



ExtremeXOS Release Notes

Software Version ExtremeXOS 21.1.1-Patch1-5



Copyright © 2018 All rights reserved.

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

Software Licensing

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at:

www.extremenetworks.com/support/policies/software-licensing

Support

For product support, phone the Global Technical Assistance Center (GTAC) at 1-800-998-2408 (toll-free in U.S. and Canada) or +1-408-579-2826. For the support phone number in other countries, visit: <http://www.extremenetworks.com/support/contact/>

For product documentation online, visit: <https://www.extremenetworks.com/documentation/>

Table of Contents

Preface.....	4
Conventions.....	4
Providing Feedback to Us.....	5
Getting Help.....	5
Related Publications.....	6
Chapter 1: Overview.....	7
New and Corrected Features in ExtremeXOS 21.1.....	7
New Hardware Supported in ExtremeXOS 21.1.....	37
Hardware No Longer Supported.....	38
VLAN Option Formatting in Commands.....	38
Circuit Emulation Service (CES) No Longer Supported.....	38
OpenFlow and SSH Included in ExtremeXOS Base Image.....	38
ExtremeXOS SSH Server Upgraded with OpenSSH v6.5.....	39
CLI Command Output Format of Ports Lists.....	39
Extreme Hardware/Software Compatibility and Recommendation Matrices.....	39
Compatibility with Extreme Management Center (Formerly NetSight).....	39
Upgrading ExtremeXOS.....	39
Supported MIBs.....	40
Tested Third-Party Products.....	40
Extreme Switch Security Assessment.....	41
Service Notifications.....	41
Chapter 2: Limits.....	42
Chapter 3: Open Issues, Known Behaviors, and Resolved Issues.....	72
Open Issues.....	72
Known Behaviors.....	76
Resolved Issues in ExtremeXOS 21.1.1-Patch1-5.....	77
Resolved Issues in ExtremeXOS 21.1.1-Patch1-2.....	81
Resolved Issues in ExtremeXOS 21.1.....	83

Preface

This section discusses the conventions used in this guide, ways to provide feedback, additional help, and other Extreme Networks publications.

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.
<i>New!</i>	New Content	Displayed next to new content. This is searchable text within the PDF.

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]
<i>Words in italicized type</i>	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 command documentation (see the Extreme Documentation page at <http://documentation.extremenetworks.com>). 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 sections.

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:

- **GTAC (Global Technical Assistance Center) 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.
- **Extreme Portal** — Search the GTAC knowledge base, manage support cases and service contracts, download software, and obtain product licensing, training, and certifications.
- **The Hub** — A forum for Extreme customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

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 RMA (Return Material Authorization) numbers

Related Publications

ExtremeXOS Publications

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

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

New and Corrected Features in ExtremeXOS 21.1
New Hardware Supported in ExtremeXOS 21.1
Hardware No Longer Supported
VLAN Option Formatting in Commands
Circuit Emulation Service (CES) No Longer Supported
OpenFlow and SSH Included in ExtremeXOS Base Image
ExtremeXOS SSH Server Upgraded with OpenSSH v6.5
CLI Command Output Format of Ports Lists
Extreme Hardware/Software Compatibility and Recommendation Matrices
Compatibility with Extreme Management Center (Formerly NetSight)
Upgrading ExtremeXOS
Supported MIBs
Tested Third-Party Products
Extreme Switch Security Assessment
Service Notifications

These release notes document ExtremeXOS 21.1.1-Patch1-5 which resolves software deficiencies.

New and Corrected Features in ExtremeXOS 21.1

This section lists the new and corrected features supported in the 21.1 software:

Virtual Extensible LAN (VXLAN) Gateway

Virtual Extensible LAN (VXLAN) is a layer 2 overlay scheme over a layer 3 network. Overlays are called VXLAN segments, and only virtual machines (VMs) within the same segment have Layer 2 connectivity. VXLAN segments are uniquely identified using an identifier called the VXLAN Network Identifier (VNI). The VNI is a 24-bit identifier; therefore, an administrative domain can support up to 16 million overlay networks.

As the scope of the MAC addresses originated by tenant VMs is restricted by the VNI, overlapping MAC addresses across segments can be supported without traffic leaking between tenant segments. When a tenant frame traverses a VXLAN overlay network, it is encapsulated by a VXLAN header that contains the VNI. This frame is further encapsulated in a UDP header and L2/L3 headers.

VXLAN can add up to a 54-byte header to the tenant VM's frame. For VXLAN to work correctly, this requires that the IP MTU be set to at least 1554 bytes on the network-side interfaces, and on all transit nodes which carry VXLAN traffic.

The role to encapsulate/decapsulate a frame is performed by a VXLAN Tunnel Endpoint (VTEP), also referred to as VXLAN gateway. A VXLAN gateway can be a Layer 2 gateway or Layer 3 gateway

depending on its capacity. A Layer 2 gateway acts as a bridge connecting VXLAN segments to VLAN segments. A Layer 3 gateway performs all that of Layer 2 gateway, and capable of routing traffic between tenant VLANs.



Note

This feature implements only Layer 2 gateway.

At tunnel initiation, a gateway looks up the destination MAC address of the frame received from the tenant VM. If the MAC address to remote VTEP IP binding is known, the gateway adds the VXLAN header and the IP/UDP header to the frame and forwards toward the DC network. A gateway node that terminates a tunnel removes the encapsulation headers from the packet and determines the bridge domain of the inner frame by examining the VNID received in the VXLAN header. The gateway then looks up the inner MAC destination address (DA) in the tenant VLAN's filtering database and decides either to flood or forward the frame to tenant ports.

The VXLAN segments with the same virtual network ID form a virtual network with one Ethernet broadcast domain.

In multicast VXLAN, the VNI is mapped to a multicast group and multicast tunnels are used to distribute broadcast, unknown unicast and multicast (BUM) tenant traffic to remote endpoints (VTEPs). This requires that the Layer 3 network should support multicast. Unicast VXLAN uses unicast tunnels, and the BUM traffic is head-end replicated at each of the remote endpoints.



Note

This feature implements only unicast VXLAN.

Supported Platforms

Summit X770 and X670-G2 series switches (standalone), and stacks that have X770 and X670-G2 slots only.

Limitations

The following capabilities are not supported in ExtremeXOS 21.1:

- Layer 3 gateways
- Multicast VXLAN
- Ability to assign more than one VNI to a virtual network
- IPv6 addresses for local and remote VTEPs
- Assigning source IP addresses for VXLAN gateway encapsulation:
 - Per virtual router
 - Per virtual network or VNI
- Support for adding more than one tenant VLAN per VNI
- A physical port being part of both a tenant VLAN and an underlay (network) VLAN
- Routing in and out of tunnels
- Integration with any controllers
- Support for heterogeneous stack environments where at least one of the stack nodes is not VXLAN capable
- More than one next hop per (network) hop
- Tagged and untagged tenant VLANs on the same port

- Multicast underlay IP network, including PIM-Bidir
- Multiple VRs

New CLI Commands

```

create virtual-network vn_name {flooding [standard | explicit-remotes]}

configure virtual-network vn_name vxlan vni [ vni | none ]

configure virtual-network vn_name [add | delete] [{vlan} vlan_name | vman
vman_name]

configure virtual-network local-endpoint [ ipaddress ipaddress { vr
vr_name } | none ]

create virtual-network remote-endpoint vxlan ipaddress ipaddress {vr
vr_name}

delete virtual-network remote-endpoint vxlan ipaddress ipaddress {vr
vr_name}

configure virtual-network vn_name [add | delete] remote-endpoint vxlan
ipaddress ipaddress {vr vr_name}

enable learning {forward-packets | drop-packets}] vxlan {vr vr_name}
ipaddress remote_ipaddress

disable learning {forward-packets | drop-packets}] vxlan {vr vr_name}
ipaddress remote_ipaddress

show virtual-network { vn_name | vxlan vni vni | [vlan vlan_name | vman
vman_name]}

show virtual-network {vn_name} remote-endpoint vxlan {vni vni} {ipaddress
ipaddress { vr vr_name } }

configure fdb { mac_addr | broadcast | unknown-unicast | unknown-
multicast } vlan vlan_name [ add | delete ] vxlan {vr vr_name }
{ipaddress} remote_ipaddress

configure virtual-network remote-endpoint vxlan ipaddress ipaddress { vr
vr_name } monitor [ on | off ]

show virtual-network { vn_name | remote-endpoint vxlan {ipaddress
ipaddress} {vr vr_name}} statistics {no-refresh}

clear counters virtual-network remote-endpoint vxlan [ all | ipaddress
ipaddress]

configure virtual-network vn_name monitor [ on | off ]

show virtual-network { vn_name | remote-endpoint remote-endpoint vxlan
{ipaddress ipaddress} {vr vr_name}} statistics {no-refresh}

```

```
clear counters virtual-network [ all | vn_name ]
```

Changed CLI Commands

Changes are underlined.

```
[ create | delete ] fdb [ mac_addr vlan vlan_name [ ports port_list |
blackhole | vxlan { vr vr_name } {ipaddress} remote_ipaddress ] |
broadcast vlan vlan_name vxlan { vr vr_name } {ipaddress}
remote_ipaddress | unknown-multicast vlan vlan_name vxlan { vr vr_name }
{ipaddress} remote_ipaddress | unknown-unicast vlan vlan_name vxlan { vr
vr_name } {ipaddress} remote_ipaddress ]
```

```
show fdb { {mac_addr | blackhole | permanent | {vlan} vlan_name | ports
port_list} {netlogin [all | mac-based-vlans]} | {vpls} {vpls_name} |
openflow | rbridge {nickname} | vxlan { vni } | virtual-network vn_name }
```

```
create vlan vlan-name {vr vr-name} {description vlan-desc} {tag [ tag |
none ] }
```

```
configure {vlan} vlan-name {tag [tag {remote-mirroring} | none] }
```

```
configure {vlan} vlan_name add ports [port_list | all] {tagged {tag {-
end tag}} | untagged | private-vlan translated }
```

```
configure {vlan} vlan_name delete ports [port_list | all] {tagged {tag}
{- end tag}}
```

Open Shortest Path First (OSPF) Exchanging Information for Virtual Extensible LAN (VXLANS)

ExtremeXOS leverages Open Shortest Path First (OSPF) to advertise and learn VTEPs dynamically in a VXLAN network. OSPFv2 advertises the triplet of VNI/Endpoint IP Address/Advertising Router ID through OSPFv2 domain using type 11 opaque link state advertisements (LSAs). The OSPFv2 VXLAN LSA link state ID uses opaque type 128. The remaining 24 bits of the field are set to the VXLAN VNI. Each locally configured VNI corresponds to a single opaque LSA advertised by the router. The OSPFv2 VXLAN LSA payload contains one top level TLV that specifies the locally configured IPv4 endpoint address on the advertising router.

Note



- The remote endpoints learned using OSPF are not saved to the configuration.
- The OSPFv2 VXLAN opaque LSA is only advertised if OSPF VXLAN extensions are enabled.
- OSPF VXLAN extensions can only be enabled when OSPFv2 is disabled.
- Local endpoint address can only be IPv4. IPv6 is not supported.

Supported Platforms

Summit X770 and X670-G2 series switches (standalone), and stacks that have X770 and X670-G2 slots only.

New CLI Commands

```
enable ospf vxlan-extensions
```

```
disable ospf vxlan-extensions
```

Changed CLI Commands

The `show ospf` command output has been changed (shown in bold):

```
show ospf

OSPF                : Enabled                MPLS LSP as Next-Hop: No
RouterId            : 192.168.170.60        RouterId Selection  : Automatic
ASBR                : No                    ABR                 : No
ExtLSA              : 0                    ExtLSAChecksum     : 0x0
OriginateNewLSA    : 190                  ReceivedNewLSA     : 102
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 : Enabled
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
```

ONEPolicy Now Supported on New ExtremeSwitching X440-G2 and X620 Series Switches

ONEPolicy, which was released in ExtremexXOS 16.1, allows you create profiles for securing and provisioning network resources based upon the role the user or device plays within the enterprise. By first defining the user or device role, network resources can be tailored to a specific user, system, service, or port-based context by configuring and assigning rules to the policy role. A policy role can be configured for any combination of Class of Service, VLAN assignment, classification rule precedence, or default behavior based upon L2, L3, and L4 packet fields. Hybrid authentication allows either policy or dynamic VLAN assignment, or both, to be applied through RADIUS authorization.

This feature is now supported on the new ExtremeSwitching X440-G2 and X620 series switches.

Supported Platforms

Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

Limitations

- When stacking switches that have different capacities, the stack goes to the lowest common level of capacities and functionality when possible. If the stack already has an existing configuration that exceeds the new lower capacity, policy disallows the ports on the new switch to become policy-enabled.
- Only 'macdest', 'macsource', or 'port' policy rules can be applied to QinQ (that is, double-tagged) packets received on an untagged VMAN port.

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 of 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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

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
-----  -----  -----  -----  -----
----  -----  -----  -----  -----
```

```

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 router 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
- Saves bandwidth on the links connecting master and backup routers

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

Platform

Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

Limitations

- Fabric Routing feature will not be supported for VRRP VR for which Virtual IP is same as interface IP (owned IP).
- Traffic sent from host destined for VIP, will be 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 will expect TTL value 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.

Platform

Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

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, (el) 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 Purview to provide additional information about flows. IPFIX can detect flows and collect flow statistics, but it cannot do deep packet payload inspections. Purview, 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 Purview is able to receive the packets for deep packet inspection.

Supported Platforms

Summit X460-G2 series switches

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
```

Border Gateway Protocol (BGP) Data Center Enhancements

The following Border Gateway Protocol (BGP) data center enhancements are now available:

- Sixty-four equal cost (ECMP) paths for BGP (previously eight).
- Support for maximum autonomous system path (AS-path) length filtering of BGP route updates.
- IPv4 peering sessions can carry IPv6 routes, and IPv6 peering sessions can carry IPv4 routes for the Unicast and Multicast sub-address families.

Supported Platforms

Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

Limitations

- Support for maximum AS-Path length is on a BGP instance basis, not per peer.
- Enabling the capability to carry IPv6 Network Layer Reachability Information (NLRI) over IPv4 peering sessions and IPv4 NLRI over IPv6 sessions does not include the ability to have mismatching next-hops. You must use outbound route-policy to specify the BGP next-hop value to be a reachable subnet for the remote router or the remote router must have a means to reach the next-hop. For IPv6 NLRI carried over IPv4 peering sessions, in the absence of route policy to set the next-hop, the next-hop is automatically set to the mapped IPv6 address based on the IPv4 address of the outgoing interface. You should either override this with policy or program the downstream router with a static route to reach the mapped address. In either case, it is assumed the VLAN interface used for peering is configured with both IPv4 and IPv6 addresses.
- The ability to carry mismatching NLRI applies to the Unicast, Multicast, and VPNv4 Sub-Address-Families (SAFIs). The VPNv6 SAFI is not supported.

New CLI Commands

```
configure bgp maximum-as-path-length max-as-path
```

Changed CLI Commands

The following command now accepts 64 ECMP paths for *max-paths*:

```
configure bgp maximum-paths max-paths
```

The following commands now allows you to enable the capability to carry NLRI of address family indicator (AFI)/SAFI combinations even if the specified AFI does not match the address family of the peering sessions:

```
enable bgp neighbor ipv4 capability ipv6-unicast
```

```
enable bgp neighbor ipv6 capability ipv4-unicast
```

Bidirectional Forwarding Detection (BFD) for the Border Gateway Protocol (BGP)

Bidirectional Forwarding Detection (BFD) protection of Border Gateway Protocol (BGP) peering sessions allows for the rapid detection of link failures such that peering sessions can be taken out of the "established" state within fractions of a second. This allows the protocol to select an alternate path (if available) to a destination immediately after the link failure, rather than waiting until the BGP hold timer expires (180 seconds by default). This feature applies to both IPv4 and IPv6 peering sessions. Both IPv6 global and link local peering sessions are supported.

Supported Platforms

Summit X460-G2, X670-G2, X770 series switches, with Core License or above.

Limitations

- The BFD setting can be applied on a per-peer basis, but the ability to set BFD on a peer-group or address-family basis is not currently supported.
- The BGP peer must be in the disabled admin state to modify its BFD setting.
- While BFD can be enabled on any BGP peering session, protection is only provided for directly connected EBGp peering sessions.

New CLI Commands

```
configure bgp {neighbor [all|remoteaddr]} {bfd [on | off]}
```

Changed CLI Commands

The show bgp neighbor command now shows BFD information (shown in bold):

```
show bgp neighbor 192.168.24.2
Peer Description      :
EBGP Peer            : 192.168.24.2      AS           : 300
Enabled              : Yes              OperStatus   : Up
Weight              : 1                Shutdown-Priority : 1024
ConnectRetry        : 120              MinAsOrig    : 30
HoldTimeCfg         : 180              KeepaliveCfg  : 60
Source Interface    : Not configured   RRClient     : No
EBGP-Multihop       : No               Remove Private AS : No
BFD                : Off              BFD Status   : Inactive
```

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 21.1. ExtremeXOS 21.1 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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

Link Aggregation Control Protocol (LACP) Fallback Option

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 feature 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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

Limitations

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

New CLI Commands

```
configure sharing portlacp 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 Sys-Pri	Actor Key	Partner MAC	Partner Sys-Pri	Partner Key	Agg Count	Actor 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	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner 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	Port Priority	Rx State	Sel Logic	Mux State	Actor Flags	Partner Port
18	5	Initialize	Fallback	Collect-Dist	A-GSCD--	1018

```

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

```

Hardware Assisted Bidirectional Forwarding Detection (BFD)

The hardware assisted Bidirectional Forwarding Detection (BFD) feature expands on the existing ExtremeXOS BFD capabilities.

Bidirectional Forwarding Detection (BFD) hardware assist support provides the functionality to run a BFD session in hardware. Effective failure detection requires BFD to run at high frequencies (using aggressive timers as low as 3 ms), which is not possible in the software mode because of CPU and ExtremeXOS restrictions.

To make BFD sessions run in the hardware, the following configuration is required.

- Unused front panel port (not available for switching the user data traffic) configured as a loopback port. The port is used internally by the BFD hardware to send control packets.
- IPforwarding is enabled on the BFD interfaces.
- Nexthop MAC address of neighbor should be known for the session creation. BFD process triggers ARP to resolve the next hop MAC address, if not configured statically.

Supported Platforms

- Summit X460-G2 series switches, standalone only

New CLI Commands

```
configure bfd hardware-assist [primary | secondary] loopback-port [ port
| none]
```

Changed CLI Commands

The following show commands are changed to show the hardware assist information (shown in bold):

```
#show bfd
Number of sessions           : 0
Sessions in Init State      : 0
```

```

Sessions in Down State           : 0
Sessions in Admin Down State     : 0
Sessions in Up State             : 0

SNMP Traps for session-down     : Disabled
SNMP Traps for session-up       : Disabled
SNMP Traps for Batch Delay      : 1000 ms
Hardware Assist Operational State : Disabled(Loopback port not configured)
Hardware Assist Primary Loopback Port : 1
Hardware Assist Secondary Loopback Port : None
Maximum # of Hardware Assist Sessions : 900

```

```

# show bfd session detail vr all
  Neighbour      : 10.10.10.1           Local      : 10.10.10.2
  Vr-Name        : bfd_vr10            Interface  : bfd_vlan10
  Session Type   : Single Hop          State      : Up
  ...
  Up Count       : 1
  Last Valid Packet Rx : 00:51:49.300000
  Last Packet Tx   : 00:51:48.8200000
  Hardware Assist : Yes

  Neighbour      : 10.10.11.1           Local      : 10.10.11.2
  Vr-Name        : bfd_vr10            Interface  : bfd_vlan11
  Session Type   : Single Hop          State      : Up
  ...
  Up Count       : 1
  Last Valid Packet Rx : 00:51:49.300000
  Last Packet Tx   : 00:51:48.8200000
  Hardware Assist : Yes

```

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

The 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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

New CLI Commands

```
configure security fips-mode [on | off]
```

```
show security fips-mode
```

CE2.0 Certification Additions

This features adds CE2.0 (previously known as MEF) certification. This certification involves the following changes:

- Removal of the preamble and interframe gap (IFG) overhead for the rate policing and shaping functions
- Support for meter out-of-profile action for setting a specified 802.1p value
- Support for ACL match criteria “ccos” for matching customer 802.1p on UNI or NNI ports

Supported Platforms

Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

New CLI Commands

configure forwarding **rate-limit** **overhead-bytes** *overhead_bytes*

Changed CLI Commands

Changes are underlined.

```
configure meter metername [{committed-rate circ committed-rate-unit {max-burst-size burst-size [Kb | Mb | Gb | packets]}]}] {out-actions [{disable-port} {drop | set-drop-precedence {dscp [dscp-value | none]}] {dot1p [dot1p-value | none]}]}] {log} {trap]}] {ports [port_group | port_list] }
```

The output of the show forwarding configuration command now shows rate limit information (shown in bold):

```
# show forwarding configuration
L2 and L3 Forwarding table hash algorithm:
  Configured hash algorithm:      crc32
  Current hash algorithm:        crc32
L3 Dual-Hash configuration: (Applies to "c", "xl"-series and 8900-40GX-xm)
  Configured setting:            on
  Current setting:              on
  Dual-Hash Recursion Level:    1
Hash criteria for IP unicast traffic for L2 load sharing and ECMP route sharing
  Sharing criteria:              L3_L4
IP multicast:
  Group Table Compression:      on
  Local Network Forwarding:    slow-path
  Lookup-Key:                   (SourceIP, GroupIP, VlanId)
External lookup tables:
  Configured Setting:           12-and-13
  Current Setting:              12-and-13
Switch Settings:
  Switching mode:               store-and-forward
L2 Protocol:
  Fast convergence:            on

Rate Limit:
  Overhead Bytes:          20
Fabric Flow Control:
  Fabric Flow Control:         auto
```

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 need be specified when making gPTP configurations. However, the protocol is still running on each member port at the physical port level.

Supported Platforms

Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 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 Speed	BW Mode	Percentage
*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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

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 method is 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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 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 | portlists]
```

```
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 is enabled, only the specified port CVID and CVIDs are allowed to egress.

Supported Platforms

Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

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]
```

Resilient Hashing

Resilient Hashing is a hardware-based capability that minimizes the remapping of flows to aggregator member ports during aggregator member changes.

In conventional hashing, physical links are used to form fat logical pipes. The static hash scheme associates a flow with a physical link. When a link fails, even flows that did not originally flow through the failed link may be assigned to a new link. This reassignment may temporarily result in out-of-order packet deliver even for the flows that were not using the failed link. In contrast, a resilient hashing scheme associates flows with physical ports. When a link fails, only the affected flows are redistributed uniformly across the remaining good physical links. Flows using functioning links remain unaffected and are not reassigned to new links.

Supported Platforms

Summit X770 and X670-G2, and on SummitStacks when at least one of the supported switches is included in the stack.

On SummitStacks, configuration of resilient hashing is not allowed unless at least one node in the stack supports resilient hashing. In a stack where one or more nodes support resilient hashing and one or more nodes do not support resilient hashing, resilient hashing is only in effect for flows received on ports on nodes where resilient hashing is supported by the hardware.

Limitations

- Resilient hashing is available only on LAGs configured to use the “custom” distribution algorithm.
- Resilient Hashing applies only to the distribution of known unicast traffic.
- Traffic originating or forwarded by the system CPU is not distributed using Resilient Hashing.

New CLI Commands

```
configure sharing master_port resilient-hashing [on | off]
```

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 | port-lists]} {resilient-hashing [on | off]} {lacp | health-check}
```

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 ASBR connected to an NSSA can be advertised within the area and can be advertised to other areas as AS-external LSAs.

Supported Platforms

Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

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
Router Alert : Disabled
ASExternal LSALimit : Disabled
Originate Default : Disabled
Graceful Restart : Both
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
100000M Cost (100G) : 10
Timeout (Count) : Disabled (0)
Grace Period : 120s

```

```
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 lsdbs 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 21.1 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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

Limitations

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

ExtremeXOS Applications Environment

ExtremeXOS 21.1 introduces an environment that allows management applications, controllable through a web interface, that 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 [ezServiceability \(File Upload/Download\)](#) on page 37).

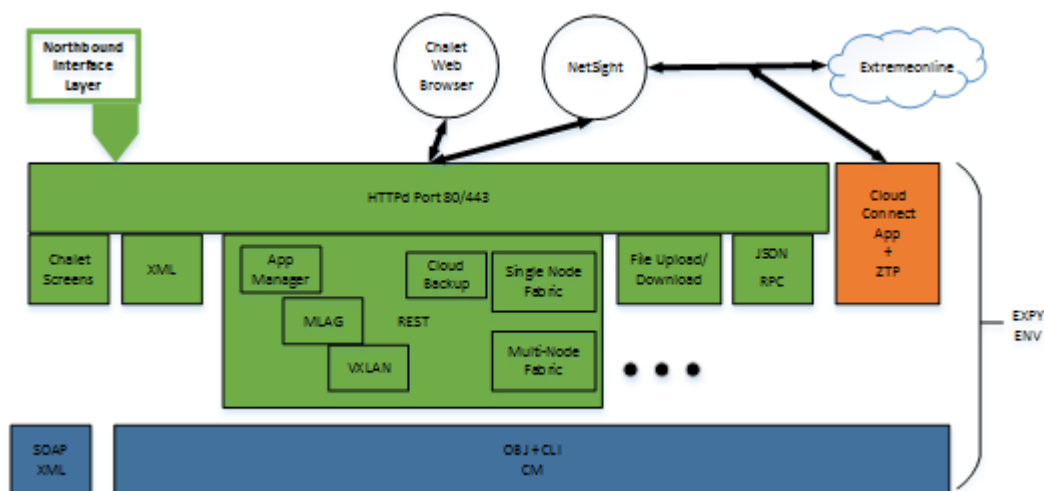


Figure 1: Application Environment Block Diagram

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 21.1:

- Service applications.
- File upload/download (see [ezServiceability \(File Upload/Download\)](#) on page 37)
- 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 X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 series switches

ezServiceability (File Upload/Download)

ezServiceability 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`, `xf`, `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.
- Retrieves/edits/returns files to a switch (provides a user-friendly way of editing files).
- Deleting files on a switch.

New Hardware Supported in ExtremeXOS 21.1

This section lists the new hardware supported in ExtremeXOS 21.1:

- ExtremeSwitching X440-G2 series switches:

X440-G2-24t-10GE4, X440-G2-24t-10GE4-DC, X440-G2-24p-10GE4, X440-G2-48t-10GE4, X440-G2-48t-10GE4-DC, X440-G2-48p-10GE4, X440-G2-12t-10GE4, X440-G-12p-10GE4, X440-G2-24x-10GE4, X440-G2-24fx-GE4, X440-G2-12t8fx-GE4, X440-G2-24t-GE4



Note

ExtremeSwitching X440-G2 10 Gigabit model switches require a license to upgrade the four SFP 1GbE ports to 10G. For more information, see *ExtremeXOS 21.1 Feature License Requirements*.

- ExtremeSwitching X620 series switches:
X620-10X, X620-8T-2X, X620-16X, X620-16T

Hardware No Longer Supported

The following hardware is no longer supported in ExtremeXOS 21.1:

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



Note

These hardware platforms *are* supported in the ExtremeXOS 16.x software.

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
```

Circuit Emulation Service (CES) No Longer Supported

Starting with ExtremeXOS 21.1, Circuit emulation service (CES) is no longer supported.

OpenFlow and SSH Included in ExtremeXOS Base Image

OpenFlow and SSH are now included in the ExtremeXOS base image starting with ExtremeXOS 21.1. A separate XMOD file is no longer required.

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 21.1, 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 21.1 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 21.1 or later to previous releases, with RSA key saved, the key becomes invalid.

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”

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 pluggable transceivers and cables.

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 21.1 is compatible with Extreme Management Center (formerly NetSight) version 7.0 and later.

Upgrading ExtremeXOS

For instructions about upgrading ExtremeXOS software, see “Software Upgrade and Boot Options” in the *ExtremeXOS 22.3 User Guide*.

Beginning with ExtremeXOS 12.1, an ExtremeXOS core image (.xos file) must be downloaded and installed on the alternate (non-active) partition. If you try to download to an active partition, the error message `Error: Image can only be installed to the non-active partition.` appears. An ExtremeXOS modular software package (.xmod file) can still be downloaded and installed on either the active or alternate partition.

Supported MIBs

The Extreme Networks MIBs are located at <http://www.extremenetworks.com/support/policies/mibs/>.

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

For detailed information on which MIBs and SNMP traps are supported, see the *Extreme Networks Proprietary MIBs* and *MIB Support Details* sections in the *ExtremeXOS 21.1 User Guide*.

Tested Third-Party Products

This section lists the third-party products tested for ExtremeXOS 21.1.1-Patch1-5.

Tested RADIUS Servers

The following RADIUS servers are fully tested:

- Microsoft—Internet Authentication Server
- Meetinghouse
- FreeRADIUS

Tested Third-Party Clients

The following third-party clients are fully tested:

- Windows 7
- Windows Vista
- Linux (IPv4 and IPv6)
- Windows XP (IPv4)

PoE Capable VoIP Phones

The following PoE capable VoIP phones are fully tested:

- Avaya 4620
- Avaya 4620SW IP telephone
- Avaya 9620
- Avaya 4602
- Avaya 9630
- Avaya 4621SW
- Avaya 4610
- Avaya 1616
- Avaya one-X
- Cisco 7970
- Cisco 7910
- Cisco 7960
- ShoreTel ShorePhone IP 212k

- ShoreTel ShorePhone IP 560
- ShoreTel ShorePhone IP 560g
- ShoreTel ShorePhone IP 8000
- ShoreTel ShorePhone IP BB 24
- Siemens OptiPoint 410 standard-2
- Siemens OpenStage 20
- Siemens OpenStage 40
- Siemens OpenStage 60
- Siemens OpenStage 80

Extreme Switch Security Assessment

DoS Attack Assessment

Tools used to assess DoS attack vulnerability:

- Network Mapper (NMAP)

ICMP Attack Assessment

Tools used to assess ICMP attack vulnerability:

- SSPing
- Twinge
- Nuke
- WinFreeze

Port Scan Assessment

Tools used to assess port scan assessment:

- Nessus

Service Notifications

To receive proactive service notification about newly released software or technical service communications (for example, field notices, product change notices, etc.), please register at: www.extremenetworks.com/support/service-notification-form

2 Limits

This chapter summarizes the supported limits in ExtremeXOS 21.1.1-Patch1-5.

Table 3 summarizes tested metrics for a variety of features, as measured in a per-system basis unless otherwise noted. These limits may change, but represent the current status. The contents of this table supersede any values mentioned in the ExtremeXOS books.

The scaling and performance information shown in **Table 3** is provided for the purpose of assisting with network design. It is recommended that network architects and administrators design and manage networks with an appropriate level of network scaling “head room.” The scaling and performance figures provided have been verified using specific network topologies using limited switch configurations. There is no guarantee that the scaling and performance figures shown are applicable to all network topologies and switch configurations and are provided as a realistic estimation only. If you experience scaling and performance characteristics that you feel are sufficiently below what has been documented, contact Extreme Networks technical support for additional assistance.

The route limits shown in **Table 3** for IPv4 and IPv6 routing protocols are software limits only. The actual hardware limits may be higher or lower than the software limits, based on platform. The hardware limits for specific platforms are specified as “IPv4/IPv6 routes (LPM entries in hardware)” in the following table.

It is not advised to have greater than 25,000 total IP routes from all routing protocols. Adverse effects can occur with routing tables larger than this, especially when a single network event or CLI command affects a significant number of routes. For example, just after such a network event, the added system load will cause a save configuration command to time out.

Table 3: Supported Limits

Metric	Product	Limit
AAA (local) —maximum number of admin and local user accounts.	All platforms	8
Access lists (meters) —maximum number of meters.	ExtremeSwitching X620, X440-G2	1,024 ingress, 256 egress
	Summit X770, X670-G2	1,024 ingress, 512 egress
Access lists (policies) —suggested maximum number of lines in a single policy file.	All platforms	300,000
Access lists (policies) —maximum number of rules in a single policy file. ^a	Summit X460-G2, X450-G2, X770, X670-G2	4,096 ingress, 1,024 egress
	ExtremeSwitching X620, X440-G2	2,048 ingress, 512 egress

Table 3: Supported Limits (continued)

Metric	Product	Limit
Access lists (policies) —maximum number of rules in a single policy file in first stage (VFP).	Summit X450-G2, X460-G2	2,048 ingress only
	Summit X670-G2, X770	1,024 ingress only
	ExtremeSwitching X620, X440-G2	512 ingress only
Access lists (slices) —number of ACL slices.	Summit X460-G2, X450-G2	16 ingress, 4 egress
	Summit X770, X670-G2	12 ingress, 4 egress
	ExtremeSwitching X440-G2, X620	8 ingress, 4 egress
Access lists (slices) —number of ACL slices in first stage (VFP).	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	4 ingress only
ACL Per Port Meters —number of meters supported per port.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	16
Meters Packets-Per-Second Capable	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	Yes
AVB (audio video bridging) —maximum number of active streams.	Summit X450-G2, X460-G2, X770, and ExtremeSwitching X620, X440-G2 Summit X670-G2	1,024 4,096
BFD sessions (Software Mode) —maximum number of BFD sessions.	Summit X460-G2, X670-G2, X450-G2, X770 (default timers—1 sec)	512
	Summit X460-G2, X670-G2, X450-G2, X770 (minimal timers—100 msec)	10 ^c
BFD sessions (Hardware Assisted) —maximum number of BFD sessions.	Summit X460-G2	900 (PTP not enabled) 425 (PTP enabled) 256 (with 3 ms transmit interval)
BGP (aggregates) —maximum number of BGP aggregates.	Summit X460-G2, X670-G2, X770 with Core license or higher Summit X450-G2, and ExtremeSwitching X440-G2, X620	256 Not supported
BGP (networks) —maximum number of BGP networks.	Summit X460-G2, X670-G2, X770 with Core license or higher Summit X450-G2, and ExtremeSwitching X440-G2, X620	1,024 Not supported
BGP (peers) —maximum number of BGP peers. Note: *With default keepalive and hold timers.	Summit X460-G2, X670-G2, X770 with Core license or higher Summit X450-G2, and ExtremeSwitching X440-G2, X620	128* Not supported
BGP (peer groups) —maximum number of BGP peer groups.	Summit X460-G2, X670-G2, X770 with Core license or higher Summit X450-G2, and ExtremeSwitching X440-G2, X620	64 Not supported
BGP (policy entries) —maximum number of BGP policy entries per route policy.	Summit X460-G2, X670-G2, X770 with Core license or higher Summit X450-G2, and ExtremeSwitching X440-G2, X620	256 Not supported

Table 3: Supported Limits (continued)

Metric	Product	Limit
BGP (policy statements) —maximum number of BGP policy statements per route policy.	Summit X460-G2, X670-G2, X770 with Core license or higher Summit X450-G2, and ExtremeSwitching X440-G2, X620	1,024 Not supported
BGP multicast address-family routes —maximum number of multicast address-family routes.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	25,000 Not supported
BGP (unicast address-family routes) —maximum number of unicast address-family routes.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	25,000 Not supported
BGP (non-unique routes) —maximum number of non-unique BGP routes.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	25,000 Not supported
BGP ECMP —maximum number of equalcost multipath for BGP and BGPv6.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	2, 4, or 8 Not supported
BGPv6 (unicast address-family routes) —maximum number of unicast address family routes.	Summit X460-G2 Summit X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	6,000 8,000 Not supported
BGPv6 (non-unique routes) —maximum number of non-unique BGP routes.	Summit X460-G2 Summit X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	18,000 24,000 Not supported
BOOTP/DHCP relay —maximum number of BOOTP or DHCP servers per virtual router. Note: User VRs not supported.	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2*, X620*	4
BOOTP/DHCP relay —maximum number of BOOTP or DHCP servers per VLAN.	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620	4
Connectivity fault management (CFM) —maximum number of CFM domains. Note: With Advanced Edge license or higher.	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620	8
CFM —maximum number of CFM associations. Note: With Advanced Edge license or higher.	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620	256

Table 3: Supported Limits (continued)

Metric	Product	Limit
<p>CFM—maximum number of CFM up end points.</p> <p>Note: With Advanced Edge license or higher.</p>	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620	32
<p>CFM—maximum number of CFM down end points.</p> <p>Note: With Advanced Edge license or higher.</p>	Summit X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620 Summit X460-G2	32 256 (non-load shared ports) 32 (load shared ports)
<p>CFM—maximum number of CFM remote end points per up/down end point.</p> <p>Note: With Advanced Edge license or higher.</p>	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620	2,000
<p>CFM—maximum number of dot1tag ports.</p> <p>Note: With Advanced Edge license or higher.</p>	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620	128
<p>CFM—maximum number of CFM segments.</p> <p>Note: With Advanced Edge license or higher.</p>	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620	1,000
<p>CFM—maximum number of MIPs.</p> <p>Note: With Advanced Edge license or higher.</p>	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X620, X440-G2	256
<p>CLEAR-Flow—total number of rules supported. The ACL rules plus CLEAR-Flow rules must be less than the total number of supported ACLs.</p>	Summit X460-G2, X770, X670-G2, X450-G2 ExtremeSwitching X440-G2, X620	4,094 1,024
<p>Data Center Bridging eXchange (DCBX) protocol Type Length Value (TLVs)—maximum number of DCBX application TLVs.</p>	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X440-G2, X620	8

Table 3: Supported Limits (continued)

Metric	Product	Limit
DHCPv6 Prefix Delegation Snooping —Maximum number of DHCPv6 prefix delegation snooped entries.	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X620, X440-G2	256 (with Underlying Protocol Ripng) 128 (with Underlying protocol OSPFv3) 1,024 (with static routes)
DHCP snooping entries —maximum number of DHCP snooping entries.	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X620, X440-G2	2,048
Dynamic ACLs —maximum number of ACLs processed per second. Note: Limits are load dependent.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 with 50 DACLs with 500 DACLs	10 5
EAPS domains —maximum number of EAPS domains. Note: An EAPS ring that is being spatially reused cannot have more than four configured EAPS domains.	Summit X670-G2, X450-G2, and X770 Summit X460-G2, and ExtremeSwitching X440-G2, X620	64 32
EAPSV1 protected VLANs —maximum number of protected VLANs.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1,000
EAPSV2 protected VLANs —maximum number of protected VLANs.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620 ExtremeSwitching X440-G2	500 Not supported
ELSM (vlan-ports) —maximum number of VLAN ports.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620	5,000
ERPS domains —maximum number of ERPS domains without CFM configured.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620	32
ERPS domains —maximum number of ERPS domains with CFM configured.	Summit X450-G2, X670-G2, X770, and ExtremeSwitching X620 Summit X460-G2	16 32
ERPSV1 protected VLANs —maximum number of protected VLANs.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1,000
ERPSV2 protected VLANs —maximum number of protected VLANs.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	500
ESRP groups —maximum number of ESRP groups.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X440-G2, X620	31

Table 3: Supported Limits (continued)

Metric	Product	Limit
ESRP domains —maximum number of ESRP domains.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	64
ESRP VLANs —maximum number of ESRP VLANs.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1,000
ESRP (maximum ping tracks) —maximum number of ping tracks per VLAN.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8
ESRP (IP route tracks) —maximum IP route tracks per VLAN.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8
ESRP (VLAN tracks) —maximum number of VLAN tracks per VLAN.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1
Forwarding rate —maximum L3 software forwarding rate.	Summit X770 Summit X670-G2 Summit X460-G2 Summit X450-G2 ExtremeSwitching X440-G2 ExtremeSwitching X620	11,000 pps 21,000 pps 25,000 pps 24,000 pps 21,000 pps 23,000 pps
FDB (unicast blackhole entries) —maximum number of unicast blackhole FDB entries.	Summit X460-G2 Summit X770, X670-G2 Summit X450-G2 ExtremeSwitching X620, X440-G2	49,152 ^f 294,912 ^f 34,816 ^f 16,384 ^f
FDB (multicast blackhole entries) —maximum number of multicast blackhole FDB entries.	Summit X460-G2, X450-G2, and ExtremeSwitching X440-G2, X620 Summit X770, X670-G2	1,024 4,096
FDB (maximum L2 entries) —maximum number of MAC addresses.	Summit X670-G2 Summit X460-G2 Summit X770 Summit X450-G2 ExtremeSwitching X620, X440-G2	294,912 ^f 98,300 ^f 294,912 ^f 68,000 ^f 16,384 ^f
FDB (Maximum L2 entries) —maximum number of multicast FDB entries.	Summit X770, X670-G2 Summit X450-G2, X460-G2, and ExtremeSwitching X620, X440-G2	4,096 1,024
Identity management —maximum number of Blacklist entries.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	512
Identity management —maximum number of Whitelist entries.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	512
Identity management —maximum number of roles that can be created.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	64

Table 3: Supported Limits (continued)

Metric	Product	Limit
Identity management —maximum role hierarchy depth allowed.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	5
Identity management —maximum number of attribute value pairs in a role match criteria.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	16
Identity management —maximum of child roles for a role.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8
Identity management —maximum number of policies/dynamic ACLs that can be configured per role.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8
Identity management —maximum number of LDAP servers that can be configured.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8
Identity management —maximum number of Kerberos servers that can be configured.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	20
Identity management —maximum database memory-size.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	512
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.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	100
Identity management —recommended number of ACL entries per identity. Note: Number of ACLs per identity based on system ACL limitation.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	20
Identity management —maximum number of dynamic ACL entries configured as an individual dynamic rule, or as an ACL entry in a policy file.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	500

Table 3: Supported Limits (continued)

Metric	Product	Limit
IGMP snooping per VLAN filters —maximum number of VLANs supported in per-VLAN IGMP snooping mode.	Summit X460-G2	1,500
	Summit X450-G2	2,048
	Summit X770, X670-G2	2,000
	ExtremeSwitching X620, X440-G2	1,000
IGMPv1/v2 SSM-map entries —maximum number of IGMPv1/v2 SSM mapping entries.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	500
IGMPv1/v2 SSM-map entries —maximum number of sources per group in IGMPv1/v2 SSM mapping entries.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	50
IGMPv2 subscriber —maximum number of IGMPv2 subscribers per port. ⁿ	Summit X770, X670-G2, X460-G2, X450-G2	4,000
	ExtremeSwitching X440-G2, X620	3,500
IGMPv2 subscriber —maximum number of IGMPv2 subscribers per switch. ⁿ	Summit X770, X670-G2	30,000
	Summit X460-G2, X450-G2	20,000
	ExtremeSwitching X620, X440-G2	17,500
IGMPv3 maximum source per group —maximum number of source addresses per group.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	250
IGMPv3 subscriber —maximum number of IGMPv3 subscribers per port. ⁿ	Summit X770, X670-G2, X460-G2, X450-G2	4,000
	ExtremeSwitching X440-G2, X620	3,500
IGMPv3 subscriber —maximum number of IGMPv3 subscribers per switch. ⁿ	Summit X460-G2, X450-G2	20,000
	Summit X770, X670-G2	30,000
	ExtremeSwitching X620, X440-G2	17,500
IP ARP entries in software —maximum number of IP ARP entries in software. Note: May be limited by hardware capacity of FDB (maximum L2 entries).	Summit X670-G2, X770	131,072 (up to) ^h
	Summit X460-G2	57,344 (up to) ^h
	Summit X450-G2	47,000 (up to) ^h
	ExtremeSwitching X440-G2, X620	20,480
IPv4 ARP entries in hardware with minimum LPM routes —maximum recommended number of IPv4 ARP entries in hardware, with minimum LPM routes present. Assumes number of IP route reserved entries is 100 or less.	Summit X460-G2	50,000 (up to) ^h
	Summit X770, X670-G2	108,000 (up to) ^h
	Summit X450-G2	39,000 (up to) ^h
	ExtremeSwitching X620	1,500
	ExtremeSwitching X440-G2	1,000

Table 3: Supported Limits (continued)

Metric	Product	Limit
IPv4 ARP entries in hardware with maximum LPM routes —maximum recommended number of IPv4 ARP entries in hardware, with maximum LPM routes present. Assumes number of IP route reserved entries is “maximum.”	Summit X460-G2	43,000 (up to) ^h
	Summit X770, X670-G2	98,000 (up to) ^h
	Summit X450-G2	29,000 (up to) ^h
	ExtremeSwitching X620	1,500
	ExtremeSwitching X440-G2	1,000
IP flow information export (IPFIX) —number of simultaneous flows.	Summit X460-G2	2,048 ingress 2,048 egress
	Summit X450-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	N/A
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. Assumes number of IP route reserved entries is 0, and number of IPv4 ARP entries present is 100 or less.	Summit X460-G2	73,000 ^h
	Summit X770, X670-G2	176,000 (up to) ^h
	Summit X450-G2	61,000 (up to) ^h
	ExtremeSwitching X440-G2, X620	3,500
IPv4 routes —maximum number of IPv4 routes in software (combination of unicast and multicast routes).	Summit X670-G2, X460-G2, X450-G2, X440-G2, X620	25,000
IPv4 routes (LPM entries in hardware) — number of IPv4 routes in hardware.	Summit X460-G2	12,000
	Summit X770, X670-G2, X450-G2	16,000
	ExtremeSwitching X620, X440-G2	480
IPv6 addresses on an interface —maximum number of IPv6 addresses on an interface.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	255
IPv6 addresses on a switch —maximum number of IPv6 addresses on a switch.	Summit X770, X670-G2, X460-G2, X450-G2	2,048
	ExtremeSwitching X620, X440-G2	510
IPv6 host entries in hardware —maximum number of IPv6 neighbor entries in hardware.	Summit X770, X670-G2	36,750 ⁱ
	Summit X460-G2	22,000 ⁱ
	Summit X450-G2	12,000 ⁱ
	ExtremeSwitching X440-G2	1,000
	ExtremeSwitching X620	1,500
IPv6 routes (LPM entries in hardware) —maximum number of IPv6 routes in hardware.	Summit X460-G2	6,000
	Summit X670-G2, X770, X450-G2	8,000
	ExtremeSwitching X620, X440-G2	240

Table 3: Supported Limits (continued)

Metric	Product	Limit
IPv6 routes with a mask greater than 64 bits in hardware —maximum number of such IPv6 LPM routes in hardware.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	256
IPv6 route sharing in hardware —route mask lengths for which ECMP is supported in hardware. Note: * >64 single path only	Summit X460-G2, X670-G2, X770, X450-G2, and ExtremeSwitching X620 ExtremeSwitching X440-G2	0-64 * Not supported
IPv6 routes in software —maximum number of IPv6 routes in software.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	25,000
IP router interfaces —maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.	Summit X460-G2, X770, X670-G2, X450-G2 ExtremeSwitching X620, X440-G2	2,048 510
IP multicast static routes —maximum number of permanent multicast IP routes.	Summit X460-G2, X670-G2, X450-G2, X770	1,024
IP unicast static routes —maximum number of permanent IP unicast routes.	Summit X460-G2, X670-G2, X450-G2, X770 ExtremeSwitching X620, X440-G2	1,024 480
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 is limited to 64 ECMP gateways per destination, while IS-IS is limited to 8. Static routes are limited to 32 next-hops. L2VPNs are limited to 16 LSPs per pseudowire on platforms that support 32 gateways, and 64 LSPs per pseudowire on platforms that support 64 gateways.	Summit X460-G2, X670-G2, X450-G2, X770, and ExtremeSwitching X620 ExtremeSwitching X440-G2	2, 4, 8, 16, or 32 N/A

Table 3: 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.	Summit X670-G2, X770, X450-G2 Summit X460-G2 ExtremeSwitching X620 ExtremeSwitching X440-G2	16,352 12,256 480 N/A
	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.	
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.	Summit X670-G2, X770 default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 if maximum gateways is 64	1,022 1,022 1,022 1,022 510 254
	Summit X460-G2, X450-G2 default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 if maximum gateways is 64	1,022 1,022 510 254
	Summit X460-G2, X450-G2 default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 if maximum gateways is 64	126 62 126 126 126 126 62 30
	ExtremeSwitching X620 default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 if maximum gateways is 64	N/A
	ExtremeSwitching X440-G2	
IP multinetting (secondary IP addresses) —maximum number of secondary IP addresses per VLAN.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	255
IS-IS adjacencies —maximum number of supported IS-IS adjacencies.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X620, X440-G2	128 N/A

Table 3: Supported Limits (continued)

Metric	Product	Limit
IS-IS ECMP—maximum number of equal cost multipath for IS-IS.	Summit X450-G2, X670-G2, X770, X460-G2 ExtremeSwitching X620, X440-G2	2, 4, or 8 N/A
IS-IS interfaces—maximum number of interfaces that can support IS-IS.	Summit X450-G2, X670-G2, X770, X460-G2 ExtremeSwitching X620, X440-G2	255 N/A
IS-IS routers in an area—recommended maximum number of IS-IS routers in an area.	Summit X450-G2, X670-G2, X770, X460-G2 ExtremeSwitching X620, X440-G2	256 N/A
IS-IS route origination—recommended maximum number of routes that can be originated by an IS-IS node.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	20,000 N/A
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.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	25,000 N/A
IS-IS IPv4 L2 routes—recommended maximum number of IS-IS Level 2 routes.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	25,000 N/A
IS-IS IPv4 L1 routes in an L1/L2 router—recommended maximum number of IS-IS Level 1 routes in an L1/L2 IS-IS router.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	20,000 N/A
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.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	10,000 N/A
IS-IS IPv6 L2 routes—recommended maximum number of IS-IS Level 2 routes.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	10,000 N/A
IS-IS IPv6 L1 routes in an L1/L2 router—recommended maximum number of IS-IS Level 1 routes in a L1/L2 router.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	10,000 N/A
IS-IS IPv4/IPv6 L1 routes in an L1 router—recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	20,000 N/A

Table 3: 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.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	20,000 N/A
IS-IS IPv4/IPv6 L1 routes in an L1/L2 router —recommended maximum number of IS-IS Level 1 routes in a Level 1/Level 2 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	20,000 N/A
Jumbo frames —maximum size supported for jumbo frames, including the CRC.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	9,216
L2 VPN: VCCV (pseudowire Virtual Circuit Connectivity Verification) VPNs per switch —maximum number of VCCV enabled VPLS VPNs.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X620, X440-G2	16 N/A
L2 VPN: VPLS MAC addresses —maximum number of MAC addresses learned by a switch.	Summit X770 Summit X670-G2 Summit X460-G2 Summit X450-G2, and ExtremeSwitching X620, X440-G2	128,000 140,000 55,000 N/A
L2 VPN: VPLS VPNs —maximum number of VPLS virtual private networks per switch.	Summit X460-G2, X770, X670-G2 Summit X450-G2, and ExtremeSwitching X620, X440-G2	1,023 N/A
L2 VPN: VPLS peers —maximum number of VPLS peers per VPLS instance.	Summit X770, X670-G2, X460-G2 Summit X450-G2, and ExtremeSwitching X620, X440-G2	64 N/A
L2 VPN: LDP pseudowires —maximum number of pseudowires per switch.	Summit X770 Summit X670-G2 Summit X460-G2 Summit X450-G2, and ExtremeSwitching X620, X440-G2	7,800 7,000 7,116 N/A
L2 VPN: static pseudowires —maximum number of static pseudowires per switch.	Summit X770 Summit X670-G2, X460-G2 Summit X450-G2, and ExtremeSwitching X620, X440-G2	15,308 7,000 N/A
L2 VPN: Virtual Private Wire Service (VPWS) VPNs —maximum number of virtual private networks per switch.	Summit X770 Summit X670-G2 Summit X460-G2 Summit X450-G2, and ExtremeSwitching X620, X440-G2	4,000 4,090 1,023 N/A

Table 3: Supported Limits (continued)

Metric	Product	Limit
<p>Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode.</p> <p>Note::</p> <ul style="list-style-type: none"> The internal lookup table configuration used is "I2-and-I3". 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 same. 	<p>Summit X770, X670-G2 Summit X460-G2 Summit X450-G2 ExtremeSwitching X620, X440-G2</p>	<p>73,000 24,000 14,000 5,000</p>
<p>Layer-3 IPv4 Multicast—maximum number of <S,G,V> entries installed in the hardware (IP multicast compression enabled).</p> <p>Note::</p> <ul style="list-style-type: none"> Limit value same for MVR senders, PIM Snooping entries, PIM SSM cache, IGMP senders, PIM cache. The internal lookup table configuration used is "more I3-and-ipmc". Assumes source-group-vlan mode as look up key. Layer 3 IPMC cache limit in mixed mode also has the same value. 	<p>Summit X460-G2 Summit X450-G2 Summit X770, X670-G2 ExtremeSwitching X620, X440-G2</p>	<p>26,000 21,000 77,500 1,500</p>

Table 3: Supported Limits (continued)

Metric	Product	Limit
Layer-3 IPv6 Multicast —maximum number of <S,G,V> entries installed in the hardware (IP multicast compression enabled). Note: <ul style="list-style-type: none"> 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. 	Summit X770, X670-G2 Summit X460-G2 Summit X450-G2 ExtremeSwitching X620, X440-G2	30,000 14,000 10,000 700
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.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	128
Load sharing —maximum number of ports per load-sharing group.	ExtremeSwitching X620, X440-G2 (standalone and stacked) Summit X770 (standalone) Summit X670-G2 (standalone) Summit X460-G2 (standalone) Summit X450-G2 (standalone) Summit X770 (stacked) Summit X670-G2 (stacked) Summit X460-G2 (stacked) Summit X450-G2 (stacked) Summit X670-G2	8 32 64
Logged messages —maximum number of messages logged locally on the system.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	20,000
MAC-based security —maximum number of MAC-based security policies.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1,024
MAC Locking —Maximum number of MAC locking stations that can be learned on a port.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	64 (static MAC locking stations) 600 (first arrival MAC locking stations)

Table 3: Supported Limits (continued)

Metric	Product	Limit
Meters —maximum number of meters supported.	Summit X460-G2, X450-G2, X670-G2, X770 ExtremeSwitching X440-G2, X620	2,048 N/A
Maximum mirroring instances	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2 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.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	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	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	128
Mirroring, one-to-many (monitor port) —maximum number of one-to-many monitor ports.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	16
MLAG ports —maximum number of MLAG ports allowed.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	768
MLAG peers —maximum number of MLAG peers allowed.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	2
MPLS RSVP-TE interfaces —maximum number of interfaces.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	32 N/A

Table 3: Supported Limits (continued)

Metric	Product	Limit
MPLS RSVP-TE ingress LSPs— maximum number of ingress LSPs.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	2,000 N/A
MPLS RSVP-TE egress LSPs— maximum number of egress LSPs.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	2,000 N/A
MPLS RSVP-TE transit LSPs— maximum number of transit LSPs.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	2,000 N/A
MPLS RSVP-TE paths— maximum number of paths.	Summit X460-G2, X770 Summit X670-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	1,000 2,000 N/A
MPLS RSVP-TE profiles— maximum number of profiles.	Summit X460-G2, X770 Summit X670-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	1,000 2,000 N/A
MPLS RSVP-TE EROs— maximum number of EROs per path.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	64 N/A
MPLS LDP peers—maximum number of MPLS LDP peers per switch.	Summit X770 Summit X670-G2, X460-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	64 128 N/A
MPLS LDP adjacencies— maximum number of MPLS LDP adjacencies per switch.	Summit X460-G2 Summit X770, X670-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	50 64 N/A
MPLS LDP ingress LSPs— maximum number of MPLS LSPs that can originate from a switch.	Summit X770, X670-G2 Summit X460-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	2,048 4,000 N/A
MPLS LDP-enabled interfaces —maximum number of MPLS LDP configured interfaces per switch.	Summit X770 Summit X670-G2, X460-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	64 128 N/A
MPLS LDP Sessions— maximum number of MPLS LDP sessions.	Summit X770 Summit X670-G2, X460-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	64 128 N/A
MPLS LDP transit LSPs— maximum number of MPLS transit LSPs per switch.	Summit X770, X670-G2, X460-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	4,000 N/A
MPLS LDP egress LSPs— maximum number of MPLS egress LSPs that can terminate on a switch.	Summit X770 Summit X670-G2, X460-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	8,000 4,000 N/A
MPLS static egress LSPs— maximum number of static egress LSPs.	Summit X460-G2 Summit X770 Summit X670-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	7,116 8,000 15,308 N/A

Table 3: Supported Limits (continued)

Metric	Product	Limit
MPLS static ingress LSPs— maximum number of static ingress LSPs.	Summit X460-G2	4,000
	Summit X770, X670-G2	2,048
	Summit X450-G2, and ExtremeSwitching X440-G2, X620	N/A
MPLS static transit LSPs— maximum number of static transit LSPs	Summit X770, X670-G2, X460-G2	4,000
	Summit X450-G2, and ExtremeSwitching X440-G2, X620	N/A
MSDP active peers—maximum number of active MSDP peers.	Summit X770, X670-G2, X460-G2	64
	Summit X450-G2, and ExtremeSwitching X440-G2, X620	N/A
MSDP SA cache entries— maximum number of entries in SA cache.	Summit X670-G2, X770	14,000
	Summit X460-G2	10,000
	Summit X450-G2, and ExtremeSwitching X440-G2, X620	N/A
MSDP maximum mesh groups —maximum number of MSDP mesh groups.	Summit X770, X670-G2, X460-G2	16
	Summit X450-G2, and ExtremeSwitching X440-G2, X620	N/A
Multicast listener discovery (MLD) snooping per-VLAN filters—maximum number of VLANs supported in per-VLAN MLD snooping mode.	Summit X460-G2	1,200
	Summit X770, X670-G2	1,200
	Summit X450-G2	512
	ExtremeSwitching X620, X440-G2	600
Multicast listener discovery (MLD)v1 subscribers— maximum number of MLDv1 subscribers per port. ¹	Summit X770, X670-G2, X450-G2, X460-G2	4,000
	ExtremeSwitching X620, X440-G2	3,500
Multicast listener discovery (MLD)v1 subscribers— maximum number of MLDv1 subscribers per switch. ¹	Summit X460-G2, X450-G2	10,000
	Summit X770, X670-G2	30,000
	ExtremeSwitching X620, X440-G2	10,000
Multicast listener discovery (MLD)v2 subscribers— maximum number of MLDv2 subscribers per port. ¹	Summit X450-G2	4,000
	SummitStack	2,000
	Summit X770, X670-G2, X460-G2	4,000
	ExtremeSwitching X620, X440-G2	3,500
Multicast listener discovery (MLD)v2 subscribers— maximum number of MLDv2 subscribers per switch. ¹	SummitStack	5,000
	Summit X460-G2, X450-G2	10,000
	Summit X770, X670-G2	30,000
	ExtremeSwitching X620, X440-G2	10,000
Multicast listener discovery (MLD)v2 maximum source per group—maximum number of source addresses per group.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	200
Multicast listener discovery (MLD) SSM-map entries— maximum number of MLD SSM mapping entries.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X440-G2, X620	500 50

Table 3: Supported Limits (continued)

Metric	Product	Limit
Multicast listener discovery (MLD) SSM-MAP entries —maximum number of sources per group in MLD SSM mapping entries.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	50
Network login —maximum number of clients being authenticated on MAC-based VLAN enabled ports.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1,024
Network login —maximum number of clients being authenticated with policy mode enabled.	Summit X450-G2, X460-G2 Summit X670-G2, X770 ExtremeSwitching X620, X440-G2	1,024 512 256
Network login —maximum number of dynamic VLANs.	Summit X460-G2, X450-G2, X670-G2, X770 ExtremeSwitching X440-G2, X620	2,000 1,024
Network login VLAN VSAs —maximum number of VLANs a client can be authenticated on at any given time.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	10
ONEPolicy Roles/Profiles —maximum number of policy roles/profiles.	Summit X450-G2, X460-G2, X670-G2, X770 ExtremeSwitching X620, X440-G2	63 64
ONEPolicy Rules per Role/Profile —maximum number of rules per role/policy.	Summit X450-G2 Summit X460-G2 Summit X670-G2, X770 ExtremeSwitching X620, X440-G2	IPv6 rules: 256 IPv4 rules: 256 L2 Rules: 184 MAC Rules: 256 IPv6 Rules: 512 IPv4 Rules: 512 L2 Rules: 440 MAC Rules: 512 IPv6 Rules: 256 L2 Rules: 184 MAC Rules: 256 IPv4 Rules: 256 IPv6 and Mac Rules: 0 Ipv4 Rules: 256 (per switch) L2 Rules: 184 (per switch)
ONEPolicy Authenticated Users per Switch —maximum number of authenticated users per switch.	Summit X450-G2, X460-G2 Summit X670-G2, X770 ExtremeSwitching X620, X440-G2	Up to 1,024 Up to 512 Up to 256

Table 3: Supported Limits (continued)

Metric	Product	Limit
ONEPolicy Authenticated Users — maximum authenticated users with a combination of TCI disabled/ enabled profiles.	Summit X450-G2, X460-G2 Summit X670-G2, X770 ExtremeSwitching X620, X440-G2	682-1,022 341-510 TCI disabled: 170 TCI enabled: 256
ONEPolicy Authenticated Users per Port —maximum number of authenticated users per port.	Summit X450-G2, X460-G2 Summit X670-G2, X770 ExtremeSwitching X620, X440-G2	Unlimited up to 1,024 Unlimited up to 512 Unlimited up to 256
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 ExtremeSwitching X620, X440-G2	952 440
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 ExtremeSwitching X620, X440-G2	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 ExtremeSwitching X620, X440-G2	256 N/A
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 ExtremeSwitching X620, X440-G2	256 256
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 ExtremeSwitching X620, X440-G2	184 184
OSPFv2/v3 ECMP —maximum number of equal cost multipath OSPFv2 and OSPFv3.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	16 4

Table 3: Supported Limits (continued)

Metric	Product	Limit
OSPFv2 areas—as an ABR, how many OSPF areas are supported within the same switch.	Summit X460-G2, X670-G2, X770	8
	Summit X450-G2, ExtremeSwitching X440-G2, X620	4
OSPFv2 external routes—recommended maximum number of external routes contained in an OSPF LSDB.	Summit X770, X670-G2, X460-G2, X450-G2	5,000
	ExtremeSwitching X440-G2, X620	2,400
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.	Summit X670-G2, X460-G2, X770	2,000
	Summit X450-G2, and ExtremeSwitching X440-G2, X620	1,000
OSPFv2 interfaces—recommended maximum number of OSPF interfaces on a switch (active interfaces only).	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	4 (with Advanced Edge licence)
	Summit X450-G2, X460-G2, X670-G2, X770	400 (with Core license or higher)
OSPFv2 links—maximum number of links in the router LSA.	Summit X460-G2, X670-G2	400
	Summit X450-G2, and ExtremeSwitching X620, X440-G2	4
	Summit X770	419
OSPFv2 neighbors—maximum number of supported OSPF adjacencies.	Summit X770, X670-G2, X460-G2	128
	Summit X450-G2, and ExtremeSwitching X440-G2, X620	4
OSPFv2 routers in a single area—recommended maximum number of routers in a single OSPF area.	Summit X770, X670-G2, X460-G2	50
	Summit X450-G2, ExtremeSwitching X440-G2, X620	4
OSPFv2 virtual links—maximum number of supported OSPF virtual links.	Summit X460-G2, X670-G2, X770	32
	Summit X450-G2, and ExtremeSwitching X440-G2, X620	4
OSPFv3 areas—as an ABR, the maximum number of supported OSPFv3 areas.	Summit X460-G2, X670-G2, X770	16
	Summit X450-G2, ExtremeSwitching X440-G2, X620	4
OSPFv3 external routes—recommended maximum number of external routes.	Summit X770, X670-G2, X460-G2, X450-G2	10,000
	ExtremeSwitching X440-G2, X620	1,200
OSPFv3 inter- or intra-area routes—recommended maximum number of inter- or intra-area routes.	Summit X770, X670-G2, X460-G2	3,000
	Summit X450-G2, ExtremeSwitching X440-G2, X620	500

Table 3: Supported Limits (continued)

Metric	Product	Limit
OSPFv3 interfaces—maximum number of OSPFv3 interfaces.	Summit X770, X670-G2, X460-G2, X450-G2 ExtremeSwitching X440-G2, X620	4 N/A
	Note:: Active interfaces limit, with Advanced Edge license. (See below for Core license limits.) Summit X770 Summit X670-G2, X460-G2 Summit X450-G2, ExtremeSwitching X440-G2, X620 Note:: With Core license or higher. (See above for Advanced Edge license limits.)	128 256 4
OSPFv3 neighbors—maximum number of OSPFv3 neighbors.	Summit X770, X670-G2, X460-G2 Summit X450-G2, ExtremeSwitching X440-G2, X620	64 4
OSPFv3 virtual links—maximum number of OSPFv3 virtual links supported.	Summit X770, X670-G2, X460-G2 with Core license or higher Summit X450-G2, ExtremeSwitching X440-G2, X620	16 4
PIM IPv4 (maximum interfaces)—maximum number of PIM active interfaces.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620, (Advanced Edge License)	512 4
PIM IPv4 (maximum interfaces)—maximum number of PIM-snooping enabled interfaces.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	512
PIM IPv4 Limits—maximum number of multicast groups per rendezvous point.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	180
PIM IPv4 Limits—maximum number of multicast sources per group.	Summit X460-G2, X670-G2, X770, X450-G2 ExtremeSwitching X440-G2, X620	5,000 1,500
PIM IPv4 Limits—maximum number of dynamic rendezvous points per multicast group.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	145
PIM IPv4 Limits—static rendezvous points.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	32
PIM IPv6 (maximum interfaces)—maximum number of PIM active interfaces.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620 (Advanced Edge License)	512 4
	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	70
PIM IPv6 Limits—maximum number of multicast groups per rendezvous point.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	70
PIM IPv6 Limits—maximum number of multicast sources per group.	Summit X460-G2, X670-G2	2,500
	Summit X450-G2	2,000
	Summit X770	2,500
	ExtremeSwitching X440-G2, X620	550

Table 3: Supported Limits (continued)

Metric	Product	Limit
PIM IPv6 Limits —maximum number of dynamic rendezvous points per multicast group.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	64
PIM IPv6 Limits —maximum number of secondary address per interface.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	70
PIM IPv6 Limits —static rendezvous points.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	32
Policy-based routing (PBR) redundancy —maximum number of flow-redirects.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	256°
Policy-based routing (PBR) redundancy —maximum number of next hops per each flow-direct.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	32°
Port-specific VLAN tags —maximum number of port-specific VLAN tags.	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	1,023 N/A
Port-specific VLAN tags —maximum number of port-specific VLAN tag ports.	Summit X770, X670-G2 Summit X460-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	6,400 4,000 N/A
Private VLANs —maximum number of subscribers. Assumes a minimum of one port per network and subscriber VLAN.	Summit X770 Summit X670-G2 Summit X460-G2 Summit X450-G2 ExtremeSwitching X440-G2 ExtremeSwitching X620	103 63 53 51 47 15
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 ExtremeSwitching X440-G2 ExtremeSwitching X620	1,024 255 510
Private VLANs —maximum number of private VLANs in an L2-only environment.	Summit X770, X670-G2, X460-G2, X450-G2 ExtremeSwitching X440-G2, X620	1,280 255
PTP/1588v2 Clock Ports	Summit X770, X460-G2, X670-G2 ExtremeSwitching X440-G2, X620	32 for boundary clock 1 for ordinary clock N/A

Table 3: Supported Limits (continued)

Metric	Product	Limit
PTP/1588v2 Clock Instances	Summit X770, X670-G2, X460-G2 ExtremeSwitching X440-G2, X620	2 combinations: <ul style="list-style-type: none"> Transparent clock + ordinary clock Transparent clock + boundary clock N/A
PTP/1588v2 Unicast Static Slaves	Summit X770, X670-G2, X460-G2 ExtremeSwitching X440-G2, X620	40 entries per clock port N/A
PTP/1588v2 Unicast Static Masters	Summit X770, X670-G2, X460-G2 ExtremeSwitching X440-G2, X620	10 entries per clock type N/A
Route policies—suggested maximum number of lines in a route policy file.	Summit X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	10,000
RIP Learned Routes—maximum number of RIP routes supported without aggregation.	Summit X770, X670-G2, X460-G2, and ExtremeSwitching X440-G2, X620	10,000
RIP interfaces on a single router—recommended maximum number of RIP routed interfaces on a switch.	Summit X670-G2, X460-G2 Summit X770, X450-G2 ExtremeSwitching X440-G2, X620	256 256 128
RIPng learned routes—maximum number of RIPng routes.	Summit X670-G2, X460-G2, X770, X450-G2 ExtremeSwitching X440-G2, X620	3,000 N/A
Spanning Tree (maximum STPDs)—maximum number of Spanning Tree Domains on port mode EMISTP.	Summit X450-G2, X770, X670-G2, X460-G2, and ExtremeSwitching X620 ExtremeSwitching X440-G2	64 32
Spanning Tree PVST+—maximum number of port mode PVST domains.	Summit X770, X670-G2, and ExtremeSwitching X620 Summit X460-G2, X450-G2, and ExtremeSwitching X440-G2	256 128
Note:: <ul style="list-style-type: none"> Maximum of 10 active ports per PVST domain when 256 PVST domains are configured. Maximum of 7 active ports per PVST domain when 128 PVST domains are configured. 		

Table 3: Supported Limits (continued)

Metric	Product	Limit
Spanning Tree —maximum number of multiple spanning tree instances (MSTI) domains.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620 ExtremeSwitching X440-G2	64 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.	Summit X770, X670-G2 Summit X460-G2, X450-G2 ExtremeSwitching X440-G2 ExtremeSwitching X620	500 600 256 600
Spanning Tree —maximum number of VLANs on all MSTP instances.	Summit X770 Summit X670-G2 Summit X460-G2, X450-G2 ExtremeSwitching X440-G2 ExtremeSwitching X620	1,024 1,000 1,024 512 1,024
Spanning Tree (802.1d domains) —maximum number of 802.1d domains per port.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1
Spanning Tree (number of ports) —maximum number of ports including all Spanning Tree domains.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620 Summit X440-G2	4,096 2,048
Spanning Tree (maximum VLANs) —maximum number of STP-protected VLANs (dot1d and dot1w).	Summit X770, and ExtremeSwitching X620 Summit X670-G2 Summit X460-G2, X450-G2 ExtremeSwitching X440-G2	1,024 560 600 500
SSH (number of sessions) —maximum number of simultaneous SSH sessions.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8
Static MAC multicast FDB entries —maximum number of permanent multicast MAC entries configured into the FDB.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1,024
Syslog servers —maximum number of simultaneous syslog servers that are supported.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	4
Telnet (number of sessions) —maximum number of simultaneous Telnet sessions.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8

Table 3: Supported Limits (continued)

Metric	Product	Limit
<p>Virtual routers—maximum number of user-created virtual routers that can be created on a switch.</p> <p>Note: Virtual routers are not supported on Summit X440 series switches.</p>	Summit X460-G2, X670-G2, X770, X450-G2 ExtremeSwitching X440-G2, X620	63 N/A
<p>Virtual router forwarding (VRFs)—maximum number of VRFs that can be created on a switch.</p> <p>Note: * Subject to other system limitations.</p>	Summit X460-G2, X670-G2, X770, X450-G2 ExtremeSwitching X440-G2, X620	960 * N/A
<p>Virtual router protocols per VR—maximum number of routing protocols per VR.</p>	Summit X460-G2, X670-G2, X770, X450-G2 ExtremeSwitching X440-G2, X620	8 N/A
<p>Virtual router protocols per switch—maximum number of VR protocols per switch.</p>	Summit X460-G2, X670-G2, X770, X450-G2 ExtremeSwitching X440-G2, X620	64 N/A
<p>VLAN aggregation—maximum number of port-VLAN combinations on any one superVLAN and all of its subVLANs.</p>	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	1,000
<p>VLANs—includes all VLANs.</p> <p>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.)</p>	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	4,094
<p>VLANs—maximum number of port-specific tag VLANs.</p>	Summit X770, X670-G2, X460-G2 ExtremeSwitching X440-G2, X620	4,093 N/A
<p>VLANs—maximum number of port-specific tag VLAN ports.</p>	Summit X460-G2 Summit X770, X670-G2 Summit X450-G2, and ExtremeSwitching X440-G2, X620	4,096 8,192 N/A
<p>VLANs (Layer 2)—maximum number of Layer 2 VLANs.</p>	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	4,094
<p>VLANs (Layer 3)—maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.</p>	Summit X460-G2, X770, X670-G2, X450-G2 ExtremeSwitching X440-G2, X620	2,048 510

Table 3: Supported Limits (continued)

Metric	Product	Limit
VLANs (maximum active port-based) —maximum active ports per VLAN when 4,094 VLANs are configured with default license.	Summit X770, X670-G2, X460-G2, X450-G2, and ExtremeSwitching X440G2 ExtremeSwitching X620	32 16
VLANs (maximum active protocol-sensitive filters) —number of simultaneously active protocol filters in the switch.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	16
VLAN translation —maximum number of translation VLANs. Assumes a minimum of one port per translation and member VLAN.	Summit X770 Summit X670-G2 Summit X460-G2 Summit X450-G2 ExtremeSwitching X620 ExtremeSwitching X440-G2	103 63 53 51 15 47
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 ExtremeSwitching X620 ExtremeSwitching X440-G2	1,024 512 255
VLAN translation —maximum number of translation VLAN pairs in an L2-only environment.	Summit X460-G2 Summit X450-G2, X770, X670-G2 ExtremeSwitching X440-G2, X620	2,046 1,024 512
VRRP (v2/v3-IPv4) (maximum instances) —maximum number of VRRP instances for a single switch, with Advanced Edge license or higher. Note: These limits are applicable for Fabric Routing configuration also.	Summit X770, X670-G2, X460-G2, X450-G2 ExtremeSwitching X440-G2, X620	511 128

Table 3: Supported Limits (continued)

Metric	Product	Limit
VRRP (v3-IPv6) (maximum instances) —maximum number of VRRP instances for a single switch, with Advanced Edge license or higher. (VRRP-VRRPv3-IPv6) Note: These limits are applicable for Fabric Routing configuration also.	Summit X770, X670-G2, X460-G2, X450-G2 ExtremeSwitching X440-G2, X620	511 128
VRRP (v2/v3-IPv4/IPv6) (maximum VRID) —maximum number of unique VRID numbers per switch. Note: With Advanced Edge license or higher	Summit X770, X670-G2, X460-G2, X450-G2 and ExtremeSwitching X440-G2, X620	31
VRRP (v2/v3-IPv4/IPv6) (maximum VRIDs per VLAN) —maximum number of VRIDs per VLAN. Note: With Advanced Edge license or higher	Summit X770, X670-G2, X460-G2, X450-G2 and ExtremeSwitching X440-G2, X620	31
VRRP (v2/v3-IPv4/IPv6) (maximum ping tracks) —maximum number of ping tracks per VLAN. Note: With Advanced Edge license or higher	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8
VRRP (maximum ping tracks) —maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8 (20 centisecond or 1 second hello interval)
VRRP (v3-IPv6) (maximum ping tracks) —maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8 (20 centisecond or 1 second hello interval)
VRRP (v2/v3-IPv4/IPv6) (maximum iproute tracks) —maximum number of IP route tracks per VLAN.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8
VRRP (v2/v3-IPv4/IPv6) —maximum number of VLAN tracks per VLAN.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8

Table 3: Supported Limits (continued)

Metric	Product	Limit
<p>VXLAN—maximum virtual networks.</p> <p>Note: Every VPLS instance/ PStag VLAN reduces this limit by 1. Assumption is all BUM (broadcast/unknown-unicast/multicast) FDB entries are pointing to the same set of RTEPs when all VNETs use explicit flooding. Depends on whether all VNETs use standard or explicit and the number of tenant VLAN ports.</p>	Summit X670-G2, X770 Summit X460-G2, X450-G2, and ExtremeSwitching X440-G2, X620	2,048–4,000 N/A
<p>VXLAN—maximum tenant VLAN plus port combinations</p> <p>Note: Every (VPLS/PStag VLAN/ TRILL access VLAN) + port reduces the limit by 1.</p>	Summit X670-G2, X770 Summit X460-G2, X450-G2, and ExtremeSwitching X440-G2, X620	4,096 N/A
<p>VXLAN—maximum static MAC to IP bindings.</p> <p>Note: Every FDB entry configured reduces this limit by 1</p>	Summit X670-G2, X770 Summit X460-G2, X450-G2, and ExtremeSwitching X440-G2, X620	64,000 N/A
<p>VXLAN—maximum RTEP IP addresses</p>	Summit X670-G2, X770 Summit X460-G2, X450-G2, and ExtremeSwitching X440-G2, X620	512 N/A
<p>VXLAN—maximum virtual networks with dynamic learning and OSPF extensions for VXLAN</p>	Summit X670-G2, X770 Summit X460-G2, X450-G2, and ExtremeSwitching X440-G2, X620	4,000 N/A
<p>XML requests—maximum number of XML requests per second.</p> <p>Note: Limits are dependent on load and type of XML request. These values are dynamic ACL data requests.</p>	Summit X450-G2, and ExtremeSwitching X440G2, X620	10 with 100 DACLs
<p>XNV authentication—maximum number of VMs that can be processed (combination of local and network VMs).</p>	Summit X460-G2, X670-G2, X770 Summit X450-G2, and ExtremeSwitching X440-G2, X620	2,048 1,024

Table 3: Supported Limits (continued)

Metric	Product	Limit
XNV database entries —maximum number of VM database entries (combination of local and network VMs).	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	16,000
XNV database entries —maximum number of VPP database entries (combination of local and network VPPs).	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	2,048
XNV dynamic VLAN —Maximum number of dynamic VLANs created (from VPPs / local VMs).	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	2,048
XNV local VPPs —maximum number of XNV local VPPs.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	2,048 ingress 512 egress
XNV policies/dynamic ACLs —maximum number of policies/dynamic ACLs that can be configured per VPP.	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	8 ingress 4 egress
XNV network VPPs —maximum number of XNV network VPPs. ^p	Summit X450-G2, X460-G2, X670-G2, X770, and ExtremeSwitching X620, X440-G2	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

Open Issues
Known Behaviors
Resolved Issues in ExtremeXOS 21.1.1-Patch1-5
Resolved Issues in ExtremeXOS 21.1.1-Patch1-2
Resolved Issues in ExtremeXOS 21.1

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 21.1.1-Patch1-5.

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

CR Number	Description
General	
xos0062584	The command <code>configure web http access-profile</code> is missing the option <code>access-profile</code> .
xos0062966	When rendezvous point receives (*, G) join, and it has (S, G, RPT) entry, the entry should be converted to (S, G), and (S, G) join should be sent upstream. However, sometimes wrong assert is triggered and this new entry is dropped after a minute, resulting in complete traffic loss. Workaround: Disable, and then re-enable PIM.
xos0063346	With multiple (greater than two) VRRP instances and host-mobility enabled, FDB flush sent during topology change from other L2 protocols does not occur.
xos0063247	Duplicate packets occur during MLAG link recovery when LACP is used for load sharing.
xos0063396	While enabling 1G links, unexpected link-down PDU is received in EAPS master.
xos0048715	IPv6 ECMP works for hardware-forwarded traffic, but does not work for slow-path traffic. Workaround: Either use BFD to keep all router neighbors alive, or configure static neighbors and static FDB entries for all router neighbors. BFD is the preferred method.
xos0062399	IPv6 BFD session for OSPFv3 flaps after disabling, and then enabling VLAN port.
SummitStack	

**Table 4: Open Issues, Platform-Specific, and Feature Change Requests (CRs)
(continued)**

CR Number	Description
xos0062386	With BGPv6, after port flap or failovers, some peers go into idle state.
xos0061909	<p>Creating an IPFIX mirroring instance to a monitor port, deleting the mirroring instance, and then recreating it again to a different monitor port, causes the following error message (similar to the one below) to appear, and IPFIX mirroring does not work:</p> <pre><Error:HAL.Mirror.Error> Slot-1: Failed to create mirroring destination for slot 2, unit 9 Entry exists</pre> <p>Workaround: If the error appears in the log, disable and delete the mirror instance, and then add it back again.</p>
ExtremeSwitching X620 Series Switches	
xos0062729	<p>On ExtremeSwitching X620 series switches, for ports with Base-T SFP optics and explicitly configured at 1,000 speed, link comes up at peer end, but link stays down at local end after either rebooting, or removing, and then re-inserting optics.</p> <p>Workaround: Manually set autoneg to "off" on the host side, which allows the link to be established.</p>
xos0062636	<p>Unexpected link switchover behavior occurs when exchanging copper and fiber cables on ExtremeSwitching X620 combo ports.</p> <p>Workaround:</p> <ul style="list-style-type: none"> When 10G combo ports are used at 1G for redundancy between fiber and copper, then set the preferred medium to copper (<code>configure ports port_list preferred-medium copper</code>), otherwise sometimes the copper link might not come up. When 10G combo ports are used at 10G for redundancy between fiber and copper, then set the preferred medium to fiber (<code>configure ports port_list preferred-medium fiber</code>), otherwise sometimes the copper link might not come up.
xos0062620	For ExtremeXOS 21.1, do not use copper DAC cables for stacking on ExtremeSwitching X620-16T switches.
Summit X670-G2 Series Switches	
xos0063170	On Summit X670-G2 series switches, greater EAPS convergence time occurs with multiple VLANs (1,000 protected VLANs).
xos0063492	When a 1G port (SX/BASET) is used as a loopback port for mirroring to a port-list, the port does not come back to active state after disabling mirroring.
Summit X460-G2 Series Switches	

**Table 4: Open Issues, Platform-Specific, and Feature Change Requests (CRs)
(continued)**

CR Number	Description
xos0063811	<p>Summit X460-G2 series switches with ExtremeXOS 15.6 through 21.1, have the following limitations for SyncE input reference frequency:</p> <ul style="list-style-type: none"> 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). <p>Workaround: For SyncE input at 10G, avoid port 52.</p> <ul style="list-style-type: none"> 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. <p>Workaround: For SyncE input at 1G, use a 1G port, not a 10G port.</p>
xos0063412	CFM fault not detected locally after disabling CCM for hardware Down MEP.
ExtremeSwitching X440-G2 Series Switches	
xos0062236	EEE becomes disabled on combo ports after peer ports are disabled, and then enabled.
xos0062583	<p>Policy: Dynamic VLAN is not removed from backup slot after issuing <code>unconfigure policy mactable</code>.</p> <p>Workaround: Reboot the switch.</p>
xos0062773	After switch boot up or restart of process dot1ag, 95% CPU utilization occurs with 32 UP MEPs (maximum value).
xos0062895	<p>On ExtremeSwitching X440-G2 stack, process nettools fails to start after rebooting with DHCPv6 client enabled. Switch reboots repeatedly and reports the following errors:</p> <pre> 10/20/2015 11:31:06.45 <Erro:EPM.crash_rate> Slot-1: Process netTools exceeded pre-configured or default crash rate 10/20/2015 11:31:06.45 <Erro:DM.Error> Slot-1: Process netTools Failed 10/20/2015 11:31:06.45 <Erro:DM.Error> Slot-1: Process netTools Failed 10/20/2015 11:31:06.45 <Erro:DM.Error> Slot-3: Node State[4] = FAIL (Not In Sync) 10/20/2015 11:31:06.46 <Erro:DM.Error> Slot-1: Node State[4] = FAIL (Process Failure) 10/20/2015 11:31:06.61 <Crit:NM.NodeStateFail> Slot-1: Slot-3 has failed for the reason of "Not In Sync". </pre>
xos0062899	<p>DHCPv6 client remains in initializing state after disabling, and then enabling, the port in the relay switch.</p> <p>Workaround: Restart process nettools, or disable, and then enable, DHCP client.</p>

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

CR Number	Description
xos0063678	In ExtremeSwitching X440-G2 stack, rebooting backup slot with CFM 32 Down MEP configuration times out with the following errors: <pre>Error: This command is not permitted on nodes that are not active 02/16/2016 15:16:34.06 <Warn:HAL.Stacking.Warning> Slot-1: Timed out waiting for 1 reboot replies. 02/16/2016 15:16:34.06 <Warn:HAL.Stacking.Warning> Slot-1: Timed out waiting for 1 reboot replies.</pre>
xos0064964	Misleading fan failure may be reported on Summit X440-G2-12t and X440-G2-12p switches due to internal fans running at 0 RPM. Fans in these models may operate at 0 RPM if the system temperature is maintained without running the fans. This is normal operation.
Summit X450-G2 Series Switches	
xos0063008	In Summit X450-G2 stack with mirroring configuration, boot up times out (after 300 seconds) while waiting for configuration checkpoint save operation to finish (FDB is still not saved).
BGP	
xos0060641	When BGP is administratively shut down, it does not send notifications to peers.
xos0063778	If an applied BGP import policy is edited such that previously permitted routes are now denied, the BGP RIB (<code>show bgp routes</code> command) still shows the newly denied route(s) as active. The routing table is, however, updated correctly to reflect the new policy.
xos0063698	A BGP route is not replaced in the routing table by a new instance of the same prefix and length containing a different metric value. This condition can occur if an applied BGP import-policy file is edited to modify the route metric.
MPLS	
xos0062996	VPLS: Traffic is not forwarded to service port (VLAN-tagged port) when CEP egress filtering is enabled on it. Workaround: Disable CEP egress filter.
NetLogin	
xos0062680	Switch fails to send Radius accounting message for dot1x user after <code>clear netlogin state port <portNumber></code> command.
Optics	
xos0062092	For Finisar LX-SFP optics, RxPower appears as "inf" instead of displaying correct value in the output of the <code>show port transceiver information detail</code> command.
SNMP	
xos0062492	Traps having tabular variables as varbinds should include the instance along with the tabular OID.
xos0062523	SNMP traps for overheat and negative temperatures incorrectly report detected problems.

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

CR Number	Description
xos0062525	extremeEdpNeighborAdded/extremeEdpRemoved traps varbinds need to include the instance along with the OID.
xos0062527	The varbinds of extremePowerSupplyGood, extremePsuPowerStatus traps need to include the instance along with the OID.

Known Behaviors

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

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

CR Number	Description
General	
xos0062115	For ExtremeSwitching X440-G2 and X620 series switches, Dot1p is not set properly in the CVID translated packet.
xos0062068	For extremeErpsRingNodeType, node type of RPL Neighbor returns value "nonRplOwner" for SNMP get.
xos0062131	In VMANs, CEP port does not transmit ELRP packets, unlike untagged/tagged ports.
xos0062466	For ExtremeSwitching X440-G2 and X620 series switches, VLAN traffic is not dropped if the port is classified both as VLAN tagged and VMAN CEP port. However, Dot1p value is set to "0" in egress for all priorities in VLAN traffic.
xos0062119, xos0063047	For ExtremeSwitching X440-G2 and X620 series switches, "qosmonitor congestion counter" for egress port does not appear in the output of the show ports qosmonitor congestion command when a port list is specified, instead of a single egress port.
xos0063413	In Chalet, when switching between earlier versions of ExtremeXOS to version ExtremeXOS 21.1, the Apps tab does not appear. Workaround: Reload the web page or clear the cache.
SSH	
xos0063327	If a switch is downgraded from ExtremeXOS 21.1 to previous releases, with RSA key saved, the key becomes invalid.
VXLAN	
xos0060213	Same port cannot be a part of network as well as tenant VLANs.
xos0063148	Rate-limit actions do not work when the port is added as VXLAN tenant on VLAN ports.
xos0059594	Egress mirroring of VXLAN traffic is not supported.

Table 5: Known Issues, Platform-Specific, and Feature Change Requests (CRs) (continued)

CR Number	Description
xos0059464	With no network ports configured and the switch receives VXLAN traffic from the access VLAN side, traffic is sent to the CPU, causing high BCMRx usage (around 50%), which in turn affects other parts of the system, such as OSPF (neighbor flap), pings etc. The frames are going to the CPU because they have the MAC DA and Destination IP address of the local switch. This behavior is no different than if the switch were a non-VXLAN-capable switch. By default all ports can terminate VXLAN traffic. If network ports are deleted with <code>configure virtual-network delete network ports portlist</code> any VXLAN traffic on these ports is sent to the CPU.
xos0062919	With VXLAN configuration, after rebooting the following error appears: <Error:HAL.Ipv4Mc.GrpTblFull> IPv4 multicast entry not added. Hardware Group Table full.
Summits and ExtremeSwitching Series Switches	
xos0063046	On ExtremeSwitching X440-G2 and Summit X460-G2 series switches, for the 1G combo ports if fiber is the preferred medium and a copper cable is inserted, and then a fiber cable is also inserted, the link switches from copper mode to fiber mode, and a link flap occurs.
Summit X460-G2 Series Switches	
xos0062225	For Summit X4460-G2 switches, when HwBFD session is enabled, configuring authentication is ineffective (session stays up despite a password mismatch between the neighbors).
Summit X670-G2 Series Switches	
xos0062486	For Summit X670-G2 series switches, configuring overhead bytes using <code>configure forwarding rate-limit overhead-bytes</code> does not work with egress ACL meter.
SummitStack	
xos0062687	For Summit X450-G2 and X620 SummitStacks, after stack reboots, the following error message appears: <Warn:DM.Warning> Slot-2: mcmgr cannot write msg_id 5 to MASTER connection 0 This error can be ignored. No functional problem has occurred.

Resolved Issues in ExtremeXOS 21.1.1-Patch1-5

The following issues were resolved in ExtremeXOS 21.1.1-Patch1-5. ExtremeXOS 21.1.1-Patch1-5 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-10, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.5.2, ExtremeXOS 15.6.5, ExtremeXOS 15.7.3, ExtremeXOS 16.1.3 and ExtremeXOS 21.1.1. For information about those fixes, see the release notes for the specific release.

Table 6: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 21.1.1-Patch1-5

CR Number	Description
General	
xos0064117	ExtremeSwitching X440-G2, and Summit X450-G2, X460-G2, X620, and X670-G2 series switches have a hardware component change in newly manufactured switches that requires an ExtremeXOS image update for support. For a detailed list of part number revisions that require the ExtremeXOS upgrade, see https://www.extremenetworks.com/extreme-hardwaresoftware-compatibility-recommendation-matrices/summit-extremeswitching-components-extremexos-software-support/ .
xos0055541	TACACS: On configuring shared secret key in encrypted form with characters "&" and "<", the <code>show configuration aaa</code> output shows a different secret key from what was actually configured.
xos0057931	After rebooting the switch multiple times, 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).
xos0062265	Some legacy commands are not recognized.
xos0062444	Kernel panic occurs in DoS protect-enabled switches when TCN SYN packets to port 80 are flooded to Management port.
xos0063331	VLAN IP address is unconfigured when modifying the VLAN name/port information from Chalet.
xos0063332	Configuration changes to VPLS are not fully retrieved by SNMP walk, which returns values for only few VPLS index.
xos0063842	Packets are being flooded in both network and access VLAN ports after port flap.
xos0063995	SNMP sysUpTime does not return correct value after failover
xos0064009	MLAG+EAPS:Traffic forwarding stops after EAPS that include ISC link converges.
xos0064054	SNMPwalk on extremeAcIStatsTable returns value with port instance instead of ifIndex.
xos0064055	Resiliency Enhancement for IPv4 and IPv6 Static Routes feature has been added.
xos0064063	Packet-Resolution match conditions need to be added as ACL match conditions.
xos0064075	The output of the <code>show fan</code> command shows fan status as "Failed" after hot re-seating a fan module.
xos0064129	Policy refresh never completes with network-zone configuration.
xos0064151	Error occurs when removing DHCP configuration from VLANs when LAG ports are added to the VLANs.
xos0064178	Hardware entries not released on disabling of ports in a LAG connecting an L2VPN router to the MPLS core when L2VPN sharing is configured and in use.

Table 6: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 21.1.1-Patch1-5 (continued)

CR Number	Description
xos0064179	MAC movement occurs in switch acting as an STP root bridge when PVST+ BPDUs are sent by peer switch using STP blocked port.
xos0064203	Incorrect next hop is chosen by BGP route after port flap.
xos0064215	The following log message appears when a subnet is reachable both using MPLS and non-MPLS: <Warn:Kern.Ipv4FIB.Warning> Slot-4: dest 0x0A420000 / 24 nexthop 0xAC11121E: Unable to add route to unit 1, rc Entry not found. Shadow problem.
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.
xos0064278	In a SummitStack or BlackDiamond chassis, FDB is not programmed in hardware after three failovers and fallback.
xos0064281	In Chalet, switch inappropriately displays logs for user accounts under enhanced security mode.
xos0064299	The hal process ends unexpectedly after executing the command debug packet capture on .
xos0064307	RADIUS accounting configuration is incorrect as shown by the command show conf aaa and is lost after upgrade.
xos0064319	Aggregated BGP route is not transmitted to upstream neighbor when highest prefix route is received from neighbor.
xos0064326	LACP flaps when the LAG port is added to vman, with the vman ether type same as LACP ether type
xos0064357	Out of sync between PIM and RTMgr process after introducing new best route.
xos0064383	In the show l2vpn detail command output, the "PW Tx Pkts" counters are not updated for VPWS sessions even though traffic is passing correctly.
xos0064395	STP digest value gets changed when adding the port in VLAN or removing the port from VLAN.
xos0064490	After upgrading from ExtremeXOS 15.2 to later release, last installed dynamic ACL rule is given more priority than previously installed rules.
xos0064519	With MVR enabled on two VLANs, IGMP report packets are looped if sent to all hosts group.
xos0064589	While learning BGP routes, some routes are not getting installed in route table when deleting and re-adding the static route.
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.
xos0062720	Unable to save configuration when ACL/CFM is configured on multiple VLANs.

Table 6: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 21.1.1-Patch1-5 (continued)

CR Number	Description
xos0063429	The output of the command <code>show fan</code> shows the fan status as empty after a hot re-seating of the fan module.
xos0064400	When switch boots up with factory default configuration, Zero Touch Provisioning (ZTP+) is enabled automatically and IP is resolved correctly using DHCP, but this causes flooding to be disabled on all ports.
xos0063693	With L2VPN sharing, traffic loss occurs after LSP failover.
xos0061018	After failover, traffic fails across VPLS configured with 64 LSPs across LAG.
xos0064312	With VXLAN, if tenant VLAN and tunnel are on different VRs, FDB is not learned on a tunnel.
xos0063844	With VXLAN, MLAG port in backup slot is not added to aggregator after reboot of switch followed by disable and enable of port.
xos0064136	Unable to configure flood rate limit as 1 packets per second.
Summit X670 Series Switches	
xos0057671	Link status goes to Ready state on port with 10/100/1000BASE-T optics after multiple reboots.
xos0063263	On Summit X670 series switches, 1000BaseSX optics are incorrectly detected as 100BaseFX.
SummitStack	
xos0061834	In SummitStacks, the command <code>synchronize stacking slot <slot no></code> does not work from master node if the target slot is in failed state.
xos0061861	A per-port meter configured on a SummitStack may not be properly configured on the backup node following a reboot.
xos0062484	EPM process crashes on master if image upgrade on a standby slot exceeds 30 minutes
Summit X670-G2 Series Switches	
xos0064574	In X670G2, IPMC cache entries are limited to 5000, when the lookup key is changed from Source-Group-Vlan to Group-Vlan mode or vice versa.
BlackDiamond X8 Series Switches	
xos0064010	The command <code>show port buffer</code> displays an incorrect port range for 100G I/O modules.
xos0060666	After failover, traffic gets flooded on the ports of service VLAN in H-VPLS core.
Summit X440 Series Switches	
xos0063627	ARP is not re-added to hardware after it is removed initially due to the table being full.

Resolved Issues in ExtremeXOS 21.1.1-Patch1-2

The following issues were resolved in ExtremeXOS 21.1.1-Patch1-2. ExtremeXOS 21.1.1-Patch1-2 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, and ExtremeXOS 16.1.3. For information about those fixes, see the release notes for the specific release.

Table 7: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 21.1.1-Patch1-2

CR Number	Description
General	
xos0061745	Ampersand used in UPM script is replaced by “& amp” in the XSF configuration.
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.
xos0063082	Updated DSCP value is not refreshed for Dynamic ACLs.
xos0063186	Kernel oops occurs when deleting private VLAN.
xos0063282	ExtremeXOS CLI restricts PVLAN subscriber VLAN from being configured as an EAPS-protected VLAN.
xos0063423	Memory leak occurs in ISIS process when exporting OSPF routes to ISIS.
xos0063465	Cannot add/delete ports to load-shared MLAG ports without disabling MLAG.
xos0063484	Enhancement added in STP flush generation mechanism to reduce hardware programming load.
xos0063495	Policy authentication fails when RADIUS request queue has stale entries.
xos0063614	Kernel crash occurs when receiving DHCP packets with invalid field values.
xos0063710	Kernel oops occurs on switch with Private VLAN and MLAG configuration after executing <code>restart ports all</code> .
xos0063736	In Syslog, username information appears as “*****” during login/logout cases.
xos0063761	Traffic is not forwarded after disable/enable sharing when traffic ingressing port is part of both port specific tag (PSTag) and service VMAN (untagged port).
xos0063791	Disabling MLAG ports on both MLAG peer switches prior to VM migration prevents proper VXLAN termination.
xos0063853	Potential memory corruption when MAC locking is configured.
xos0063870	Kernel oops occurs due to memory overrun in user kernel interface.
xos0063956	ACL slice is not freed up after changing IGMP snooping filter from per-VLAN to per-port mode.

Table 7: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in ExtremeXOS 21.1.1-Patch1-2 (continued)

CR Number	Description
xos0063967	XNV dynamic VLAN is not created on MLAG peer where MLAG port is down.
xos0063968	HAL process ends unexpectedly after changing/reverting service VLAN tag.
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.
xos0063948	Clearflow delta values are randomly not calculated properly.
xos0063463	Static FDB created on PStag VLAN port is incorrectly displayed in <code>show configuration</code> command.
xos0063494	OSPFv3 process ends unexpectedly on BFD-enabled switches, if there are frequent link flaps for a long duration.
xos0063814	UPM process ends unexpectedly with Signal 11 occasionally when UPM timers are configured.
xos0063849	VXLAN: The commands <code>disable ospf</code> and <code>disable OSPF vxlan-extensions</code> does not flush learned RTEPs.
xos0064045	Need support for tagged and untagged VXLAN tenant VLANs on the same port.
xos0064122	The command <code>show tech-support</code> terminates prematurely when 40G or 100G optics are present in the switch.
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.
xos0063872	After multiple executions of <code>run failover</code> with redirect-flow configuration, IPv4 ping fails.
xos0063928	Sysuptime in sFlow packets is invalid.
Summit X460-G2 Series Switches	
xos0063927	Error "Deferred L2 notification code out of sync unit 0" repeatedly appears in log.
Summit X670-G2 Series Switches	
xos0063807	On Summit X670-G2 series switches, egress ACL rule actions do not take effect on ports 64-72.
SummitStack	
xos0061777	Standby nodes do not come back up to operational state after they go into failed state.
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.
xos0063904	FDB process ends unexpectedly in backup node of SummitStack configured as MLAG peer when certain FDB entries are not flushed properly after age-out.

Resolved Issues in ExtremeXOS 21.1

The following issues were resolved in ExtremeXOS 21.1. ExtremeXOS 21.1 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, 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)

CR Number	Description
General	
xos0050771	The command <code>show access-list dynamic counters</code> does not display the complete MAC address of VMs and it may not be possible to read the counters correctly from the output.
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.
xos0056829	Switches do not re-send the Group Specific Query following the <code>last_member_query_interval</code> (1 second).
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.
xos0057269	SNMP trap <code>extremelpSecurityViolation</code> is sent with incorrect VLAN description.
xos0057374	Switch odometer value is reinitialized when Master Switch Fabric Module (MSM) fails to read the value.
xos0057672	The process <code>rtmgr</code> ends unexpectedly when disabling GRE tunnels.
xos0058669	DHCPv6 client: After changing the client identifier type, and then restarting the port, old IPv6 addresses are not released, causing the <code>show vlan</code> command to show multiple IPv6 addresses.
xos0058750	Neighbor discovery packets are duplicated in L2 VLANs when IPv6 addresses are configured for other VLANs that do not have any ports.
xos0059942	SSH connection ends when show commands produce lengthy output.
xos0060092	Fetching values using SNMP for "extremePortQosStatsTable" does not work correctly.
xos0060643	Commands for downloading and installing images should use active/inactive options when specifying partitions (in addition to current primary/secondary options).
xos0061085	Kernel oops occurs while deleting VR with enable BGP export and IPARP proxy configurations.
xos0061173	L2PT packets are dropped when ingress port is configured with software learning.
xos0061198	Disabling VPN-VRF affects traffic on another VPN-VRF.
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.

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

CR Number	Description
xos0061247	Configuring IPv6 Syslog target in a specific format produces an incomplete command error, even though the command is complete.
xos0061331	Bootprelay for VRF is not supported. Commands to configure bootprelay should reflect this.
xos0061445	After creating and enabling an STPD, the command <code>configure "Default" add ports 1 tagged stpd "s1"</code> adds ports to the Default VLAN, but not with STPD domain, even though the error "Command Aborted and no changes were made" appears.
xos0061465	IPv6 source address that is not configured on any VLAN in the given VR is accepted as from source IP. Issue does not occur with IPv4.
xos0061507	SNMPget on EXTREME-SOFTWARE-MONITOR table returns value with incorrect OID.
xos0061517	LACP adjacency fails while forwarding the PDU with l2pt profile over L2VPN tunnels when MPLS PHP is enabled.
xos0061565	The TCL function, "clock scan," generates errors with default time zone configuration.
xos0061656	Nodes remain in the "FDBSync" state due to temp-flooding while rebooting the stack.
xos0061788	The process devmgr ends unexpectedly during snmpwalk when continuous EMS logs are sent to the switch console.
xos0062017	DHCP trusted port configuration is lost after disabling, and then re-enabling LAG.
xos0062018	For IPv6 routes with mask lengths greater than 64-bits, IPv6 unicast packets destined for the switch CPU can be dropped if another IPv6 route is present with a matching prefix and mask length less than or equal to 64-bits. This issue affects Summit X460-G2, X670-G2, and X770 switches.
xos0062133	STP flush event does not happen after ports are quickly disabled, and then enabled.
xos0062145	With QoS configuration, ACL process signal 11 ends unexpectedly after rebooting.
xos0062240	Port that was administratively disabled becomes up after enabling rx pause.
xos0062271	CLI memory leak occurs when executing show commands with include option through script.
xos0062277	The command <code>show vlan vlan_list</code> does not show information for dynamic VLANs nor the Default VLAN. Error appears.
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).
xos0062427	EDP process ends unexpectedly when CDP packets without portId TLV are received.
xos0062441	The process rtMgr ends unexpectedly when IPv6 static route is deleted.

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

CR Number	Description
xos0062472	Source MAC addresses learned through CDP packets received on EAPS-blocked ports cause traffic to be dropped.
xos0062570	In SummitStacks, executing the command <code>enable sflow ports</code> all enables sFlow inappropriately on stacking ports.
xos0062705	Kernel oops can occur after clearing IPMC FDB in a stack.
xos0062789	Disabling learning on LAG ports does not flush FDB entries.
xos0062879	Transceiver information shows same Rx power value for 4x10G partition ports even though some ports are in ready state.
xos0063089	Kernel oops triggered infrequently during continuous addition/deletion of ARP entries for long durations.
xos0063359	The process <code>rtmgr</code> might end unexpectedly after executing <code>disable bgp</code> , and then <code>enable bgp</code> , or after <code>disable port</code> , and then <code>enable port</code> , or after rebooting a switch containing BGP routes.
xos0063368	In an MLAG configured switch, FDBs are not installed in hardware after reboot if there are frequent MACMoves between MLAG port and ISC.
xos0063134	Traffic stops after disabling, and then enabling LAG portst having pstag with static FDB
xos0063245	With IGMP per-VLAN mode, VRRP flaps occur after adding tagged ports to VLANs.
xos0063457	Configuration for adding network VLAN port in STP for subscriber is not saved.
xos0063521	A few IBGP routes are not updated in routing table when <code>disable bgp</code> and <code>enable bgp</code> commands are executed in quick succession.
Summit X460-G2 Series Switches	
xos0061486	Combo ports have unsupported autonegotiation and half-duplex settings.
xos0062425	On Summit X460-G2 series switches, the primary port is incorrectly set as 40 when it should be 41. Under certain conditions, this can cause a kernel crash.
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.
xos0063071	Add support for ONEPolicy IP socket classification.
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.
xos0061704	With SSH2 enabled, reboot the switch and force some other standby node to become the master node. Key becomes invalid on new master node.
Summit X670-G2	

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

CR Number	Description
xos0061791	On SummitStacks containing master and standby nodes of different switches, the standby node may go to failed state after a node reboot.
xos0062166	On Summit X670-G2 series switches configured with L3VPN, executing the <code>clear iparp</code> command causes the switch to reboot with Kernel Oops.
xos0063204	Traffic stops on LAG ports when frequently modifying the sharing group.
SummitStack	
xos0057835	In SummitStacks, clear-flow sampling period is incorrectly calculated.
xos0061799	Precedence order between policy port rules and policy MAC-based rules is not preserved following a master/backup Failover.
xos0061841	FDB entries are not learned again after limit learning is unconfigured, and then configured again, with PSTAG configuration in SummitStacks.
xos0061957	HAL process ends unexpectedly during failover when switches have ACL policy without meter action.
xos0062084	Rebooting modules with only policy configurations clears their policy port configurations when they rejoins the stack.
xos0062123	Port groups do not appear in the <code>show configuration</code> command. However, they do appear in the <code>show ports group</code> command.
xos0062238	On a stacked system, configuration of a user-defined CoS value's <code>etsysCos8021dPriority</code> using the MIB can return success when the set actually failed (as seen by a subsequent <code>get</code>).
xos0062291	Applying the same policy MAC admin rule to multiple ports produces the following error message: <code>hardware configuration of rule failed for policy</code> .
xos0062367	ACL process ends unexpectedly on repeated refresh of ACL policy with clear-flow action.
xos0062504	You can set a GTPP "peer delay current interval" outside of the correct range of -3 to 17..
xos0062522	In SummitStack switches, standby slots go to failed state when a very large number of log messages are continuously generated in the switch.
xos0063242	Stacks configured as DHCP clients do not respond to pinging after failover.
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.
xos0063490	CFM stays down after slot reboot on a stack.
ACL	
xos0054348	Cannot delete flow names after deleting, and then creating, the flow while the ACL is installed.
xos0054714	When ACLs are applied in both ingress and egress directions, you cannot see egress direction using SNMP. When a policy has more than one counter, using SNMP, you can only check the updates from the first counter, and subsequent counters do not appear.

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

CR Number	Description
xos0059924	The output of the command <code>show access-list meter ports</code> displays additional meter name when only one meter is applied using ACL policy.
xos0060716	Need support for new ACL action "redirect-vlan" to redirect matched packets to all ports in specified VLANs.
xos0061922	Dynamic ACLs applied as "any" fail to install in hardware after upgrading ExtremeXOS from any release other than EXOS 15.3.
xos0062156	ACL Manager API slice type can be off by one.
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.
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.
xos0063547	Process ACL ends unexpectedly after applying a policy file with source zone as a match condition.
AVB	
xos0062494	Source MAC addresses learned through MVRP packets on a blocked port (STP) cause traffic to be dropped.
BGP	
xos0058441	After creating a BGP peering session between link local IPv6 addresses with the scope ID specified, deleting the VLAN containing link local IPv6 address, and then issuing the command <code>show configuration bgp</code> , switch reboots with <code>Epm application wdg timer warning</code> error message.
xos0060641	When BGP is administratively shut down, it does not send notifications to peers.
xos0060680	Switch stops responding after executing <code>clear bgp neighbor all counters</code> on a switch without BGP configuration.
xos0060749	Configuring, enabling, disabling, or deleting BGP neighbors using link local address results in the following error message: <pre>Error: cmBackendXmlParseEnd Failed to convert "bgpCfgPeerRemoteAddr" value "fe80::204:96ff:fe97:efef/ brian-to-112"</pre>
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-RI

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

CR Number	Description
xos0061505	After a topology change in the network, BGP routes requiring two levels of recursive lookup are programmed in hardware with incorrect next hops.
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.
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.
xos0063173	Process dcbgp ends unexpectedly with signal 11 after issuing the command <code>show bgp neighbor</code> .
Chalet	
xos0060354	ExtremeXOS Chalet using IPv6 does not work with HTTPS.
xos0062016	Command line process memory leak occurs when accessing switches with Chalet.
xos0063255	In Chalet, VLANs are sorted incorrectly.
ClearFlow	
xos0062629	Clearflow rule does not work properly if there is dot(.) in the ACL counter.
EAPS	
xos0061038	Loops occur in EAPS-protected VLANs, after peer reboot, if a VLAN's port is also protected by ELSM.
xos0061385	EAPS process ends unexpectedly after deleting EAPS shared-port configuration.
ELRP	
xos0062460	The <code>show configuration</code> 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.
ESRP	
xos0061965	Configuring ESRP member VLANs (VRRP-enabled) produces errors.
FDB	
xos0059481	Static FDB is programmed incorrectly in hardware after a stack failover.
IGMP	
xos0062914	The process mcmgr ends unexpectedly after receiving corrupted IGMPv3 join packets on MLAG ports.
LAG\MLAG	
xos0062428	Member ports with a modified speed configuration that is different than the master port should not be allowed in LAG.

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

CR Number	Description
xos0063365	Frequent MLAG bulk syncs observed due to checksum mismatch between MLAG peers when ISC port was added as an untagged port to a tagged VLAN and VRRP was running between the peers.
MPLS	
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.
xos0061943	MPLS process ends unexpectedly when get-next is done with incomplete OID for mplsXCIndex.
xos0062045	LLDP packets are tunnelled over L2VPn.
xos0062300	CEP CVID Ranges, other than first VLAN, fail when access port is a trunk.
xos0062301	Packet drops occurs between customer edge switches when VMAN and CVID tag are the same.
xos0062380	Switch rejects incorrect LSP configurations as expected, but this operation still uses LSP indexes in hardware.
xos0062754	VPLS traffic egresses out with dot1q tag when secondary EtherType is configured.
xos0063271	Layer 3 packets in non-default virtual routers are slow-path forwarded after disabling MPLS in the peer switch.
xos0063478	Traffic drop occurs while adding new member port to the existing LAG group and PSTAG is configured on the port.
OpenFlow	
xos0060531	Deleting VMAN deletes the VLAN configuration, but not associated OpenFlow logical ports.
Optics	
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.
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.
xos0060264	The output of the <code>show port transceiver info</code> command for optics inserted in 40G/100G ports might be abnormally lengthy if the same command is executed from two different CLI sessions simultaneously.
xos0062719	Allow use of 3rd-party optics without any additional license.
xos0063120	Error message "CFP2 modules >= 18 W unsupported" incorrectly appears for Finisar Corp CFP2 LR4 optics.
OPSFV2	
xos0061855	Configured OSPF neighbor is not retained after rebooting.

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

CR Number	Description
xos0063380	Error message appears after rebooting switch with OSPF configuration: "Error while loading "ospfInterface": ERROR: 0.0.0.0 is not a valid configured neighbor for interface".
Power	
xos0062113	The <code>show power</code> command output does not display power usage for PSUs with part numbers starting with "800515".
QoS	
xos0061027	For SummitStacks, creating or deleting non-default QoS profiles may cause some ports to flap.
xos0062050	QoS committed rate configurations for port groups are not loaded properly after a save and reboot.
Security	
xos0057679	Account user name and password are not encrypted in logs when cli-config-logging is enabled.
xos0058808	Rarely, MAC addresses of authenticated clients learned on NetLogin-enabled ports are not programmed in hardware.
xos0060909	In UPM profiles the variable EVENT.TIME incorrectly has the current time rather than the time when the event was queued/triggered.
xos0061433	extremeNetloginUserLogoutTrap is received with errors.
xos0061597	After authenticating a NetLogin client, executing the command <code>clear netlogin state port</code> , and then checking that ID-mgmt has deleted the clients, log displays UnDscvrld EMS message, which does not indicate the identity user.
xos0061652	Netlogin Dot1x: Authenticated value should be "Yes, Local" for clients with invalid password authenticated on auth failure VLAN.
xos0061781	Identity manager entries become stale when clients are moved from one port to another in sub-VLANs.
xos0061797	Dot1x client moves to authentication failure VLAN if authentication failed due to incorrect supplicant password or framework failure, such as error in VLAN movement, etc.; even if web-based NetLogin is enabled.
xos0061820	Dot1x clients move to authentication failure VLAN when web-based NetLogin is enabled globally.
xos0061868	With protocol order as MAC dot1x, web-based UPM profile is not executed for the client, which is authenticated as MAC.
xos0062366	After rebooting, DHCP binding entries are not restored using vr-default.
xos0062674	UPM profile fails to set the variables received from the RADIUS server using VSA 212.
xos0062965	Policy process ends unexpectedly with signal 6 when master node goes down.
xos0063090	NetLogin client does not move into authfail VLAN when user is absent from local database.

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

CR Number	Description
xos0063248	NTP MD5 authentication with NTP server is failing.
xos0063445	NetLogin: FDB is not in synch when changing VLAN VSA's dynamically.
xos0063506	Traceroute MAC address in CFM domain does not return information about destination switch.
SNMP	
xos0059964	SNMP poll for MIB dot3StatsDuplexStatus always returns unknown(1) when ports are configured with auto-negotiation on.
xos0060792	SNMP authentication failure log message and trap is inappropriately generated when switch detects "Not In Time Windows" error.
xos0061379	Switch temperature value retrieved using SNMP get operation is incorrect.
xos0061886	SNMP master process ends unexpectedly with signal 6 with certain sequence of snmpbulkget and snmpget.
xos0061945	SnmpSubagent crash occurs when snmpset executed on the last row in EAPSMbrVlanEntry.
xos0063349	Switch stops responding to SNMP requests if SNMP get for multiple OIDs is continuously initiated.
STP	
xos0062701	HAL timeout occurs while rebooting a stack with STP configuration.
TWAMP	
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.
VLANs	
xos0054039	IP multicast traffic is not forwarded on PSTAG VLANs when it shares ports with other IGMP snooping-enabled VLANs or other L3 VLANs.
xos0060184	After configuring MVRP registration forbidden, the command is accepted and registration is forbidden. However, the show configuration mrvp command does not display this configuration and this configuration is not saved after a reboot.
xos0062255	CEP CVID configurations is missing after adding/deleting the port from sharing.
xos0063207	Error occurs while adding LAG ports as tagged in one VMAN and untagged in another VMAN, even though the VMAN EtherType is primary.
xos0063257	Saving configuration fails/times-out when VLANs added to a mirror filters are renamed.
xos0063274	VLAN packets are egressing with VMAN ethertype when an egress port is deleted from a VMAN that is also part of a VLAN.