ports on this switch.	Summit X440, X460, X460-G2, X450-G2 Summit X670, X670-G2 Summit X430	1,024 4,096 100*
BFD sessions—maximum number of software BFD sessions.	All platforms (default timers—1 sec) BlackDiamond X8 and 8800 (minimal timers—50 msec) All Summits, except X460-G2 (minimal timers—100 msec) Summit X460-G2	512 10 [c] 10 [c] 900 (if PTP feature not enabled) 425 (with PTP enabled) 256 (with 3 ms transmit interval)
BGP (aggregates)—maximum number of BGP aggregates.	All platforms (except E4G-200, X430, X440, and X450-G2) with Core license or higher	256
BGP (networks)—maximum number of BGP networks.	All platforms (except E4G-200, X430, X440, and X450-G2) with Core license or higher BlackDiamond X8 series	1024
BGP (peers)—maximum number of BGP peers. Note: *With default keepalive and hold timers.	BlackDiamond X8 series, xl-series, 8000 series All Summits, except X450-G2, X480, X440, X430, E4G-200 E4G-400H Summit X480	512 128* 128* 512
BGP (peer groups)—maximum number of BGP peer groups.	BlackDiamond 8900 series BlackDiamond 8800 BlackDiamond X8 series Summit X480 Summit X770, X670-G2, X670v-48t, X670, X460-G2, X460 (with Core license or higher)	128 64 128 128 64
BGP (policy entries)—maximum number of BGP policy entries per route policy.	All platforms (except E4G-200, X430, X440, and X450-G2) with Core license or higher	256
BGP (policy statements)—maximum number of BGP policy statements per route policy.	All platforms (except E4G-200, X430, X440, and X450-G2) with Core license or higher	1,024
BGP multicast address-family routes—maximum number of multicast address-family routes.	BlackDiamond 8900 xl-series BlackDiamond X8 series BlackDiamond X8-xl series Summit X480 E4G-400 Summit X460, X460-G2, X670, X670-G2, X770	524,256 (up to) [b] 1,048,544 (up to) [i] 524,256 (up to) [b] 25,000 25,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
BGP (unicast address-family routes)—maximum number of unicast address-family routes.	BlackDiamond 8900 xl-series BlackDiamond X8 series BlackDiamond X8-xl series Summit X480 E4G-400 Summit X460, X460-G2, X670, X670-G2, X770	524,256 (up to)[b] 1,048,544 (up to)[i] 524,256 (up to)[b] 25,000 25,000
BGP (non-unique routes)—maximum number of non-unique BGP routes.	BlackDiamond 8900 xl-series, X8 xl-series BlackDiamond X8 series BlackDiamond X8 xl-series Summit X460, X460-G2, X670, X670-G2, X770 Summit X480 E4G-400	1,200,000 24,000 2,000,000 25,000 1,000,000 25,000
BGP ECMP—maximum number of equalcost multipath for BGP and BGPv6.	All platforms, except Summit X430, X440, and E4G-200 BlackDiamond 8800 G48Te2 (for BGPv6) Summit X450-G2	2, 4, or 8 N/A N/A
BGPv6 (unicast address-family routes)—maximum number of unicast address family routes.	BlackDiamond 8900 xl-series, BlackDiamond X8 xl-series BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond X8 series Summit X460, X460-G2 Summit X480 Summit X670, X670-G2, X770 E4G-400	20,000 6,000 240 8,000 6,000 20,000 8,000 6,000
BGPv6 (non-unique routes)—maximum number of non-unique BGP routes.	BlackDiamond 8900 xl-series BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond X8 series, X8 xl-series Summit X460, X460-G2 Summit X670, X670-G2, X770 E4G-400	24,000 18,000 720 24,000 18,000 24,000 18,000
BOOTP/DHCP relay—maximum number of BOOTP or DHCP servers per virtual router.	All platforms, except Summit X430	4
BOOTP/DHCP relay—maximum number of BOOTP or DHCP servers per VLAN.	All platforms, except Summit X430	4
Connectivity fault management (CFM)—maximum number or CFM domains. Note: With Advanced Edge license or higher.	All platforms	8
CFM—maximum number of CFM associations. Note: With Advanced Edge license or higher.	All platforms	256

Table 5: Supported Limits (continued)

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

Table 5: Supported Limits (continued)

Metric	Product	Limit
DHCPv6 Prefix Delegation Snooping—Maximum number of DHCPv6 prefix delegation snooped entries.	All platforms	256 (with Underlying Protocol Ripng) 128 (with Underlying protocol OSPFv3) 1,024 (with static routes)
DHCP snooping entries—maximum number of DHCP snooping entries.	All Summits BlackDiamond X8	2,048 6,000
Dynamic ACLs—maximum number of ACLs processed per second. Note: Limits are load dependent.	Summit X480, X670 with 50 DACLs with 500 DACLs	 10 5
	BlackDiamond X8 BlackDiamond 8800	N/A N/A
EAPS domains—maximum number of EAPS domains. Note: An EAPS ring that is being spatially reused cannot have more than four configured EAPS domains.	BlackDiamond 8000 series, X8 series Summit X670-G2, X450-G2, and X770 Summit X670, X480, X460, X460-G2, X440, E4G-200, E4G-400 Summit X430	64 64 32 4
EAPSv1 protected VLANs—maximum number of protected VLANs.	BlackDiamond 8000 series, X8 series Summit series, E4G-200, E4G-400	2,000 1,000
EAPSv2 protected VLANs—maximum number of protected VLANs.	BlackDiamond 8000 series, X8 series All Summits (except X430, X440), E4G-200, E4G-400	2,000 500
ELSM (vlan-ports)—maximum number of VLAN ports.	BlackDiamond 8000 series, X8 series All Summits, E4G-200, E4G-400	5,000
ERPS domains—maximum number of ERPS domains without CFM configured.	BlackDiamond 8800 series, X8 series Summit series (except X430), E4G-200, E4G-400 Summit X430	32 32 4
ERPS domains—maximum number of ERPS domains with CFM configured.	BlackDiamond 8800 series, X8 series E4G-200, E4G-400 Summit X460, X460-G2 Summit X430 Summit X440, X770, X670, X670-G2, X480, X450-G2	16 32 32 4 16
ERPSv1 protected VLANs—maximum number of protected VLANs.	BlackDiamond 8800 series, X8 series All Summits, E4G-200, E4G-400	2,000 1,000
ERPSv2 protected VLANs—maximum number of protected VLANs.	BlackDiamond 8800 series, X8 series All Summits (except X430), E4G-200, E4G-400	2,000 500

Table 5: Supported Limits (continued)

Metric	Product	Limit
ESRP groups—maximum number of ESRP groups.	All platforms	7
ESRP domains—maximum number of ESRP domains.	All platforms	64
ESRP VLANs—maximum number of ESRP VLANs.	BlackDiamond 8800 BlackDiamond X8 All Summits E4G-200. E4G-400	1,000 2,048 1,000 1,000
ESRP (maximum ping tracks)—maximum number of ping tracks per VLAN.	All platforms (except Summit X430)	8
ESRP (IP route tracks)—maximum IP route tracks per VLAN.	All platforms (except Summit X430)	8
ESRP (VLAN tracks)—maximum number of VLAN tracks per VLAN.	All platforms (except Summit X430)	1
Forwarding rate—maximum L3 software forwarding rate.	BlackDiamond 8000 series BlackDiamond X8 series Summit X770 Summit X670-G2 Summit X670 Summit X480 Summit X460-G2 Summit X460 Summit X450-G2 Summit X440 E4G-200 E4G-400	10,000 pps 20,000 pps 11,000 pps 21,000 pps 14,829 pps 14,509 pps 25,000 pps 5,222 pps 24,000 pps 5,418 pps 8,718 pps 5,536 pps
FDB (unicast blackhole entries)—maximum number of unicast blackhole FDB entries.	BlackDiamond 8900 series 8900 c-series 8900 xl-series 8900-40G6X-xm	 32,000 524,288 (up to)[b] 128,000
	BlackDiamond 8000 e-series BlackDiamond 8800 c-series BlackDiamond X8 a-series modules BlackDiamond X8-100G4X modules. BlackDiamond X8 xl-series module E4G-200, E4G-400 Summit X440, X430 Summit X480 Summit X460 Summit X460-G2 Summit X670 VIM4-40G4x, X480 VIM3-40G4X Summit X770, X670-G2 Summit X670, X670v-48t Summit X450-G2	8,000 32,000 128,000 384,000 384,000[d] 32,000 16,000 524,288 (up to)[b] 32,000 49,152[e] 128,000 294,912 130,000[e] 34,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
FDB (multicast blackhole entries)—maximum number of multicast blackhole FDB entries.	BlackDiamond 8000 series, X8 series Summit X480, X460-G2, X460, X440, X430, X450-G2 Summit X770, X670, X670-G2, X670v-48t, X480 VIM3-40G4X E4G-200, E4G-400	1,024 1,024 4,096 1,024
FDB (maximum L2 entries)—maximum number of MAC addresses.	BlackDiamond 8000 c-series BlackDiamond 8000 e-series BlackDiamond 8000 (system), except 8900 xl-series BlackDiamond 8900 xl-series BlackDiamond X8 a-series modules BlackDiamond X8-100G4X modules BlackDiamond X8 xl-series E4G-200, E4G-400 Summit X440, X430 Summit X480 Summit X480 (40G4X) Summit X460 Summit X670-G2 Summit X460-G2 Summit X670 Summit X770 Summit X450-G2	32,768[f] 8,192[f] 128,000[f] 524,488 (up to)[b] 128,000[f] 384,000[f] 1,048,576 (up to)[b g] 32,000[f] 16,000[f] 524,488 (up to)[b f] 128,000[b f] 32,000[f] 294,912[f] 98,300[f] 128,000[e f] 294,912[f] 68,000
FDB (Maximum L2 entries)—maximum number of multicast FDB entries.	BlackDiamond X8, 8800 Summit X770, X670, X670-G2 Summit X480, X460, X460-G2, X430, X440, X450-G2 E4G-200, E4G-400	1,024 4,096 1,024 1,024
FIP Snooping VLANs	BlackDiamond X8 BlackDiamond 8800 (8900-40G6X-c only) Summit X670	768
FIP Snooping Virtual Links (FPMA mode) per port group	BlackDiamond X8 BlackDiamond 8800 (8900-40G6X-c only) Summit X670	1,908
FIP Snooping FCFs (with perimeter port) per port group	BlackDiamond X8 BlackDiamond 8800 (8900-40G6X-c only)	238
FIP Snooping FCFs (with Enode-to-FCF port)	BlackDiamond X8 BlackDiamond 8800 (8900-40G6X-c only) Summit X670	212
Identity management—maximum number of Blacklist entries.	All platforms, except Summit X430 Summit X430	512 N/A
Identity management—maximum number of Whitelist entries.	All platforms, except Summit X430 Summit X430	512 N/A
Identity management—maximum number of roles that can be created.	All platforms, except Summit X430 Summit X430	64 N/A

Table 5: Supported Limits (continued)

Metric	Product	Limit
Identity management—maximum role hierarchy depth allowed.	All platforms, except Summit X430 Summit X430	5 N/A
Identity management—maximum number of attribute value pairs in a role match criteria.	All platforms, except Summit X430 Summit X430	16 N/A
Identity management—maximum of child roles for a role.	All platforms, except Summit X430 Summit X430	8 N/A
Identity management—maximum number of policies/dynamic ACLs that can be configured per role.	All platforms, except Summit X430 Summit X430	8 N/A
Identity management—maximum number of LDAP servers that can be configured.	All platforms, except Summit X430 Summit X430	8 N/A
Identity management—maximum number of Kerberos servers that can be configured.	All platforms, except Summit X430 Summit X430	20 N/A
Identity management—maximum database memory-size.	All platforms, except Summit X430 Summit X430	512 N/A
Identity management—recommended number of identities per switch. Note: Number of identities per switch is for a default identity management database size (512 Kbytes) across all platforms.	All platforms, except Summit X430 Summit X430	100 N/A
Identity management—recommended number of ACL entries per identity. Note: Number of ACLs per identity based on system ACL limitation.	All platforms, except Summit X430 Summit X430	20 N/A
Identity management—maximum number of dynamic ACL entries configured as an individual dynamic rule, or as an ACL entry in a policy file.	All platforms, except Summit X430 Summit X430	500 N/A

Table 5: Supported Limits (continued)

Metric	Product	Limit
IGMP snooping per VLAN filters—maximum number of VLANs supported in per-VLAN IGMP snooping mode.	BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 a-series modules BlackDiamond X8 b-series modules E4G-200, E4G-400 Summit X440 Summit X460, X670, X440 Summit X460-G2 Summit X450-G2 Summit X480 Summit X770, X670-G2	2,000 448 1,000 4,000 1,000 1,000 4,000 1,000 448 1,000 1,500 2,048 4,000 2,000
IGMPv1/v2 SSM-map entries—maximum number of IGMPv1/v2 SSM mapping entries.	All platforms	500
IGMPv1/v2 SSM-MAP entries—maximum number of sources per group in IGMPv1/v2 SSM mapping entries.	All platforms	50
IGMPv2 subscriber—maximum number of IGMPv2 subscribers per port. [n]	BlackDiamond 8800 c-series, 8900 c-series, X8 series Summit X430, X460, E4G-200, E4G-400, X440 Summit X480, X670, X670v-48t Summit X770, X670-G2, X460-G2, X450-G2	2,000 1,000 2,000 4,000
IGMPv2 subscriber—maximum number of IGMPv2 subscribers per switch. [n]	BlackDiamond 8800 c-series, 8900 c-series, X8 series Summit X430, X440, E4G-200 Summit X770, X670-G2 Summit X460, X460-G2, X480, X670, E4G-400, X670v-48t, X450-G2	20,000 10,000 30,000 20,000
IGMPv3 maximum source per group—maximum number of source addresses per group.	All platforms	250
IGMPv3 subscriber—maximum number of IGMPv3 subscribers per port. [n]	BlackDiamond 8800 e-series BlackDiamond 8800 c-series BlackDiamond 8900 series BlackDiamond X8 series Summit X480, X670, X670v-48t, E4G-200, X440 Summit X770, X670-G2, X460-G2, X450-G2 Summit X460, E4G-400	1,000 2,000 5,000 3,000 1,000 4,000 2,000
IGMPv3 subscriber—maximum number of IGMPv3 subscribers per switch. [n]	BlackDiamond 8800 e-series BlackDiamond 8800 c-series BlackDiamond 8900 series BlackDiamond X8 series Summit X670, X670v-48t, X480, E4G-200, X440 Summit X460, X460-G2, E4G-400, X450-G2 Summit X770, X670-G2	10,000 20,000 30,000 20,000 10,000 20,000 30,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
IP ARP entries in software—maximum number of IP ARP entries in software. Note: May be limited by hardware capacity of FDB (maximum L2 entries).	BlackDiamond X8-100G4X modules Summit X670-G2, X770 Summit X670, X480, X460, X440, X430 Summit X460-G2 Summit X450-G2 E4G-200, E4G-400	229,374 (up to)[h] 131,072 (up to)[h] 20,480[h] 57,344 (up to)[h] 47,000 (up to)[h] 20,480
IP ARP entries in software with distributed mode on—maximum number of IP ARP entries in software with distributed mode on.	BlackDiamond 8000 series with 8900-MSM128, MSM-48c, or MSM-96 and only 8900 xl-series I/O modules BlackDiamond 8000 series with any I/O modules that are not 8900 xl-series BlackDiamond X8 series All other platforms	260,000 100,000 172,000 N/A
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 Per BlackDiamond 8900-G48X-xl or 8900-G48T-xl, up to 130,000 per system Per BlackDiamond 8000 c-series, up to 18,000 per system Per BlackDiamond 8900-40G6X-xm, up to 22,000 per system Per BlackDiamond X8 a-series, up to 28,000 per system Per BlackDiamond X8 xl-series, up to 172,000 per system All other platforms	32,500[b] 16,250[b] 8,000 8,000 12,000 172,000 N/A
IPv4 ARP entries in hardware with minimum LPM routes—maximum recommended number of IPv4 ARP entries in hardware, with minimum LPM routes present. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series switches, assumes number of IP route reserved entries is 100 or less.	BlackDiamond 8800 c-, xm-series BlackDiamond 8000 e-series BlackDiamond 8900 xl-series BlackDiamond X8 a-series BlackDiamond X8-100G4X modules BlackDiamond X8 xl-series E4G-200 E4G-400 Summit X440 Summit X670, X480 (40G4X) Summit X460, X480 Summit X460-G2 Summit X770, X670-G2 Summit X450-G2	8,000 1,000[i] 16,000 16,000 182,000 (up to)[h m] 294,000 (up to)[i] 8,000 16,000 412 8,000 16,000 50,000 (up to)[h] 108,000 (up to)[h] 39,000 (up to)[h]
IPv4 ARP entries in hardware with maximum LPM routes—maximum recommended number of IPv4 ARP entries in hardware, with maximum LPM routes present. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series, assumes number of IP route reserved entries is “maximum.”	BlackDiamond 8800 c-, xm-series BlackDiamond 8000 e-series BlackDiamond 8900 xl-series BlackDiamond X8 a-series BlackDiamond X8-100G4X modules BlackDiamond X8 xl-series E4G-200 E4G-400 Summit X440 Summit X460, X480 Summit X670, X480 VIM3-40G4X Summit X770, X670-G2 Summit X460-G2 Summit X450-G2	6,000[i] 500[i] 12,000[i] 12,000[i] 172,000 (up to)[h j] 290,000 (up to)[i] 6,000[i] 12,000[i] 380 12,000[i] 6,000[i] 98,000 (up to)[h] 43,000 (up to)[h] 29,000 (up to)[h]

Table 5: Supported Limits (continued)

Metric	Product	Limit
IP flow information export (IPFIX)—number of simultaneous flows.	BlackDiamond 8900 xl-series modules BlackDiamond 8900 c-series modules BlackDiamond X8 b-series modules Summit X460-24t/x/p, X460-G2 Summit X480, X460-48t/x/p E4G-400	4,096 ingress, 4,096 egress 4,096 ingress, 4,096 egress 2,048 ingress, 2,048 egress 2,048 ingress, 2,048 egress 4,096 ingress, 4,096 egress 2,048 ingress, 2,048 egress
IPv4 remote hosts in hardware with zero LPM routes—maximum recommended number of IPv4 remote hosts (hosts reachable through a gateway) in hardware when LPM routing is not used. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series, assumes number of IP route reserved entries is 0, and number of IPv4 ARP entries present is 100 or less.	BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 a-series BlackDiamond X8-100G4X andX8 xl-series E4G-200 E4G-400 Summit X440 Summit X460 Summit X460-G2 Summit X480 Summit X670, X480 VIM3-40G4X Summit X770, X670-G2 Summit X450-G2	18,000[i] 1,000[i] 40,000[b] 22,000[i] 28,000[i] 311,000 (up to)[h] 18,000[i] 20,000[i] 448 20,000[i] 73,000[h] 40,000[b] 22,000[i] 176,000 (up to)[h] 61,000 (up to)[h]
IPv4 routes—maximum number of IPv4 routes in software (combination of unicast and multicast routes).	BlackDiamond 8900 xl-series with 8900-MSM128, MSM-48c, or MSM-96 All other BlackDiamond 8000 series hardware BlackDiamond X8 series BlackDiamond X8 with BDX X8 xl-series Summit X440 Summit X460, X670, X770, X670-G2, X460-G2, X450-G2 Summit X480 E4G-200, E4G-400	524,256 (up to)[b] 25,000 25,000 1,048,544 (up to)[i] 256 25,000 524,256 (up to)[b] 25,000
IPv4 routes (LPM entries in hardware)— number of IPv4 routes in hardware.	BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 series BlackDiamond BDX X8 xl-series E4G-200, E4G-400 Summit X440 Summit X460, X460-G2 Summit X480 Summit X480 VIM3-40G4X Summit X670 Summit X770, X670-G2, X450-G2	12,000 480 524,256 (up to)[b] 16,000[e] 16,000[e] 1,048,544 (up to) 12,000 32 12,000 524,256 (up to)[b o] 16,000[o] 12,000 16,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
IPv6 addresses on an interface—maximum number of IPv6 addresses on an interface.	All platforms	255
IPv6 addresses on a switch—maximum number of IPv6 addresses on a switch.	BlackDiamond 8000 series BlackDiamond X8 series E4G-200, E4G-400 Summit X440 Summit X460, X480 Summit X770, X670, X670-G2, X460-G2, X450-G2	512 2,048 512 254 512 2,048
IPv6 host entries in hardware—maximum number of IPv6 neighbor entries in hardware.	BlackDiamond 8800 c-, xm-series BlackDiamond 8000 e-series BlackDiamond 8900-10G24X-c modules BlackDiamond 8900-G96T-c modules BlackDiamond 8900 xl-series BlackDiamond X8 a-series BlackDiamond X8-100G4X BlackDiamond X8 xl-series E4G-200 E4G-400 Summit X440 Summit X460, X670, X480 VIM3-40G4X Summit X770, X670-G2 Summit X480, X670v-48t Summit X460-G2 Summit X450-G2	3,000[i] 250[i] 2,000[i] 4,000[i] 8,192 (up to)[b i] 3,000[i] 49,000[i h] 49,000[i l] 2,000[i] 3,000[i] 192[i] 3,000[i] 36,750[i] 6,000[i] 22,000[i] 12,000[i]
IPv6 routes (LPM entries in hardware)—maximum number of IPv6 routes in hardware.	BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 xm-series BlackDiamond 8900 xl-series BlackDiamond X8 series BlackDiamond X8 xl-series E4G-200, E4G-400 Summit X440 Summit X460, X460-G2 Summit X480 Summit X670, X480 (VIM3-40G4X), X670-G2, X770, X450-G2	6,000 240 8,000 245,760 (up to)[b] 8,000 524,288 (up to)[l] 6,000 16 6,000 245,760 (up to)[b] 8,000
IPv6 routes with a mask greater than 64 bits in hardware—maximum number of such IPv6 LPM routes in hardware.	BlackDiamond 8000 c-, e-, xm-series BlackDiamond 8000 xl-series BlackDiamond X8 series BlackDiamond X8 xl-series E4G-200, E4G-400 Summit X480 Summit X440, X460, X460-G2, X670, X670-G2, X770, X480 (VIM3-40G4X), X450-G2	256 245,760 (up to)[k] 256 524,288 (up to)[l] 256 245,760 (up to)[k] 256

Table 5: Supported Limits (continued)

Metric	Product	Limit
IPv6 route sharing in hardware—route mask lengths for which ECMP is supported in hardware. Note: * >64 single path only	Summit X460, X480, X670, X670V-48t E4G-200, E4G-400 BlackDiamond 8800 (all I/O modules, except G48Te2) Summit X460-G2, X670-G2, X770 BlackDiamond X8 a-series BlackDiamond X8-100G4X modules BlackDiamond X8 xl-series, Summit X450-G2 Summit X440, X430 BlackDiamond 8800 G48Te2	0–128 0–128 0–128 0–64 * 0–128 0–64 * 0–128[l] N/A N/A
IPv6 routes in software—maximum number of IPv6 routes in software.	BlackDiamond 8900 xl-series with 8900-MSM128, MSM-48c, or MSM-96 All other BlackDiamond 8000 series hardware BlackDiamond X8 series BlackDiamond X8 with xl-series Summit X460, X460-G2, X670, X670-G2, X770, X450-G2, E4G-200, E4G-400 Summit X480 Summit X440	245,760 (up to)[k] 25,000 25,000 524,288 (up to)[l] 25,000 245,760 (up to)[k] 256
IP router interfaces—maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.	Summit X460-G2, X670, X770, X670-G2, X450-G2, and BlackDiamond X8 BlackDiamond 8800 Summit X440 Summit X480, X460 E4G-200, E4G-400	2,048 512 254 512 512
IP multicast static routes—maximum number of permanent multicast IP routes.	All platforms (except Summit X430, X440) Summit X430, X440	1,024 32
IP unicast static routes—maximum number of permanent IP unicast routes.	All platforms (except Summit X430, X440) Summit X430, X440	1,024 32
IP route sharing (maximum gateways)—Configurable maximum number of gateways used by equal cost multipath OSPF, BGP, IS-IS, static routes, or L2VPNs. Routing protocol OSPF is limited to 16 ECMP gateways per destination. Routing protocols BGP and IS-IS are limited to 8 ECMP gateways per destination. Static routes are limited to 64 next-hops. L2VPNs are limited to 16 LSPs per pseudowire on platforms that support 32 gateways, and 64 LSPs per pseudowire on platforms that support 64 gateways.	All platforms, except Summit X430, X440, X670, and BlackDiamond X8 Summit X670, BlackDiamond X8 Summit X430, X440 BlackDiamond 8800 G48Te2 (for IPv6)	2, 4, 8, 16, 32 2, 4, 8, 16, or 32, or 64 N/A N/A

Table 5: Supported Limits (continued)

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

Table 5: Supported Limits (continued)

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

Table 5: Supported Limits (continued)

Metric	Product	Limit
IS-IS routers in an area—recommended maximum number of IS-IS routers in an area.	Summit X480 All other platforms, except Summit X440, X430	128 256
IS-IS route origination—recommended maximum number of routes that can be originated by an IS-IS node.	BlackDiamond 8000 series, BlackDiamond X8 series BlackDiamond X8 xl-series, 8900 xl-series E4G-400 E4G-200 Summit X450-G2, X460, X460-G2, X670, X670-G2, X770, X480	20,000 30,000 25,000 20,000 20,000
IS-IS IPv4 L1 routes in an L1 router—recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router.	BlackDiamond 8000 series, BlackDiamond X8 series BlackDiamond X8 xl-series, 8900 xl-series Summit X480 Summit X450-G2, X460, X460-G2, X670, X670-G2, X770 E4G-200, E4G-400	25,000 120,000 50,000 25,000 25,000
IS-IS IPv4 L2 routes—recommended maximum number of IS-IS Level 2 routes.	BlackDiamond 8000 series BlackDiamond X8 series BlackDiamond X8 xl-series, 8900 xl-series Summit X480 Summit X450-G2, X460, X460-G2, X670, X670-G2, X770 E4G-200, E4G-400	20,000 25,000 120,000 50,000 25,000 25,000
IS-IS IPv4 L1 routes in an L1/L2 router—recommended maximum number of IS-IS Level 1 routes in an L1/L2 IS-IS router.	BlackDiamond 8000 series, X8 series, 8900 xl-series E4G-200, E4G-400 Summit X450-G2, X460, X460-G2, X480, X670, X670-G2. X770	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, X8 series BlackDiamond X8 xl-series, 8900 xl-series Summit X480 Summit X450-G2, X460, X460-G2, X670, X670-G2, X770, E4G-400	10,000 40,000 25,000 10,000
IS-IS IPv6 L2 routes—recommended maximum number of IS-IS Level 2 routes.	BlackDiamond 8000 series, X8 series BlackDiamond X8 xl-series, 8900 xl-series Summit X480 Summit X450-G2, X460, X460-G2, X670, X670-G2, X770 E4G-200, E4G-400	10,000 40,000 15,000 10,000 10,000
IS-IS IPv6 L1 routes in an L1/L2 router—recommended maximum number of IS-IS Level 1 routes in a L1/l2 router.	BlackDiamond 8000 series, X8 series BlackDiamond X8 xl-series, 8900 xl-series Summit X480 Summit X450-G2, X460, X460-G2, X670, X670-G2, X770, E4G-400	10,000 15,000 15,000 10,000
IS-IS IPv4/IPv6 L1 routes in an L1 router—recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	BlackDiamond 8000 series, X8 series BlackDiamond X8 xl-series, 8900 xl-series Summit X480 Summit X450-G2, X460, X460-G2, X670, X670-G2. X770 E4G-200, E4G-400	20,000 60,000 40,000 20,000 20,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
IS-IS IPv4/IPv6 L2 routes in an L2 router—recommended maximum number of IS-IS Level 2 routes in a Level 2 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	BlackDiamond 8000 series, X8 series BlackDiamond X8 xl-series, 8900 xl-series Summit X480 Summit X450-G2, X460,X460-G2, X670, X670-G2, X770 E4G-200, E4G-400	20,000 60,000 40,000 20,000 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 BlackDiamond X8 series BlackDiamond 8900 xl-series Summit X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 E4G-200, E4G-400	20,000
Jumbo frames—maximum size supported for jumbo frames, including the CRC.	All platforms	9,216
L2 VPN: VCCV (pseudowire Virtual Circuit Connectivity Verification) VPNs per switch—maximum number of VCCV enabled VPLS VPNs.	All platforms, except Summit X440, X430, and X450-G2	16
L2 VPN: VPLS MAC addresses—maximum number of MAC addresses learned by a switch.	BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 a-series modules BlackDiamond X8-100G4X modules BlackDiamond X8 xl-series E4G-200, E4G-400 Summit X460 Summit X480 Summit X670, Summit X670V-48t, Summit X770 Summit X480 (40G VIM) Summit X670-G2 Summit X460-G2 Summit X450-G2	512,000 128,000 128,000 384,000 1,048,576 [g] 32,000 32,000 512,000 128,000 121,000 140,000 55,000 N/A
L2 VPN: VPLS VPNs—maximum number of VPLS virtual private networks per switch.	BlackDiamond 8900 xl-series BlackDiamond 8900-40G6x-xm BlackDiamond X8 series E4G-200, E4G-400 Summit X460, X460-G2, X480, X670, X670V-48t, X480 (40G VIM), X770, X670-G2	1,023
L2 VPN: VPLS peers—maximum number of VPLS peers per VPLS instance.	BlackDiamond 8900 xl-series, 8900-40G6x-xm, X8 series Summit X770, X670-G2, X670v-48t, X480, X460-G2 Summit X670, X460 E4G-200, E4G-400	64 64 32 32

Table 5: Supported Limits (continued)

Metric	Product	Limit
L2 VPN: LDP pseudowires—maximum number of pseudowires per switch.	BlackDiamond 8900 xl-series, X8 series BlackDiamond 8900-40G6X-xm E4G-200, E4G-400 Summit X770 Summit X670-G2, X670v-48t, X480 Summit X670 Summit X460-G2 Summit X460 Summit X450-G2	7,000 3,000 1,000 7,800 7,000 3,000 7,116 1,000 N/A
L2 VPN: static pseudowires—maximum number of static pseudowires per switch.	BlackDiamond 8900 xl-series, X8 series BlackDiamond 8900-40G6X-xm Summit X460, X480, X670V-48t Summit X770 Summit X480-40G, Summit X670 Summit X670-G2, X460-G2 E4G-200 E4G-400 Summit X450-G2	7,116 3,020 7,116 15,308 3,020 7,000 2,764 6,860 N/A
L2 VPN: Virtual Private Wire Service (VPWS) VPNs—maximum number of virtual private networks per switch.	Summit X460 Summit X480, X770 Summit X480-40G VIM, X670 Summit X670V-48t BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 series Summit X670-G2 Summit X460-G2 E4G-200, E4G-400 Summit X450-G2	1,000 4,000 2,047 4,000 4,000 2,047 4,000 4,090 1,023 1,000 N/A
Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode. Note: • The internal lookup table configuration used is "l2-and-l3". • IPv6 and IPv4 L2 IPMC scaling is the same for this mode. • Layer-2 IPMC forwarding cache limits — (IGMP/MLD/PIM snooping) in mixed-mode are the same.	BlackDiamond 8800 e-series switches BlackDiamond 8800 c- and xl-series switches BlackDiamond 8800 xm-series switches BlackDiamond X8 series switches E4G-200, E4G-400 Summit X480, X460 Summit X670, X670V Summit X440 Summit X770, X670-G2 Summit X460-G2 Summit X430 Summit X450-G2	2,000 8,000 15,000 15,000 8,000 8,000 15,000 5,000 73,000 24,000 5,000 14,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
Layer-3 IPv4 Multicast—maximum number of <S,G,V> entries installed in the hardware (IP multicast compression enabled). Note: • Limit value same for MVR senders , PIM Snooping entries. PIM SSM cache,IGMP senders, PIM cache. • The internal lookup table configuration used is ""more l3-and-ipmc". • Assumes source-group-vlan mode as look up key. • Layer 3 IPMC cache limit in mixed mode also has the same value. "	BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 a-series BlackDiamond X8-100G4X and X8 xl-series E4G-200 E4G-400 Summit X440 Summit X480 Summit X460 Summit X670 Summit X770, X670-G2 Summit X450-G2 Summit X460-G2	6,000 500 6,000 12,000 3,000 6,000 59,000 3,000 6,000 192 12,000 6,000 3,000 77,500 21,000 26,000
Layer-3 IPv6 Multicast—maximum number of <S,G,V> entries installed in the hardware (IP multicast compression enabled). Note: • Limit value same for MLD sender per switch,PIM IPv6 cache. • The internal lookup table configuration used is ""more l3-and-ipmc". • Assumes source-group-vlan mode as look up key.	BlackDiamond 8800 c-series BlackDiamond 8800 e-series BlackDiamond 8900 c-series BlackDiamond 8900-40G6X-xm BlackDiamond 8900 xl-series BlackDiamond X8 a-series BlackDiamond X8-100G4X and X8 xl-series E4G-200 E4G-400 Summit X440 Summit X460, X480, X670 Summit X770, X670-G2 Summit X450-G2 Summit X460-G2	1,000 250 1,000 1,000 3,000 3,000 30,000 1,500 3,000 90 3,000 30,000 10,000 14,000

Table 5: 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) With distributed IP ARP mode on	 128 64
	BlackDiamond 8000 series 8900-40G6X-xm using address-based custom algorithm With distributed IP ARP mode off (default) With distributed IP ARP mode on	 128 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) With distributed IP ARP mode on	 127 63
	SummitStack with X670 with L2, L3 or L3_L4 algorithm configured for any group	127
	All other SummitStack configurations and Summit series switches	128
	BlackDiamond X8 series using address-based custom algorithm With distributed IP ARP mode off (default) With distributed IP ARP mode on	 384 384
	BlackDiamond X8 series with L2, L3 or L3_L4 algorithm configured for any group With distributed IP ARP mode off (default) With distributed IP ARP mode on	 127 63
Load sharing—maximum number of ports per load-sharing group. Note: *For custom algorithm ** For L2 and L3 algorithms Note: For a mix of Summit X770 and Summit X670 series switches in a stack, the limits are the Summit X670 limits.	BlackDiamond X8 series Summit X460-G2 (standalone)	64 32
	Summit X670 (standalone)	32 * 16 **
	Summit X670 (stacked) Summit X670-G2 (stacked)	64 * 16 **
	Summit X770 (standalone) Summit X670-G2 (standalone) Summit X460-G2 (standalone) Summit X450-G2 (standalone)	32
	Summit X770 (stacked) Summit X670-G2 (stacked) Summit X460-G2 (stacked) Summit X450-G2 (stacked)	64
	All other Summit series, SummitStacks, E4G cell site routers, and BlackDiamond 8000 series switches	8
Logged messages—maximum number of messages logged locally on the system.	All platforms	20,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
MAC address learning rate—hardware learning rat.	E4G-200	22 msec
MAC-based security—maximum number of MAC-based security policies.	All platforms	1,024
MAC Locking—Maximum number of MAC locking stations that can be learned on a port.	All platforms	64 (static MAC locking stations) 600 (first arrival MAC locking stations)
Meters—maximum number of meters supported.	All platforms	2,048
Maximum mirroring instances Note: The Summit X430 can only support one egress mirroring instance.	All platforms Note: Only two or four mirroring instance will be active at a time depending on the mirroring filter added to it. There are four hardware resource slots. Each single instance uses one such slot, while each ingress plus egress instance uses two slots. So this allows the you to use a total of four slots, while there are no more then two egress instances. The maximum possible combination for mirroring instances: 1 4 ingress 2 3 ingress + 1 egress 3 2 ingress + 2 egress 4 2 (ingress + egress) 5 1 (ingress + egress) + 2 ingress 6 1 (ingress + egress) + 1 egress + 1 ingress	16 (including default mirroring instance)
Mirroring (filters)—maximum number of mirroring filters. Note: This is the number of filters across all the active mirroring instances.	BlackDiamond 8000 series BlackDiamond X8 series All Summit series E4G cell site routers	128
Mirroring, one-to-many (filters)—maximum number of one-to-many mirroring filters. Note: This is the number of filters across all the active mirroring instances	BlackDiamond 8000 series BlackDiamond X8 series All Summit series E4G cell site routers	128
Mirroring, one-to-many (monitor port)—maximum number of one-to-many monitor ports.	All platforms	16

Table 5: Supported Limits (continued)

Metric	Product	Limit
MLAG ports—maximum number of MLAG ports allowed.	BlackDiamond 8000 series BlackDiamond X8 series All Summit series (except X430) E4G cell site routers	768
MLAG peers—maximum number of MLAG peers allowed.	All platforms, except Summit X430	2
MPLS RSVP-TE interfaces—maximum number of interfaces.	All platforms, except Summit X450-G2, X440, and X430	32
MPLS RSVP-TE ingress LSPs—maximum number of ingress LSPs.	All platforms, except Summit X450-G2, X440, and X430	2,000
MPLS RSVP-TE egress LSPs—maximum number of egress LSPs.	All platforms, except Summit X450-G2, X440, and X430	2,000
MPLS RSVP-TE transit LSPs—maximum number of transit LSPs.	All platforms, except Summit X450-G2, X440, and X430	2,000
MPLS RSVP-TE paths—maximum number of paths.	All platforms, except Summit X450-G2, X440, X430, and X670-G2 Summit X670-G2	1,000 2,000
MPLS RSVP-TE profiles—maximum number of profiles.	All platforms, except Summit X440, X430, X670-G2, and X450-G2 Summit X670-G2	1,000 2,000
MPLS RSVP-TE EROs—maximum number of EROs per path.	All platforms, except Summit X450-G2, X440, and X430 Summit X450-G2	64 N/A
MPLS RSVP-TE fast reroute—MPLS RSVP-TE fast reroute (FRR) switching time.	E4G-200	50 msec
MPLS LDP peers—maximum number of MPLS LDP peers per switch.	BlackDiamond 8900 xl-series, 8900-40G6x-xm BlackDiamond X8 series E4G-400, E4G-200 Summit X460, Summit X670 Summit X670-G2, X460-G2 Summit X480, Summit X480 (40G VIM), X670V-48t, X770, X670v-48t	64 64 32 32 128 64
MPLS LDP adjacencies—maximum number of MPLS LDP adjacencies per switch.	BlackDiamond 8900 xl-series BlackDiamond 8900-40G6x-xm BlackDiamond X8 series E4G-200, E4G-400 Summit X460, X480, X670, X460-G2 Summit X670V-48t, X480 (40G VIM), X770, X670-G2	50 64 50 50 50 64

Table 5: Supported Limits (continued)

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

Table 5: Supported Limits (continued)

Metric	Product	Limit
MPLS static ingress LSPs—maximum number of static ingress LSPs.	BlackDiamond 8900 xl-series BlackDiamond 8900-40G BlackDiamond X8 Summit X460, X480, X460-G2 Summit x480-40G, X670, x670V-48t, X770, X670-G2 E4G-200 E4G-400	4,000 2,048 2,048 4,000 2,048 2,048 4,000
MPLS static transit LSPs—maximum number of static transit LSPs	BlackDiamond 8900 xl-series BlackDiamond 8900-40G BlackDiamond X8 Summit X460, X480, X670V-48t, X770, X670-G2, X460-G2 Summit X480-40G, X670 E4G-200 E4G-400	4,000 3,000 4,000 4,000 3,000 2,700 4,000
MSDP active peers—maximum number of active MSDP peers.	BlackDiamond 8000 series, 8900 series, X8 series Summit X460, X480, X670, E4G-400, X460-G2 Summit X770, X670-G2 Summit X450-G2 (Advanced Edge License)	64 16 64 N/A
MSDP SA cache entries—maximum number of entries in SA cache.	BlackDiamond 8000 series, 8900 series, X8 series Summit X480, X670, E4G-400 Summit X670-G2, X770 Summit X460-G2 Summit X450-G2 Summit X460	16,000 8,000 14,000 10,000 8,000 6,000
MSDP maximum mesh groups—maximum number of MSDP mesh groups.	BlackDiamond 8000 series, 8900 series, X8 series Summit X460, X480, X670, E4G-400 Summit X770, X670-G2, X460-G2 Summit X450-G2	16 4 16 N/A
Multicast listener discovery (MLD) snooping per-VLAN filters—maximum number of VLANs supported in per-VLAN MLD snooping mode.	BlackDiamond e-series BlackDiamond 8800 c-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 a-series BlackDiamond X8 xl-series Summit X460, X450-G2, E4G-400 Summit X460-G2 Summit X480 Summit X440 Summit X670, E4G-200 Summit X770, X670-G2 Summit X450-G2	250 1,000 500 2,000 500 500 2,000 1,000 1,200 2,000 250 500 1,200 512
Multicast listener discovery (MLD)v1 subscribers—maximum number of MLDv1 subscribers per port.[n]	BlackDiamond 8800 c-series BlackDiamond xl-series, X8 Series Summit X440 Summit X460, X480, X670, E4G-400 Summit X770, X670-G2, X450-G2, X460-G2	500 1,500 750 1,500 4,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
Multicast listener discovery (MLD)v1 subscribers—maximum number of MLDv1 subscribers per switch. [n]	BlackDiamond 8800 series, X8 series Summit X440 Summit X460, X480, X670, E4G-400, X460-G2, X450-G2 Summit X770, X670-G2	10,000 5,000 10,000 30,000
Multicast listener discovery (MLD)v2 subscribers—maximum number of MLDv2 subscribers per port. [n]	BlackDiamond 8800 c-series BlackDiamond xl series BlackDiamond X8 series Summit X440, X450-G2, SummitStack Summit X460, X480, X670, E4G-400, Summit X770, X670-G2, X450-G2, X460-G2	500 2,500 2,000 1,000 2,000 4,000
Multicast listener discovery (MLD)v2 subscribers—maximum number of MLDv2 subscribers per switch. [n]	BlackDiamond 8800 series, BlackDiamond xl series Summit X440, SummitStack Summit X460, X480, X670, E4G-400, X460-G2, X450-G2 Summit X770, X670-G2	10,000 5,000 10,000 30,000
Multicast listener discovery (MLD)v2 maximum source per group—maximum number of source addresses per group.	All platforms, except Summit X430	200
Multicast listener discovery (MLD) SSM-map entries—maximum number of MLD SSM mapping entries.	All platforms	500
Multicast listener discovery (MLD) SSM-MAP entries—maximum number of sources per group in MLD SSM mapping entries.	All platform	50
Network login—maximum number of clients being authenticated on MAC-based VLAN enabled ports.	BlackDiamond 8000 series (clients per module/per system) BlackDiamond X8 series Summit series	1,024
Network login—maximum number of clients being authenticated with policy mode enabled.	Summit X450-G2, X460-G2 Summit X670-G2, X770	1,024 512
Network login—maximum number of dynamic VLANs.	All platforms	2,000
Network login VLAN VSAs—maximum number of VLANs a client can be authenticated on at any given time.	All platforms	10
ONEPolicy Roles/Profiles—maximum number of policy roles/profiles.	Summit X450-G2, X460-G2, X670-G2, X770 All other platforms	63 N/A

Table 5: Supported Limits (continued)

Metric	Product	Limit
ONEPolicy Rules per Role/ Profile—maximum number of rules per role/policy.	Summit X450-G2	IPv6 rules: 256 IPv4 rules: 256 L2 rules: 184 MAC rules:2 56
	Summit X460-G2	IPv6 rules: 512 IPv4 rules: 512 L2 rules: 440 MAC rules: 512
	Summit X770	IPv6 Rules:256 L2 Rules:184 MAC Rules:256 IPv4 Rules:256
	All other platforms	N/A
ONEPolicy Authenticated Users per Switch—maximum number of authenticated users per switch.	Summit X450-G2, X460-G2 Summit X670-G2, X770 All other platforms	Up to 1,024 Up to 512 N/A
ONEPolicy Authenticated Users— maximum authenticated users with a combination of TCI disabled/ enabled profiles.	Summit X450-G2, X460-G2 Summit X670-G2, X770 All other platforms	682–1,022 341–510 N/A
ONEPolicy Authenticated Users per Port—maximum number of authenticated users per port.	Summit X450-G2, X460-G2 Summit X670-G2, X770 All other platforms	Unlimited up to 1,024 Unlimited up to 512 N/A
ONEPolicy Permit/Deny Traffic Classification Rules Types— total maximum number of unique permit/deny traffic classification rules types (system/stack).	Summit X450-G2, X460-G2, X670-G2, X770 All other platforms	Up to 952 N/A
ONEPolicy Permit/Deny Traffic Classification Rules Types— maximum number of unique MAC permit/deny traffic classification rules types (macsource/macdest).	Summit X450-G2, X460-G2, X670-G2, X770 All other platforms	256 N/A
ONEPolicy Permit/Deny Traffic Classification Rules Types— maximum number of unique IPv6 permit/deny traffic classification rules types (ipv6dest).	Summit X450-G2, X460-G2, X670-G2, X770 All other platforms	256 N/A

Table 5: Supported Limits (continued)

Metric	Product	Limit
ONEPolicy Permit/Deny Traffic Classification Rules Types—maximum number of unique IPv4 permit/deny traffic classification rules (typesipsource / ipdest / ipfrag / udpsourceportIP / udpdestportIP / tcpsourceportIP / tcpdestportIP / ipttl / iptos / iptype).	Summit X450-G2, X460-G2, X670-G2, X770 All other platforms	256 N/A
ONEPolicy Permit/Deny Traffic Classification Rules Types—maximum number of unique Layer 2 permit/deny traffic classification rules (ethertype/ port).	Summit X450-G2, X460-G2, X670-G2, X770 All other platforms	184 N/A
OSPFv2/v3 ECMP—maximum number of equal cost multipath OSPFv2 and OSPFv3.	All platforms, except Summit X450-G2, X440, X430, and E4G-200) BlackDiamond 8800 G48Te2 (for IPv6) Summit X450-G2 E4G-200	16 N/A 4 8
OSPFv2 areas—as an ABR, how many OSPF areas are supported within the same switch.	All platforms (except X430, X440)	8
OSPFv2 external routes—recommended maximum number of external routes contained in an OSPF LSDB.	BlackDiamond 8000 series BlackDiamond 8900 xl-series BlackDiamond X8 series BlackDiamond X8 xl-series Summit X460, X670, X770, X670-G2, X460-G2, X450-G2 Summit X480 E4G-200, E4G-400	20,000 130,000 20,000 130,000 5,000 130,000 5,000
OSPFv2 inter- or intra-area routes—recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB with one ABR in OSPF domain.	BlackDiamond 8000 series, 8900 xl-series, X8 series Summit X460, X670, X670-G2, X460-G2 Summit X480, X770 Summit X450-G2 E4G-400	7,000 2,000 7,000 1,000 2,000
OSPFv2 interfaces—recommended maximum number of OSPF interfaces on a switch (active interfaces only).	All platforms (except X430), with Advanced Edge license. All platforms (except X430 and X440) with Core license or higher	4 400
OSPFv2 links—maximum number of links in the router LSA.	All platforms, except Summit X450-G2, X770, and X430 Summit X450-G2 Summit X770	400 4 419

Table 5: Supported Limits (continued)

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

Table 5: Supported Limits (continued)

Metric	Product	Limit
PIM IPv4 (maximum interfaces)—maximum number of PIM active interfaces.	All platforms, except Summit X430, X450-G2, and X440 Summit X440 Summit X450-G2	512 253 4
PIM IPv4 (maximum interfaces)—maximum number of PIM-snooping enabled interfaces.	All platforms, except Summit X430	512
PIM IPv4 Limits—maximum number of multicast groups per rendezvous point.	All platforms, except Summit X430	180
PIM IPv4 Limits—maximum number of multicast sources per group.	BlackDiamond 8800 (E-series modules) BlackDiamond 8800 (C-series modules) BlackDiamond 8800 (xl-series modules) BlackDiamond X8 Summit X460-G2, X670-G2, X770, X450-G2 Summit X460, X480 Summit X670-48x Summit X670-48t Summit X440	1,000 3,000 4,000 3,000 5,000 1,200 1,000 4,000 175
PIM IPv4 Limits—maximum number of dynamic rendezvous points per multicast group.	All platforms, except Summit X430	145
PIM IPv4 Limits—static rendezvous points.	All platforms, except Summit X430	32
PIM IPv6 (maximum interfaces)—maximum number of PIM active interfaces.	All platforms, except Summit X450-G2 and X430 Summit X450-G2	512 4
PIM IPv6 Limits—maximum number of multicast group per rendezvous point.	All platforms, except Summit X430	70
PIM IPv6 Limits—maximum number of multicast sources per group.	BlackDiamond 8000 BlackDiamond X8 Summit X460-G2, X670-G2 Summit X460, X480 Summit X670 Summit X440 Summit X450-G2 Summit X770	1,280 1,500 2,500 800 2,000 175 2,000 2,500
PIM IPv6 Limits—maximum number of dynamic rendezvous points per multicast group.	All platforms, except Summit X430	64
PIM IPv6 Limits—maximum number of secondary address per interface.	All platforms, except Summit X430	70
PIM IPv6 Limits—static rendezvous points.	All platforms, except the Summit X430	32

Table 5: Supported Limits (continued)

Metric	Product	Limit
Policy-based routing (PBR) redundancy—maximum number of flow-redirects.	All platforms	256 [o]
Policy-based routing (PBR) redundancy—maximum number of next hops per each flow-direct.	All platforms	32 [o]
Port-specific VLAN tags—maximum number of port-specific VLAN tags.	All platforms, except Summit X450-G2, X440, and X430	1,023
Port-specific VLAN tags—maximum number of port-specific VLAN tag ports.	BlackDiamond X8 and 8800 xl-series Summit X480 Summit X460-48t Summit X460-24x, X670-48x Summit X670V-48t Summit X670V-48t stack Summit X770, X670-G2 Summit X460-G2 E4G-400 E4G-200	8,090 3,800 7,200 3,400 3,600 7,200 6,400 4,000 3,400 3,800
Private VLANs—maximum number of subscribers. Assumes a minimum of one port per network and subscriber VLAN.	BlackDiamond 8800 c-, e-, xl-series with eight modules of 48 ports 8900-G96T-c modules	383
	BlackDiamond X8 series Summit X770 Summit X670-G2, X670v-48t Summit X670 Summit X480 Summit X460-G2, X460 Summit X440 Summit X430 Summit X450-G2 E4G-200 E4G-400	767 103 63 47 23 53 25 27 51 11 33
Private VLANs—maximum number of private VLANs with an IP address on the network VLAN. **Note:** This limit is dependent on the maximum number of private VLANs in an L2-only environment if the configuration has tagged and translated ports.	Summit X770, X670-G2, X460-G2, X450-G2 Summit X670, X480, X460, X460, X480 Summit X440 E4G-200, E4G-400	1,024 512 127 512

Table 5: Supported Limits (continued)

Metric	Product	Limit
Private VLANs—maximum number of private VLANs in an L2-only environment.	BlackDiamond 8800 c-, e-series BlackDiamond 8900 series, X8 series E4G-200 E4G-400 Summit X440 Summit X480, Summit X670 Summit X460 Summit X770, X670-G2, X460-G2, X450-G2 Summit X430	384 2,046 597 1,280 127 597 820 1,280 255
PTP/1588v2 Clock Ports	Summit X770, X460-G2, X670-G2, and E4G-200, E4G-400 cell site routers	32 for boundary clock 1 for ordinary clock
PTP/1588v2 Clock Instances	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	2 combinations: • Transparent clock + ordinary clock • Transparent clock + boundary clock
PTP/1588v2 Unicast Static Slaves	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	40 entries per clock port
PTP/1588v2 Unicast Static Masters	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	10 entries per clock type
Route policies—suggested maximum number of lines in a route policy file.	All platforms	10,000
RIP Learned Routes—maximum number of RIP routes supported without aggregation.	All platforms, except Summit X430	10,000
RIP neighbors—maximum number of RIP neighbors.	E4G-200	256
RIP interfaces on a single router—recommended maximum number of RIP routed interfaces on a switch.	BlackDiamond 8000 series, X8 series BlackDiamond 8900 xl-series Summit X440 Summit X460, X670-G2, X460-G2 Summit X480 Summit X670, X770, X450-G2 E4G-400	256 384 128 256 384 256 256
RIPng learned routes—maximum number of RIPng routes.	BlackDiamond 8000 series, X8 series BlackDiamond 8900 xl-series Summit X480 Summit X460, X670, X670-G2, X460-G2, X770, X450-G2 E4G-200	3,000 5,000 5,000 3,000 3,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
Spanning Tree (maximum STPDs)—maximum number of Spanning Tree Domains on port mode EMISTP.	All platforms (except Summit X430, X440) Summit X440 Summit X430	64 32 16
Spanning Tree PVST+—maximum number of port mode PVST domains. Note: • Maximum of 10 active ports per PVST domain when 256 PVST domains are configured. • Maximum of 7 active ports per PVST domain when 128 PVST domains are configured.	BlackDiamond X8 and 8900 series Summit X670, X770, X670-G2 Summit X460, X480, X440, X460-G2 Summit X430 Summit X450-G2 E4G-400	256 256 128 50 128 128
Spanning Tree—maximum number of multiple spanning tree instances (MSTI) domains.	All platforms (except Summit X430, X440) Summit X430 Summit X440	64 5 32
Spanning Tree—maximum number of VLANs per MSTI. Note: Maximum number of 10 active ports per VLAN when all 500 VLANs are in one MSTI.	BlackDiamond X8, 8800, 8900 MSM 128/XL Summit X770, X670-G2, X670v-48t, X670 Summit X480, X460-G2, X460, X450-G2 E4G-200 E4G-400 Summit X440 Summit X430	500 500 600 500 600 250 100
Spanning Tree—maximum number of VLANs on all MSTP instances.	BlackDiamond X8, 8800, 8900 MSM 128/XL Summit X770 Summit X670-G2, X670v-48t, X670, X480 Summit X460-G2, X460, X450-G2 E4G-200 E4G-400 Summit X440 Summit X430	1,000 1,024 1,000 1,024 1,000 1,024 500 200
Spanning Tree (802.1d domains)—maximum number of 802.1d domains per port.	All platforms	1
Spanning Tree (number of ports)—maximum number of ports including all Spanning Tree domains.	All platforms (except Summit X430, X440) Summit X430 Summit X440	4,096 1,024 2,048

Table 5: Supported Limits (continued)

Metric	Product	Limit
Spanning Tree (maximum VLANs)—maximum number of STP-protected VLANs (dot1d and dot1w).	BlackDiamond X8, 8800, 8900 MSM 128/XL Summit X770 Summit X670-G2, X670v-48t, X670, X480 Summit X460-G2, X460, X450-G2 E4G-200 E4G-400 Summit X440 Summit X430	1,024 1,024 560 600 500 600 500 128
SSH (number of sessions)—maximum number of simultaneous SSH sessions.	All platforms	8
Static MAC multicast FDB entries—maximum number of permanent multicast MAC entries configured into the FDB.	BlackDiamond 8000 c-, e-, xl-series BlackDiamond X8 series All Summits E4G-200, E4G-400	1,024
Syslog servers—maximum number of simultaneous syslog servers that are supported.	All platforms	4
Telnet (number of sessions)—maximum number of simultaneous Telnet sessions.	All platforms	8
TRILL—trees rooted from switch.	BlackDiamond X8 Summit X670, X770	1
TRILL—computed trees.	BlackDiamond X8 Summit X670, X770	1
TRILL—TRILL VLANs.	BlackDiamond X8 Summit X670, X770	4
TRILL—forwarding VLANs.	BlackDiamond X8 Summit X670, X770	4,095
TRILL—forwarding ports.	BlackDiamond X8 Summit X670, X770	All
TRILL—RBridge FDB entries.	BlackDiamond X8 Summit X670 Summit X770	128,000 128,000 288,000
TRILL—ECMP RBridge next hops.	BlackDiamond X8 Summit X670, X770	8
TRILL—neighbor adjacencies.	BlackDiamond X8 Summit X670, X770	32
TRILL—nodes.	BlackDiamond X8 Summit X670, X770	256
TRILL—links.	BlackDiamond X8 Summit X670, X770	2,000

Table 5: Supported Limits (continued)

Metric	Product	Limit
Virtual routers—maximum number of user-created virtual routers that can be created on a switch. Note: Virtual routers are not supported on Summit X440 series switches.	BlackDiamond 8000 c-, xl-, xm-series BlackDiamond X8 series Summit X460, X460-G2, X480, X670, X670-G2, X770, X450-G2 E4G-200, E4G-400	63
Virtual router forwarding (VRFs)—maximum number of VRFs that can be created on a switch. Note: * Subject to other system limitations.	All platforms, except Summit X440 and X430	960 *
Virtual router protocols per VR—maximum number of routing protocols per VR.	All platforms, except Summit X440, X430	8
Virtual router protocols per switch—maximum number of VR protocols per switch.	All platforms, except Summit X440, X430	64
VLAN aggregation—maximum number of port-VLAN combinations on any one superVLAN and all of its subVLANs.	All platforms (except Summit X430, X440) Summit X440, X430	1,000 256
VLANs—includes all VLANs. Note: ExtremeXOS supports only 4,092 user-configurable VLANs. (VLAN 1 is the default VLAN, and 4,095 is the management VLAN, and you may not configure them.)	All platforms	4,094
VLANs—maximum number of port-specific tag VLANs.	BlackDiamond 8800 xl-series only, BlackDiamond X8 series BlackDiamond X8 xl-series Summit X460, X770, X480, E4G-400, X670-G2, X460-G2 Summit X670, X670V-48t E4G-400 E4G-200	1,023 4,093 4,093 1,023 4,093 2,047
VLANs—maximum number of port-specific tag VLAN ports.	BlackDiamond 8800 xl-series only BlackDiamond X8 BlackDiamond X8 xl-series E4G-400, E4G-200 Summit X460, X670, X670V-48t, X460-G2 Summit X770, X670-G2 Summit X480	4096 4096 32,767 4096 4096 8,192 16,383

Table 5: Supported Limits (continued)

Metric	Product	Limit
VLANs (Layer 2)—maximum number of Layer 2 VLANs.	All platforms	4,094
VLANs (Layer 3)—maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.	BlackDiamond X8 Summit X460-G2, X670, X770, X670-G2, X450-G2 Summit X440 Summit X480, X460 E4G-200, E4G-400	2,048 2,048 254 512 512
VLANs (maximum active port-based)—maximum active ports per VLAN when 4,094 VLANs are configured with default license.	BlackDiamond X8, 8800 series Summit X770, X670-G2, X670v-48t, X670, X480, X460-G2, X460, X450-G2 E4G-200 E4G-400 Summit X440 Summit X430	32 32 12 32 7 2
VLANs (maximum active protocol-sensitive filters)—number of simultaneously active protocol filters in the switch.	All platforms, except Summit X450-G2 Summit X450-G2	15 16
VLAN translation—maximum number of translation VLANs. Assumes a minimum of one port per translation and member VLAN.	BlackDiamond 8000 a-, c-, e-, xl series with eight modules of 48 ports 8900-G96T-c modules	 383 767
	Summit X770 Summit X670-G2, X670v-48t Summit X670 Summit X480 Summit X460-G2 Summit X460 E4G-200 E4G-400 Summit X440 Summit X430 Summit X450-G2	103 63 47 53 53 57 11 33 25 27 51
VLAN translation—maximum number of translation VLAN pairs with an IP address on the translation VLAN. Note: This limit is dependent on the maximum number of translation VLAN pairs in an L2-only environment if the configuration has tagged and translated ports.	Summit X770, X670-G2, X450-G2 Summit X670, X480, X460 Summit X440 E4G-200, E4G-400	1,024 512 127 512

Table 5: Supported Limits (continued)

Metric	Product	Limit
VLAN translation—maximum number of translation VLAN pairs in an L2-only environment.	BlackDiamond 8800 c-, e-series BlackDiamond 8900 xl-series BlackDiamond X8 series Summit X460 Summit X430 Summit X480, X670, X670-G2, X460-G2 Summit X450-G2, X770 Summit X440 E4G-400, E4G-200	384 2,046 2,046 2,000 512 2,046 1,024 127 2,000
VRRP (v2/v3-IPv4) (maximum instances)—maximum number of VRRP instances for a single switch, with Advanced Edge license or higher.	BlackDiamond X8 BlackDiamond 8800 MSM-48c and MSM-96 BlackDiamond 8900 xl-series 8900-MSM128 Summit X770, X670, X670-G2, X460-G2, X480, X450-G2 Summit X460 Summit X440 E4G-200, E4G-400	511 511 511 511 255 32 128
VRRP (v3-IPv6) (maximum instances)—maximum number of VRRP instances for a single switch, with Advanced Edge license or higher. (VRRP-VRRPv3-IPv6)	BlackDiamond X8, BlackDiamond 8800 MSM-48c and MSM-96 BlackDiamond 8900 xl-series 8900-MSM128 Summit X770, X670, X670-G2, X460-G2, X450-G2 Summit X460, X480 Summit X440 E4G-200, E4G-400	511 511 511 511 255 15 255
VRRP (v2/v3-IPv4/IPv6) (maximum VRID)—maximum number of unique VRID numbers per switch.	All platforms with Advanced Edge license or higher, except Summit X430	31
VRRP (v2/v3-IPv4/IPv6) (maximum VRIDs per VLAN)—maximum number of VRIDs per VLAN.	All platforms with Advanced license or higher, except for Summit X430	31
VRRP (v2/v3-IPv4/IPv6) (maximum ping tracks)—maximum number of ping tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8
VRRP (maximum ping tracks)—maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher.	All platforms, except the Summit X440	8 (20 centisecond or 1 second hello interval)
	Summit X440 Hello interval: 20 centiseconds Hello interval: 1 second	1

Table 5: Supported Limits (continued)

Metric	Product	Limit
VRRP (v3-IPv6) (maximum ping tracks)—maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher.	All platforms, except the Summit X440	8 (20 centisecond or 1 second hello interval)
	Summit X440 Hello interval: 20 centiseconds Hello interval: 1 second	1
VRRP (v2/v3-IPv4/IPv6) (maximum iproute tracks)—maximum number of IP route tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8
VRRP (v2/v3-IPv4/IPv6)—maximum number of VLAN tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8
XML requests—maximum number of XML requests per second. Note: Limits are dependent on load and type of XML request. These values are dynamic ACL data requests.	BlackDiamond 8800 c-series with 100 DACLs with 500 DACLs	 10 3
	BlackDiamond 8800 c-series with 100 DACLs with 500 DACLs	 10 3
	Summit X480, X670 with 100 DACLs with 500 DACLs	 4 1
	Summit X450-G2 with 100 DACLs	10
XNV authentication—maximum number of VMs that can be processed (combination of local and network VMs).	All platforms, except Summit X430 and X450-G2 Summit X450-G2	2,048 1,024
XNV database entries—maximum number of VM database entries (combination of local and network VMs).	All platforms, except Summit X430	16,000
XNV database entries—maximum number of VPP database entries (combination of local and network VPPs).	All platforms, except Summit X430	2,048
XNV dynamic VLAN—Maximum number of dynamic VLANs created (from VPPs / local VMs).	All platforms, except Summit X430	2,048

Table 5: Supported Limits (continued)

Metric	Product	Limit
XNV local VPPs—maximum number of XNV local VPPs.	All platforms, except Summit X430	2,048 ingress 512 egress
XNV policies/dynamic ACLs—maximum number of policies/ dynamic ACLs that can be configured per VPP.	All platforms, except Summit X430	8 ingress 4 egress
XNV network VPPs—maximum number of XNV network VPPs.[p]	All platforms, except Summit X430	2,048 ingress 512 egress

[a] The table shows the total available.
[b] Limit depends on setting configured for "configure forwarding external-tables".
[c] When there are BFD sessions with minimal timer, sessions with default timer should not be used.
[d] Based on in “none more-l2” mode.
[e] Based on forwarding internal table configuration "more l2”.
[f] Effective capacity varies based on actual MAC addresses and VLAN IDs used and hash algorithm selected.
[g] Based on “l2-only mode”.
[h] Based on forwarding internal table configuration "more l3-and-ipmc”.
[i] Based on forwarding external table configuration "l3-only ipv4".
[j] The limit depends on setting configured with configure iproute reserved-entries.
[k] Based on forwarding external table configuration "l3-only ipv4”.
[l] Based on forwarding external table configuration "l3-only ipv6”.
[m] The IPv4 and IPv6 multicast entries share the same hardware tables, so the effective number of IPv6 multicast entries depends on the number of IPv4 multicast entries present and vice-versa.
[n] If IGMP and MLD are simultaneously configured on the switch, the number of effective subscribers supported would be appropriately lessened.
[o] Sum total of all PBR next hops on all flow redirects should not exceed 4,096.
[p] The number of XNV authentications supported based on system ACL limitations.

3 Open Issues, Known Behaviors, and Resolved Issues

This chapter lists open software issues, limitations in ExtremeXOS system architecture (known issues), and resolved issues in ExtremeXOS.

Open Issues

The following are new open issues for supported features found in ExtremeXOS 16.2.

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

CR Number	Description
General	
xos0048715	IPv6 ECMP works for hardware-forwarded traffic, but does not work for slow-path traffic. Note: Either use BFD or ping protection to keep all router neighbors alive, or configure static neighbors and static FDB entries for all router neighbors. BFD or ping protection are the preferred methods.
BlackDiamond 8000 Series Switches	
xos0063621	After 5 million internet routes are received from an upstream peer, BlackDiamond 8000 series switches experience a link flap, and then CPU utilization goes up to 78%. Note: Flush and relearn routes.
SummitStacks	
xos0063739	On E4G-400 and Summit X440 stacks, with 256 Down MEPs, issuing the command **`restart process dot1ag`** generates the following error: `<Erro:cm.sys.actionErr> Slot-2: Error while loading "maintenancePoint": MP 12 Creation Failed due to HAL problem.`

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

CR Number	Description

xos0064105	On first-time boot up without any configuration on Summit X670v/X480 stacks, backup and standby nodes do not come to operational state and the following error message appears:
	`04/13/2016 06:04:02.95 <Erro:HAL.Port.Error> Unable to get media type from slot 2 port 41 error -1`
	`04/13/2016 06:04:02.95 <Erro:HAL.VLAN.Error> failed to setup qosprofile bandwidth port 2:41 unit 9 port 41 qos 0 minbw 0 maxbw 10000000 (Conduit failure)`
	`04/13/2016 06:04:02.95 <Erro:HAL.VLAN.Error> failed to setup qosprofile bandwidth port 2:41 unit 9 port 41 qos 1 minbw 0 maxbw 10000000 (Conduit failure)`
	`04/13/2016 06:04:02.95 <Erro:HAL.VLAN.Error> failed to setup qosprofile bandwidth port 2:41 unit 9 port 41 qos 2 minbw 0 maxbw 10000000 (Conduit failure)`
	Also, all slots are continuously rebooting, except master node. Eventually, master and backup come to "in sync" state automatically, and then standby nodes go to "present" state instead of "operational" state.

xos0061909	Creating an IPFIX mirroring instance to a monitor port, deleting the mirroring instance, and then recreating it again to a different monitor port, causes an error message, similar to the one below, to appear, and IPFIX mirroring does not work: `Erro:HAL.Mirror.Error&gt Slot-1: Failed to create mirroring destination for slot 2, unit 9 Entry exists` Note: If the error appears in the log, disable and delete the mirror instance, and then add it back again.
Summit X450-G2	
xos0064736	During severe congestion on Summit X450-G2 stacks, master slot reboots due to EPM watchdog expiration from memory depletion and stuck kernel.

Known Behaviors

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

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

CR Number	Description
BlackDiamond 8000 Series Switches	

xos0062138	Enabling, and then disabling VRRP fabric routing generates a HAL port error log messages similar to the following:
	`7/18/2015 11:07:17.75 <Erro:HAL.Port.Error> MSM-A: Failed to configure static mac move behavior on port 1:2.`
	`07/18/2015 11:07:17.76 <Erro:HAL.Port.Error> MSM-A: Failed to configure static mac move behavior on port 1:3.`
	`07/18/2015 11:07:17.76 <Erro:HAL.Port.Error> MSM-A: Failed to configure static mac move behavior on port 1:4. 07/18/2015 11:07:17.76 <Erro:HAL.Port.Error> MSM-A: Failed to configure static mac move behavior on port 1:5.`

Resolved Issues in ExtremeXOS 16.2

The following issues were resolved in ExtremeXOS 16.2. ExtremeXOS 16.2 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, ExtremeXOS 16.1, ExtremeXOS 16.1.2, and ExtremeXOS 16.1.3. For information about those fixes, see the release notes for the specific release.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2

CR Number	Description
Summit X430 Series Switches	
xos0059486	On Summit X430 series switches, optics are not detected after repeated removal and reinsertion of optics when CPU is busy.
xos0064084	In Summit X430 series switches, the command `show power details` displays fan status as "empty".
Summit Series Switches	
xos0059007	QSFP+ to SFP+ adapter support is added to work with all optical SFP+ transceivers with the exception of LRM and passive copper direct attach cables.
xos0059484	Gigabit Ethernet compliance value is shown as "UNKNOWN" for BXU/D optics. Also, DDMI values do not appear. Media Type appears as "NONE" in `show ports configuration` command.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0059508	Link does not come up when 1G BX U/D optics are inserted into 10G SFP+ ports. Note: Configure the port for speed 1,000 with auto-negotiation "on" using the command `configure port <> auto on speed 1000 duplex full`. For Summit X670V switches, remove and then re-insert the optics. For all Summit switches, except the X670V, the corrected behavior persists even after rebooting. For Summit X670V switches after rebooting, the transceiver must be removed, and then re-inserted.
xos0062570	In SummitStacks, executing the command enable sflow ports all enables sFlow inappropriately on stacking ports.
xos0062821	ACL rules installed are not mapped to single virtual group even though ACL action-resolution mode is highest-priority.
Summit X440 Series Switches	
xos0060466	RX CRC errors with traffic loss occur when "CBL, SFPP, PASSIVE 10GB-C10-SFPP, 10M" from Tyco is inserted into Summit X440-48t-10G switches SFP+.
xos0062362	On Summit X440-24t series switches, process rtmgr pid 1572 ends unexpectedly with signal 11 error after disabling BGP with compressed routes.
xos0062621	On Summit 440-8p switches, the `show fan` command output displays that the fan is unsupported.
xos0063627	ARP is not re-added to hardware after it is removed initially due to the table being full.
xos0064050	While running diagnostics on a Summit X440 10G model switch with revision 10 and diagnostics test version 6.0 or above, "Test loopback phy fiber" and "Test snake interface" fail.
Summit X440-G2 Series Switches	
xos0063317	When policy is enabled one less than the max number of netlogin users can be handled per slot.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description

xos0062775	With 32 Up MEPs configured, saving and rebooting produces the following errors:
	`10/03/2015 18:14:06.81 <Noti:EPM.system_stable> System is stable. Change to warm reset mode`
	`10/03/2015 18:14:10.06 <Info:HAL.IPv4ACL.Info> Synching ACLs to Switch`
	`10/03/2015 18:14:10.35 <Info:HAL.IPv4ACL.Info> Done synching ACLs to Switch`
	`10/03/2015 18:15:11.73 <Erro:HAL.FDB.Error> pibInstallCPUFilter(1:33, 12, 33) - pibAclWrapInstall(1, 2, 1020, 5242950) returned "Table full".`
	`10/03/2015 18:15:11.77 <Erro:HAL.FDB.Error> pibInstallCPUFilter(1:33, 13, 33) - pibAclWrapInstall(1, 2, 1021, 5242950) returned "Table full".`
	`10/03/2015 18:15:13.26 <Erro:HAL.FDB.Error> pibInstallCPUFilter(1:5, 12, 33) - pibAclWrapInstall(1, 1, 1028, 5242950) returned "Table full".`
	`10/03/2015 18:15:13.30 <Erro:HAL.FDB.Error> pibInstallCPUFilter(1:5, 13, 33) - pibAclWrapInstall(1, 1, 1029, 5242950) returned "Table full".`
	`10/03/2015 18:15:15.47 <Info:HAL.Card.Info> Switch is operational`

Summit X450-G2 Series Switches	
xos0060129	On Summit X450-G2 series switches, 10/100/1000BASE-T SFP+ optics do not link to similar optics when in the SFP/SFP+ ports. They do link or partially link when connected to a regular triple speed copper port.
Summit X460 Series Switches	
xos0063206	Cannot add L2 entries in hardware due to a full L2 table caused by hash collisions.
xos0063595	On Summit X460 series switches, the command to configure stacking ports does not show the native option.
xos0063948	Clearflow delta values are randomly not calculated properly.
Summit X460-G2 Series Switches	

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0060018	With a 0.5M, 40G QSFP MOLEX passive copper cable inserted, disabling the port where the optic is inserted, rebooting, and then enabling the port, the port stays in the ready state and doesn't come up as enabled. Note: Remove and re-insert the optics. The port then comes up as enabled.
xos0061486	Combo ports have unsupported autonegotiation and half-duplex settings.
xos0062855	On the Summit X460-G2 series switches, VPLS packets are forwarded with two tags when the service VLAN ports are also members of an untagged VMAN.
xos0062913	On Summit X460-G2 series switches, copper combo port does not advertise its flow control capabilities to peers.
xos0063495	Policy authentication fails when RADIUS request queue has stale entries.
xos0063811	Summit X460-G2 series switches with ExtremeXOS 15.6 through 21.1, have the following limitations for SyncE input reference frequency: • Network clock does not lock with input SyncE source port 52 (both at 10G and 1G speed) on all 48-port models (X460-G2-48t, 48x and 48p). Note: For SyncE input at 10G, avoid port 52. • When the 10G ports operate at 1G speed, the network clock does not lock. Models with Ethernet BASE-T or Power over Ethernet (PoE) ports may lock on initial configuration, but do not lock after a save and reboot. Note: For SyncE input at 1G, use a 1G port, not a 10G port.
Summit X480 Series Switches	
xos0054290	When a semicolon is missing from ACL matching criteria, the ACL process ends unexpectedly.
xos0060614	When exporting OSPF routes into BGP, and long policy name causes process dcbgp pid 1575 to end unexpectedly with signal 11.
xos0061251	For Summit X480, X460, X440, X430 series switches, and E4G-200 and E4G-400 cell site routers, Dot1p examination functionality does not work correctly for untagged ports when diffserv examination is enabled on those ports.
Summit X670 Series Switches	

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0060656	Link does not come up when 1G BX U/D optics are inserted into 10G SFP+ ports. Note: Configure the port for speed 1,000 with auto-negotiation "on" using the command configure port <> auto on speed 1000 duplex full. For Summit X670V switches, remove and then re-insert the optics. For all Summit switches, except the X670V, the corrected behavior persists even after rebooting. For Summit X670V switches after rebooting, the transceiver must be removed, and then re-inserted.
xos0061167	Links become active without a connection with tri-speed Base-T SFP installed.
xos0061559	Enabling OpenFlow on VLANs causes double-wide ACL slice to be used even though it can fit in single-wide slice.
xos0063052	Traffic loss occurs on computer connected to Summit X670v-48t switches when the connected switch port is oversubscribed in 100 MB mode.
xos0063137	Known unicast traffic is not shared between the stacking high-gig trunk ports.
Summit X670-G2 Series Switches	
xos0061791	On SummitStacks containing master and standby nodes of different switches, the standby node may go to failed state after a node reboot.
xos0062166	On Summit X670-G2 series switches configured with L3VPN, executing the `clear iparp` command causes the switch to reboot with Kernel Oops.
xos0063204	Traffic stops on LAG ports when frequently modifying the sharing group.
xos0063807	On Summit X670-G2 series switches, egress ACL rule actions do not take effect on ports 64-72.
Summit X770 Series Switches	

xos0053867	Internal errors occur when looking at egress port qosmonitor after issuing any QoS command such as: `enable diffser examination port all`, `enable diffser replacement port all`, `disable diffser examine port all`, or `disable diffser replacement port all`:
	
	`Slot-1 Stack.34 # sh port 1:93 2:73 qosmonitor`
	`Error: Internal error --> no response to stats request for port 1:93.`

SummitStack	
xos0061799	Precedence order between policy port rules and policy MAC-based rules is not preserved following a master/backup Failover.
xos0062217	In SummitStacks with eight nodes and sFlow configuration, "Hardware L3 Table full" error messages appear when the stacks have a large number of Layer 3 entries.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0062700	When upgrading from ExtremeXOS 15.7 or earlier to 16.1, image download fails if image was installed in backup node first and master node second.
xos0062949	HAL process ends unexpectedly in stack after executing the following commands: `debug hal configure stacking pdu-trace mask 0xf`, `debug hal configure stacking pdu-trace capture cap_file`
xos0063344	With MLAG and LAG configurations, when a stack node comes up after a reboot, FDB entries flooded from other slots are programmed on incorrect ports internally.
xos0061777	Standby nodes do not come back up to operational state after they go into failed state.
xos0062367	ACL process ends unexpectedly on repeated refresh of ACL policy with clear-flow action.
xos0062522	In SummitStack switches, standby slots go to failed state when a very large number of log messages are continuously generated in the switch.
xos0062800	Stack node fails because of license mismatch for 3rd-party optics.
xos0063242	Stacks configured as DHCP clients do not respond to pinging after failover.
xos0063490	CFM stays down after slot reboot on a stack.
xos0063349	Switch stops responding to SNMP requests if SNMP get for multiple OIDs is continuously initiated.
BlackDiamond 8800 Series Switches	
xos0063333	In BlackDiamond 8800 series switches, optics information is not detected and ports remain in "Ready" state after reboot.
xos0063510	ACL rule to deny packets matching L4 match condition stops working if a rule with VID as match condition is appended without an L4 match condition.
xos0064579	Support for dual hash needs to be enabled for BlackDiamond 8800 c-series modules.
xos0057419	On BlackDiamond 8800 series switches, C-series modules reboot after enabling dot1p replacement on ports.
xos0061260	When learning a large number of BGP routes, routes are learned first in route table, rather than LOCAL-RIB. No routes are learned to RIB until all routes are learned in route table. All prefixes are learned in route table, and only then are all prefixes are learned to LOCAL RIB.
xos0062009	In BlackDiamond 8800 series switches with XL modules, clearing FDBs when there is a loop causes the FDBs to lose synchronization across slots or switching units.
xos0062535	On BlackDiamond 8900-10G24Xc I/O modules, packets are dropped by ACL rules with "redirect-port-list" action.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0063872	After multiple executions of `run failover` with redirect-flow configuration, IPv4 ping fails.
BlackDiamond X8 Series Switches	
xos0060666	After failover, traffic gets flooded on the ports of service VLAN in H-VPLS core.
xos0061639	Packets ingressing on VLAN-bridged interfaces (Layer 2 VLANs) are not forwarded when the destination MAC address is the same as the switch MAC address, and the switch has at least one Layer 3 interface.
xos0061902	BlackDiamond X8 series switches use VLAN instance as index instead of router interface (rtif) for ARP entries.
xos0062200	BGP is not converging when there is change in network in scaling environment. When the best path goes down, BGP RIB converges. but route table is still showing old peer as gateway.
xos0062262	UDP profile configurations on VLANs do not take effect in BlackDiamond X8 series switches.
xos0062306	Packets get reflected with same tag on port-specific, tag-enabled VLANs after failover. Issue happens only on switches having both port-specific tags and MPLS RSVP-TE configurations.
xos0062477	BlackDiamond X8 series switchesâ€™ management ports flap and show "Detected Tx unit hung" error messages.
xos0062499	Multicast packets are dropped in Layer 2 bridged VLANs.
xos0063546	The following error message may appear when BlackDiamond X8 modules are rebooted after configuring port partition: `<Erro:Kern.Card.Error> Slot-2: _setSchedMode: u:p=1:006 schedMode=1 err=Invalid parameter`
xos0063928	In BlackDiamond X8 series switches, Sysuptime in sFlow packets is invalid.
xos0063958	BDX8-40G12X-XL module goes into reboot cycle if any physical loop is configured or a network loop is present on the module.
xos0064010	The command `show port buffer` displays an incorrect port range for 100G I/O modules.
General	
xos0063554	The following vulnerability in OpenSSL exists that impacts ExtremeXOS (CVE-2015-3197): A malicious client can negotiate SSLv2 ciphers that have been disabled on the server and complete SSLv2 handshakes even if all SSLv2 ciphers have been disabled, provided that the SSLv2 protocol was not also disabled via SSL_OP_NO_SSLv2. This issue affects OpenSSL versions 1.0.2 and 1.0.1.
xos0061745	Ampersand used in UPM script is replaced by "& amp" in the XSF configuration.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0061841	FDB entries are not learned again after limit learning is unconfigured, and then configured again, with PSTAG configuration in SummitStacks.
xos0061943	MPLS process ends unexpectedly when get-next is done with incomplete OID for mplsXCIndex.
xos0062850	When upgrading ExtremeXOS to 15.7 or later releases, the web HTTP access is enabled even though it is disabled in the configuration.
xos0063028	RADIUS configuration with shared-secret of 32 character is lost after reboot.
xos0063248	NTP MD5 authentication with NTP server is failing.
xos0063257	Saving configuration fails/times-out when VLANs added to a mirror filters are renamed.
xos0063271	Layer 3 packets in non-default virtual routers are slow-path forwarded after disabling MPLS in the peer switch.
xos0062255	CEP CVID configurations is missing after adding/deleting the port from sharing.
xos0062720	Unable to save configuration when ACL/CFM is configured on multiple VLANs.
xos0057931	After rebooting the switch multiple times, the following error log message appears: `<Erro:cm.loadErr> Failed to load configuration: timed out (after 150 seconds) while waiting for all applications to get ready to load configuration on OPERATIONAL ( eaps is still not ready yet).`
xos0061788	The process devmgr ends unexpectedly during snmpwalk when continuous EMS logs are sent to the switch console.
xos0062271	CLI memory leak occurs when executing show commands with include option through script.
xos0064043	Unable to use a configuration file that has been copied from an existing configuration file.
xos0064179	MAC movement occurs in switch acting as an STP root bridge when PVST+ BPDUs are sent by peer switch using STP blocked port.
xos0062366	After rebooting, DHCP binding entries are not restored using vr-default.
xos0057231	An FDB entry created by ARP with "i" flag set is not removed from the FDB table after a static entry for the same IP address is added with a different MAC value.
xos0063521	A few IBGP routes are not updated in routing table when `disable bgp` and `enable bgp` commands are executed in quick succession.
xos0064203	Incorrect next hop is choosen by BGP route after port flap.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0062441	The process rtMgr ends unexpectedly when IPv6 static route is deleted.

xos0060319	Trying to create the maximum number of Layer 3 VLANs causes the following error messages to appear:
	`02/20/2015 11:37:34.18 <Erro:HAL.VLAN.Error> MSM-A: pibVlanInstallBatchFilter(): Internal error while translating IPv6MC Ctrl Port filter id on Slot  2, unit 13.`
	`02/20/2015 11:37:34.18 <Erro:HAL.VLAN.Error> MSM-A: Failed to install IPV6 Link-Local MC Control Packet Filter for the port 2:24 (Conduit failure)`
	`02/20/2015 11:37:34.18 <Erro:HAL.VLAN.Error> MSM-A: pibVlanInstallBatchFilter(): Internal error while translating IPv6MC Ctrl Port filter id on Slot  2, unit 13.`
	`02/20/2015 11:37:34.18 <Erro:HAL.VLAN.Error> MSM-A: Failed to install IPV6 Link-Local MC Control Packet Filter for the port 2:48 (Conduit failure)`

xos0062265	Some legacy commands are not recognized.
xos0062277	The command show `vlan` *`vlan_list`* does not show information for dynamic VLANs nor the Default VLAN. Error appears.
xos0062427	EDP process ends unexpectedly when CDP packets without portId TLV are received.
xos0058750	Neighbor discovery packets are duplicated in L2 VLANs when IPv6 addresses are configured for other VLANs that do not have any ports.
xos0062240	Port that was administratively disabled becomes up after enabling rx pause.
xos0061506	In Summit X440-G2 and X460-G2 series switches, the combo port comes up as active even though when link peer port is down.
xos0063255	In Chalet, VLANs are sorted incorrectly..
xos0064447	Creation of user accounts through XML does not work.
AAA	
xos0064307	RADIUS accounting configuration is incorrect as shown by the command `show conf aaa` and is lost after upgrade.
ACL	
xos0062537	HAL crash occurs when redirect-port-list action contains more than 64 ports.
xos0062619	SSH access-profile using policy does not work with IPv6 addresses.
xos0063082	Updated DSCP value is not refreshed for Dynamic ACLs.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0063172	ACL action "redirect-port-list" does not take effect when another slice has a rule to match all packets with deny action.
xos0063240	ACL process ends unexpectedly when switch has clear-flow ACL rule with count interval greater than snmptrap generation timer.
xos0064054	SNMPwalk on extremeAclStatsTable returns value with port instance instead of ifIndex.
xos0064223	Need to add an ACL match condition for matching next-hop addresses during the look-up cycle of a packet, so that actions can be taken based on the next-hop a packet is destined for.
xos0064490	After upgrading from ExtremeXOS 15.2 to later release, last installed dynamic ACL rule is given more priority than previously installed rules.
xos0062145	With QoS configuration, ACL process signal 11 ends unexpectedly after rebooting.
xos0063547	Process ACL ends unexpectedly after applying a policy file with source zone as a match condition.
xos0064129	Policy refresh never completes with network-zone configuration.
AVB	
xos0062494	Source MAC addresses learned through MVRP packets on a blocked port (STP) cause traffic to be dropped.
BGP	
xos0055051	When applying an import policy to BGP, cost configured in the policy is not applied to route tables. This issue is not resolved after multiple policy refreshes nor after multiple disabling, and then enabling BGP.
xos0058441	After creating a BGP peering session between link local IPv6 addresses with the scope ID specified, deleting the VLAN containing link local IPv6 address. and then issuing the command `show configuration bgp`, switch reboots with "`Epm application wdg timer warning`" error message.
xos0061129	In a multi-peer setup with many routes (over 150K), a few routes from the preferred peer do not become active in the BGP RIB. Disabling, and then re-enabling peer, restores all routes.
xos0061411	Route table installs sub-optimal BGP routes (next-hop) to kernel, while the BGP RIB shows different paths when same routes are received from two different peers in local-RIB Note: Disable, and then enable peer or disable, and then enable BGP.
xos0061505	After a topology change in the network, BGP routes requiring two levels of recursive lookup are programmed in hardware with incorrect next hops.
xos0063173	Process dcbgp ends unexpectedly with signal 11 after issuing the command `show bgp neighbor`.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0064319	Aggregated BGP route is not transmitted to upstream neighbor when highest prefix route is received from neighbor.
xos0064589	While learning BGP routes, some routes are not getting installed in route table when deleting and re-adding the static route.
xos0062260	BGP process ends unexpectedly when local address or password is changed for BGP neighbor, and then you immediately execute a BGP show/configuration command.
xos0064496	BGP route policy performs improper community delete operation.
CFM	
xos0063506	Traceroute MAC address in CFM domain does not return information about destination switch.
ClearFlow	
xos0062629	Clearflow rule does not work properly if there is dot(.) in the ACL counter.
Clocking	
xos0063370	PTP Delay Response correction field contains high value (random value) when PTP Delay response packets are passing through combo ports.
xos0062504	You can set a GPTP "peer delay current interval" outside of the correct range of -3 to 17..
Devices	
xos0063429	The output of the command `show fan` shows the fan status as empty after a hot re-seating of the fan module.
xos0062879	Transceiver information shows same Rx power value for 4x10G partition ports even though some ports are in ready state.
xos0064075	The output of the `show fan` command shows fan status as "Failed" after hot re-seating a fan module.
DHCP/BOOTP	
xos0058668	After rebooting DHCPv6, client remains in rebooting state.
xos0058669	DHCPv6 client: After changing the client identifier type, and then restarting the port, old IPv6 addresses are not released, causing the `show vlan` command to show multiple IPv6 addresses.
xos0061219	Parallel-mode-enabled DHCP offer is sent using primary IPv4 address to the client for multiple offers received from server for different IPv4 addresses.
xos0062017	DHCP trusted port configuration is lost after disabling, and then re-enabling LAG.
xos0064151	Error occurs when removing DHCP configuration from VLANs when LAG ports are added to the VLANs.
EAPS	

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0063282	ExtremeXOS CLI restricts PVLAN subscriber VLAN from being configured as an EAPS-protected VLAN.
EDP	
xos0062472	Source MAC addresses learned through CDP packets received on EAPS-blocked ports cause traffic to be dropped.
ELRP	
xos0062460	The `show configuration` command output shows incorrect ELRP configuration.
xos0062618	ELRP forgets the disabled port information if the port is deleted from another VLAN that also has ELRP enabled. As a result, the disabled port stays disabled unless manually enabled.
EMS	
xos0063736	In Syslog, username information appears as "*****" during login/logout cases.
ESRP	
xos0061965	Configuring ESRP member VLANs (VRRP-enabled) produces errors.
FDB	
xos0062789	Disabling learning on LAG ports does not flush FDB entries.
xos0063368	In an MLAG configured switch, FDBs are not installed in hardware after reboot if there are frequent MACMoves between MLAG port and ISC.
xos0059481	Static FDB is programmed incorrectly in hardware after a stack failover.
Identity Management	
xos0061781	Identity manager entries become stale when clients are moved from one port to another in sub-VLANs.
IGMP	
xos0062914	The process mcmgr ends unexpectedly after receiving corrupted IGMPv3 join packets on MLAG ports.
ISIS	
xos0063423	Memory leak occurs in ISIS process when exporting OSPF routes to ISIS.
Layer 2	
xos0064682	Enabling egress VMAN CEP filtering on a CEP port sends a tagged packet, even though it should be forwarded as untagged.
Layer 3	
xos0062710	On BlackDiamond 8800 or BlackDiamond X8 series switches, with Distributed IP ARP mode on, ECMP routes are sometimes not installed when gateways flap.
Multicast	

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0063245	With IGMP per-VLAN mode, VRRP flaps occur after adding tagged ports to VLANs.
xos0064519	With MVR enabled on two VLANs, IGMP report packets are looped if sent to all hosts group.
xos0062705	Kernel oops can occur after clearing IPMC FDB in a stack.
xos0064357	Out of sync between PIM and RTMgr process after introducing new best route.
LACP/LAG	
xos0062428	Member ports with a modified speed configuration that is different than the master port should not be allowed in LAG.
xos0063134	Traffic stops after disabling, and then enabling LAG porst having pstag with static FDB.
xos0064326	LACP flaps when the LAG ports are added to VMANAs, with the VMAN ether type same as LACP ether type.
MLAG	
xos0064009	MLAG+EAPS:Traffic forwarding stops after EAPS that include ISC link converges.
xos0064067	Traffic loss occurs in MLAG setup when ingress port and ISC port reside on different hardware units, and when the internal port number for both of these ports is the same.
xos0056368	Kernel errors occur after disabling sharing configuration on ISC ports of MLAG. For example: "exvlan: handleVsmKernelRequest:8545: handleVsmKernelRequest Invalid Ingress port: 1000008 got"
NetLogin	
xos0061433	extremeNetloginUserLogoutTrap is received with errors.
xos0061868	With protocol order as MAC dot1x, web-based UPM profile is not executed for the client, which is authenticated as MAC.
xos0062674	UPM profile fails to set the variables received from the RADIUS server using VSA 212.
xos0063090	NetLogin client does not move into authfail VLAN when user is absent from local database.
Optics	
xos0063120	Error message "CFP2 modules >= 18 W unsupported" incorrectly appears for Finisar Corp CFP2 LR4 optics.
OSPF	
xos0061855	Configured OSPF neighbor is not retained after rebooting.
Policy	
xos0062965	Policy process ends unexpectedly with signal 6 when master node goes down.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
QoS	
xos0062050	QoS committed rate configurations for port groups are not loaded properly after a save and reboot.
xos0061027	For SummitStacks, creating or deleting non-default QoS profiles may cause some ports to flap.
RSVP-TE	
xos0062380	Switch rejects incorrect LSP configurations as expected, but this operation still uses LSP indexes in hardware.
Sflow/Netflow	
xos0063418	No mapping for Modid errors occur when sFlow is enabled on the port.
SNMP	
xos0057269	SNMP trap extremeIpSecurityViolation is sent with incorrect VLAN description.
xos0061507	SNMPget on EXTREME-SOFTWARE-MONITOR table returns value with incorrect OID.
xos0057212	SNMP traps not sent after changing or saving configuration, even though respective traps are enabled.
xos0063332	Configuration changes to VPLS are not fully retrieved by SNMP walk, which returns values for only few VPLS index.
SSH	
xos0059942	SSH connection ends when show commands produce lengthy output.
STP	
xos0057785	STP domain tag is removed when all ports are deleted from STP auto-bind enabled-VLANs.
xos0062133	STP flush event does not happen after ports are quickly disabled, and then enabled.
xos0062290	Due to ExtremeXOS reflection RSTP BPDU support, upstream bridges believe that they are receiving their own BPDUs (contain the bridge's ID), thus causing multisource events during topology changes, which can cause slow convergence times when lp is configured (upwards of 30 seconds).
xos0063457	Configuration for adding network VLAN port in STP for subscriber is not saved.
xos0063484	Enhancement added in STP flush generation mechanism to reduce hardware programming load.
xos0064395	STP digest value gets changed when adding the port in VLAN or removing the port from VLAN.
xos0062701	HAL timeout occurs while rebooting a stack with STP configuration.
VLANs	

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 16.2 (continued)

CR Number	Description
xos0063331	VLAN IP address is unconfigured when modifying the VLAN name/ port information from Chalet.
xos0063186	Kernel oops occurs when deleting private VLAN.
VMANs	
xos0063274	VLAN packets are egressing with VMAN ethertype when an egress port is deleted from a VMAN that is also part of a VLAN.
xos0063207	Error occurs while adding LAG ports as tagged in one VMAN and untagged in another VMAN, even though the VMAN EtherType is primary.
VPLS/HVPLS	
xos0064033	In BlackDiamond X8 and Summit X670 series switches, traffic gets software forwarded after disabling/enabling members of a shared group and recreating the shared group after deletion.
xos0059596	Can add more than one LSP a pseudo-wire when it is associated with a VPWS.
xos0061092	Traffic forwarding on VPLS-serviced VMAN stops after link flap.
xos0062045	LLDP packets are tunnelled over L2VPn.
xos0062754	VPLS traffic egresses out with dot1q tag when secondary EtherType is configured.
xos0063842	Packets are being flooded in both network and access VLAN ports after port flap.
VR/VRF/L3VPN	
xos0062128	L3VPN traffic is not forwarded after executing `disable port` and `enable port` in MPLS core network.
xos0052723	With L3VPN configured (also: OSPF, BGP, MPLS, LSP) and routes are being advertised and installed in the VRF routing table, after restarting process OSPF, VPN routes are not installed.
xos0061198	Disabling VPN-VRF affects traffic on another VPN-VRF.