ExtremeWare Release Note

Software Version 7.0.1 Build 11

Extreme Networks, Inc.
3585 Monroe Street
Santa Clara, California 95051
(888) 257-3000
http://www.extremenetworks.com

Published: March 2003
Part Number: 120165-00 Rev 02

Author: Rich Small

Editor: Rich Small

Production: Rich Small

Special Thanks: Paul

Contents

Chapter 1 Overview

New Features in ExtremeWare 7.0 **11**
- Features Added or Enhanced in ExtremeWare 7.0.1b11 11
- Features Added or Enhanced in ExtremeWare 7.0.0b68 13
- Features Added or Enhanced in ExtremeWare 7.0.0b61 16

Supported Hardware **17**
- BlackDiamond Component Support 18
- Alpine Component Support 19
- Summit Component Support 20
- GBIC Support 20
  - *Mini-GBIC Support* *20*

Chapter 2 Upgrading to ExtremeWare 7.0.1

Staying Current **21**

Upgrading ExtremeWare **21**
- Upgrading Switches to ExtremeWare 7.0.0 22
  - *Save the Current Configuration* *22*
  - *Upgrade the BootROM to Version 7.6* *23*
  - *Upgrade to ExtremeWare 6.1.9* *23*
  - *Upgrade to ExtremeWare 6.2.2b56* *23*
  - *Upgrade the BootROM to Version 7.8* *24*
  - *Upgrade to ExtremeWare 7.0.1b11* *24*
  - *Upgrade ATM, MPLS, ARM, PoS, T1, E1, or T3 Modules* *25*
- Upgrading an Alpine 3802 to ExtremeWare 7.0.1 25

Downgrading Switches **26**

Chapter 3 Supported Limits

Supported Limits **27**

Chapter 4 Clarifications, Known Behaviors, and Resolved Issues

Clarifications and Known Behaviors **33**
System Related – All Systems 33
GBIC Type in the show ports configuration Command Output *33*
Do Not Telnet to Port 80 and Continuously Press Keys *33*
Smart Redundancy Enabled in Saved Configuration *34*
Microsoft Load Balancing *34*
Telnet and the show ports Command *34*
The show configuration Output *34*
Configure Slots or VLANs Before Uploading a Configuration *34*
LACP not Supported *34*
Upgrading to ExtremeWare 7.0 and Bi-Directional Rate Shaping *34*
Upgrading to ExtremeWare 7.0 and Debug-Trace *34*
Upgrading to ExtremeWare 7.0 and OSPF *34*
Routing Traffic Through the MGMT Port *35*
Configuring the Timezone *35*
Blank Space in show port info detail Command Output *35*
Using an ExtremeWare 7.0 Configuration with an Earlier Image *35*
Console Response with a Large Number of ARP Entries *35*
Configuring 1000Base-T Ports for 10,000 Mbps *35*
The show log chronological Command *35*
BOOTP-Dependent Routes in Downloaded Configuration not Created *35*
Enable Flow Statistics Ping-Checking *35*
UDP Echo Transmit Rate *36*
The disable learning Command and Flooding *36*
Port Mirroring *36*
Setting Auto-negotiation Off on a Gigabit Port *36*
Enabled IdleTimeouts and Console Connections *36*
User Accounts *36*
TFTP Download of Configuration Files *36*
Port Tag Limitation *36*
BlackDiamond 37
MPLS Hello Packets *37*
Slot Failure Messages During a Broadcast Storm *37*
Hot-Inserting an MSM Disrupts MPLS and ARM Modules *37*
No Image Information Reported to SNMP with One MSM *37*
MPLS and CPU DoS Protect *37*
Duplicate Precedence Rules *37*
BlackDiamond 6816 MSM C and D Diagnostics Messages not in Syslog *37*
Synchronize a Newly Installed MSM64i *38*
Disabling CLI Paging from the Slave MSM64i *38*
Limited Commands Mode and the reboot Command *38*
The unconfig switch all Command *38*
Dynamic Memory Scanning and Mapping Module Support *38*
Extended Diagnostics *38*
BlackDiamond 6816 MIB Value for Input Power Voltage *39*
Backplane Traffic *39*
QoS *39*
Alpine 39
Alpine 3802 show switch Output Shows Incorrect PSU Placement *39*

System Health Check Events Might Not Be Logged 39
Configuring Two Multilink groups on One T1 or E1 Module 39
Limited Commands Mode 39
T1 and E1 Error Message 39
VDSL Modules in a Half-Duplex Link 39
Summit 40
Health Check Error Messages 40
Limited Commands Mode 40
Summit48i Redundant PHY 40
Summit48i Single Fiber Signal Loss 40
SNMP Results for Power Sources 40
Summit48si MIB value for Input Power Voltage 40
Command Line Interface (CLI) 40
Only US Character Set Supported 40
The show iproute Command 40
Serial and Telnet Configuration 41
Displaying Management Port with show port config 41
Auto Negotiation and 1000BaseT Ports 41
Switching and VLANs 41
Saving ip-mtu Settings 41
FDB 41
Configure Less Than 400 Ports in a VLAN 41
Cannot Delete "mgmt-1" VLAN 41
VLAN priority and STP, EDP 41
Default Routes or Static Routes 42
Configuring a Protocol Filter with 'ffff' 42
Deleting Protocols from a VLAN 42
MAC Based VLANs and DHCP Relay 42
VLAN to VLAN Access Profiles 42
Load Sharing 42
Spanning Tree 43
MAC Security 43
Mirroring 44
QoS 44
The qosprofile Accepts a Value Greater than 100% 44
Re-Ordering Access List Precedence Numbers 44
Access List FDB Entries not Cleaned Up 44
Access Lists Using the IP Deny Any Rule 44
Access Lists and IP Fragmentation 44
QoS Configuration Bandwidth Parameters 44
Access List Precedence Intervals 45
Creating Access Lists from Multiple Sessions 45
QoS and dot1p 45
5,120 Access Lists and SNMP 45
Monitoring QoS and the show port qos Command 45
Ingress QoS 45
Ingress QoS Not Supported on Other Modules 45
The show ports ingress stats Command Truncates 45
Bi-Directional Rate Shaping 46

Locking and Unlocking Learning 46
Loopback Port Must be on Same Module 46
1000Base-T Ports as Loopback Ports 46
Changing the Configuration of a Loopback Port 46
EAPS 46
WAN Modules Not Currently Supported with EAPS 46
Do Not Configure a Hello Time of 0 46
A Large EAPS Configuration with a Link Transition 46
Changing the Protected VLAN Tag 46
EAPS Performance Statistics 47
EAPS and STP or EMISTP 47
EAPS Secondary Port Recovery 47
ESRP and EAPS Secondary Port 48
Incorrect show vlan Output 48
ESRP 48
Configure a Neighbor Timeout Less than 6 Times Hello Timer 48
Transition Incorrectly Logged 48
Dual Master Recovery Not Logged 48
A Flapping Redundant Link Might Cause ESRP to Fail Over 48
ESRP and Ingress Rate Shaping 48
ESRP and Protocol-Based VLANs 48
ESRP and Load Sharing 49
Hot-Swapping a Module with 5,000 ACLs 49
Traffic Convergence Time 49
ESRP PDUs on Ports 49
Multiple ESRP VLANs 49
ESRP Interoperability 49
Mixing Clients and Routers on an ESRP-Enabled VLAN 49
ESRP and Bi-Directional Rate Shaping 49
VRRP 50
The show tech-support Command Through Telnet 50
Increase Advertisement Interval When CPU is Busy 50
Backup Transition Creates Duplicate Packets 50
Changing the Advertisement Interval 50
Changing the Priority 50
The track-diagnostic and track-environment Features Not Supported 50
IP Unicast Routing 50
Deleting a Static Entry Using SNMP 50
The show iproute Output 50
Traffic Crosses Layer 3 Boundary 51
Moving a sub-VLAN Client 51
No Static ARP Entries 51
VLAN Aggregation and ESRP 51
ARP Entry Age 51
Multinetting and Client Default Gateways 51
Multinetting and the Show VLAN Stats Command 51
Multinetting and VRRP 51
RIP Routing 51
RIP V2 Authentication 51

RIP in Conjunction with other Routing Protocols 51
OSPF 52
Default Route Entries in the IP FDB 52
Disable OSPF Before Adding or Removing External Area Filters 52
BGP 52
Multi Exist Descriminator Not Compared 52
Route Dropped if Switch's AS is First AS in Path 52
BGP Set Community Inadvertantly Advertised 52
Do Not Use configure access-profile Command to Set Community 52
Best Routes 52
BGP Loops 52
Redistributing BGP Routes to OSPF 53
Removed encrypted Option from enable bgp neighbor password Command 53
IP Multicast Routing 53
Use the always Parameter to Guarantee Advertisement 53
Cisco Interoperation 53
Traffic Rate Exceeding Last Hop Threshold 53
IPX Routing 53
Tuning 53
IPX and Round-Robin Loadsharing 53
IPX Performance Testing Using Traffic Generators 53
IPX and Bi-Directional Rate Shaping 54
Security and Access Policies 54
Simulated Mode Creates ACL 54
Network Login Design Guidelines and Limitations 54
Configure RADIUS with Existing VLAN for Network Login 54
RADIUS and the BlackDiamond 54
RADIUS and Telnet 55
TACACS+ and RADIUS 55
Network Login and Saving the Configuration 55
The show netlogin Command Output 55
Flow Redirection 55
Enumeration Mode Redirects ICMP Packets 55
Cache Servers Set To "Down" Under Sustained High Traffic Loads 55
Health Checking Cannot be Disabled 55
NAT 56
Vista 56
VLAN Ports Tagging Information Incorrect 56
Blackhole Flag Missing 56
Multicast Address Display 56
Configuration Statistics PSU Display 56
Closing Internet Explorer 4.0 56
Vista and RADIUS 56
Configuration Options with Large Number of Interfaces 56
SNMP 57
Modular Switch get Error 57
SNMP v1 Traps 57
SNMP and ACLs 57
Adding or Deleting a Trapreceiver 57

Incrementing the intfIf Value 57
WinSCP2 Not Supported 57
SNMP ifAdminStatus MIB Value 57
Trap Receivers as Broadcast Entry 57
Bridge MIB Attributes 57
SNMP Time-out Setting 58
SNMP Access Profile 58
SNMP and Auto-negotiation Settings 58
SNMP and the BGP MIB 58
SNMP and the FDB MIB 58
Extreme Fan Traps 58
Extreme Power Supply Traps 58
DHCP 58
Diagnostics and Troubleshooting 59
The show diagnostics backplane-utilization Command Available 59
Spurious Message When system-down is Configured 59
The use configuration Command 59
Output of the show diagnostics Command 59
Configure Auto-Recovery to online or Alarm-Level to traps 59
Error Count Not Accurate 59
Configuring Diagnostics Mode Off 59
Disable Remote Syslog Before Enabling IPARP Debug-Tracing 59
Rebooting Using SNMP or RMONII With Reboot Loop Protection 60
Configuring a New Threshold for Reboot Loop Protection 60
The card-down Option 60
Do Not Use a Count of One for Reboot Loop Protection 60
Documentation 60
Issues Resolved in ExtremeWare 7.0.1b11 **60**
General 60
BlackDiamond 60
Summit 61
Issues Resolved in ExtremeWare 7.0.0b68 **61**
BlackDiamond 61
Alpine 61
ESRP 61
Spanning Tree 61
QoS 61
IS-IS 62
Issues Resolved in ExtremeWare 7.0.0b61 **62**
General 62
BlackDiamond 62
Summit 63
IP Multicast Routing 63
RIP Routing 63
EAPS 63
FDB 64
BGP 64

OSPF 64
Spanning Tree 64
ESRP 64
VLANs 65
IS-IS 65
NetFlow 65
SNMP 65

1 Overview

These Release Notes document ExtremeWare 7.0.1 build 11. ExtremeWare 7.0.1 introduces new hardware products and software features.

NOTE

You can only load ExtremeWare 7.0 on a switch running ExtremeWare 6.2.2 (or later). To install ExtremeWare 7.0, see "Upgrading ExtremeWare" on page 21.

This chapter contains the following sections:

- "New Features in ExtremeWare 7.0" on page 11
- "Supported Hardware" on page 17

For information on issues resolved from previous releases, you can obtain previous versions of release notes through a login account on the Extreme Networks Support web site at http://www.extremenetworks.com/support/support.asp.

New Features in ExtremeWare 7.0

Following are descriptions of features introduced or enhanced in ExtremeWare 7.0. These features are documented in detail in the *ExtremeWare Software User Guide* or the *ExtremeWare Software Command Reference Guide*, unless otherwise noted.

Numbers in parentheses are for internal use and can be ignored.

Features Added or Enhanced in ExtremeWare 7.0.1b11

The following features were added or enhanced in ExtremeWare 7.0.1b11:

- You can configure 802.3x flow control on your "3" series module. Because these modules are oversubscribed to the module switch fabric, traffic can congest. Flow control allows you to stop incoming traffic when too much congestion occurs.

  Flow control sends a PAUSE frame to the transmitter when traffic approaches the congestion threshold for a specific queue. The PAUSE frame is sent *before* the queue overflows, so throughput is slightly reduced when flow control is enabled. Flow control is auto-negotiated and is disabled if both ports do not support it.

Flow control is disabled by default. To enable 802.3x flow control, use the following command:

```
enable flow-control ports [<portlist> | all]
```

To disable 802.3x flow control, use the following command:

```
disable flow-control ports [<portlist> | all]
```

Use the `all` keyword to specify all configured "3" series ports.

To see the flow control configuration, use the `show ports configuration` command. DSBL indicates that flow control is disabled on that port. ENBL indicates that flow control is enabled while auto-negotiation is off for that port. If flow control and auto-negotiation are both enabled, the negotiated flow control value is displayed.

DiffServ examination is enabled by default on all "3" series ports; DiffServ examination is disabled by default on all other ports.

This feature is not documented in the *ExtremeWare 7.0.0 Command Reference Guide* or the *ExtremeWare 7.0.0 User Guide*.

- You can now configure ingress QoS on your "3" series module. Ingress QoS is used to prioritize traffic received on "3" series ports.

  Congestion can cause ingress traffic to be dropped on oversubscribed "3" series I/O modules. Ingress QoS allows received traffic with different VLAN priority values, different Diffserv codepoints (IP TOS), or from different VLANs to be routed to up to 8 different ingress queues. This allows for specified traffic types to be queued separately so they remain unaffected by congestion in the main ingress queue. Default ingress QoS settings avoid discards of higher priority traffic (identified by diffserv codepoint or VLAN priority) in the presence of ingress congestion.

  To configure ingress QoS, use the following command:

  ```
  configure qostype ingress priority [ vlan | diffserv | dot1p ] <qos-priority (0-15)>
  ```

  The `diffserv` parameter specifies the priority based on DiffServ information. The default is 3.

  The `vlan` parameter specifies the priority of VLAN ID-based input queue selection. The default is 2.

  The `dot1p` parameter specifies the priority based on dot1p information. The default is 1.

  The `priority` range is 0-15 (15 is the highest priority). Each queue selection criteria must have a unique priority (no ties).

  To restore the default ingress QoS settings, use the following command:

  ```
  unconfigure qostype ingress priority
  ```

  To configure a VLAN to use a particular ingress QoS profile, use the following command:

  ```
  config vlan <vlan name> qosprofile ingress [<IQOS profile> | none]
  ```

  The `vlan name` parameter specifies a VLAN name.

  The `IQOS profile` parameter specifies an ingress QoS profile, such as iqp1.

  The `none` parameter specifies that traffic from this VLAN is not associated with any ingress queue based on VLAN ID. This is the default setting.

  To view the ingress QoS settings, use the following command:

  ```
  show qostype ingress priority
  ```

  To view the real-time ingress statistics, use the following command:

  ```
  show ports {<portlist>} ingress stats {detail}
  ```

  The output indicates the following:

  — Port Number

— Link Status—The current status of the link. Options are:

- Ready (R): The port is ready to accept a link.
- Active (A): The link is present at this port.
- Disabled (D): The link is disabled at this port.
- Not Present (NP): The link is not present at this port.

— High Priority Bytes—Sum, per port, of the bytes forwarded for received high-priority packets. Reserved for a future release.

— Low Priority Bytes—Sum, per port, of the bytes forwarded for received low-priority packets. For this release, all packets are considered low priority in this context.

— Received Total Bytes—The total number of bytes that were received by the port.

— Receive Bytes Dropped—Total number of bytes dropped for this port.

— Total Percent Dropped—Percentage of incoming bytes dropped due to oversubscription congestion. Displayed with a precision of 1/100 of a percent.

— Transmit XOFF—Total number of XOFF flow control packets sent from this port.

When `detail` is specified, the following additional information is displayed per ingress queue:

— Queue—One of 8 ingress queue names for this port.

— High Priority Bytes—Sum, per ingress queue, of the bytes forwarded for received high-priority packets. Reserved for a future release.

— Low Priority Bytes—Sum, per ingress queue, of the bytes forwarded for received low-priority packets. For this release, all packets are considered low priority in this context.

— Total Percent Dropped—Percentage of incoming bytes on this queue dropped due to oversubscription congestion. This is determined using cumulative counters, so is not a rate. This will be displayed with a precision of 1% and is accurate within 3%.

— Byte Rates—The following three rate values will always either add up to 0% or 100%:

- High Priority Percentage—The ratio of high priority traffic forwarded on this queue to the total bytes received on this queue.
- Low Priority Percentage—The ratio of low priority traffic forwarded on this queue to the total bytes received on this queue.
- Dropped Percentage—Percentage of receive bytes dropped by this queue relative to the total number of bytes input to this queue.

In addition, the `configure slot module` command now includes the new modules.

This feature is not documented in the *ExtremeWare 7.0.0 Command Reference Guide* or the *ExtremeWare 7.0.0 User Guide*.

- The output of the `show version` and `show switch` commands has been updated to add more image information (PD2-125850501). The ExtremeWare MIB has also been updated.

Features Added or Enhanced in ExtremeWare 7.0.0b68

The following features were added or enhanced in ExtremeWare 7.0.0b68:

- You can now configure multiple T1 and E1 ports on the same module in the same VLAN when using BCP (PD2-80806101). To achieve this, the T1 and E1 modules maintain a subset of the switch's FDB entries. The SMMi and WAN module FDBs are synchronized via occassional SMMi flooding of dynamic entries. Static entries are synchronized as you enter them.

The following features are no longer supported on T1 or E1 modules:

— T1 port mirroring

— Static Load sharing

— Software-Controlled Reduntant Ports

— ACLs on a per port basis

— Per Port Egress QOS

— Traffic Grouping for source ports

— BiDirectional Rate Shaping

— DLCS

— MAC address and protocol-based VLANs that include T1 ports

— VLAN aggregation

In addition, layer 2 multicast traffic is treated as broadcast traffic by the T1 and E1 modules.

This feature is not documented in the *ExtremeWare 7.0.0 Command Reference Guide* or the *ExtremeWare 7.0.0 User Guide*.

- You can now block the SQL Slammer DoS attack. SQL Slammer causes high CPU utilization on the next-hop switch servicing multicast requests as IGMP sender entries are quickly populated into the multicast sender list. This leads to a high number of multicast entries in the IGMP snooping entry table, and a message similar to the following in the system log (PD2-118292101):

  ```
  <WARN:HW> tBGTask: Reached maximum otp ExtraMC index allocation
  ```

  To block and clean up after this attack:

  a Block the attack by creating an ACL to block port 1434 using the following command:

  ```
  create access-list UDP dest any ip-port 1434 source any ip-port any
  ```

  b Remove affected SQL servers from the network (you can simply disable the port connecting the server).

  c Clean up the existing IGMP snooping entries and IPMC cache using the following commands:

  ```
  igmp snooping
  clear ipmc cache
  ```

  d Disable IGMP snooping on the affected switches. Disabling IGMP snooping affects routing protocols using multicast addresses and multicast traffic on that switch.

  This feature is not documented in the *ExtremeWare 7.0.0 Command Reference Guide* or the *ExtremeWare 7.0.0 User Guide*.

- BGP no longer sends withdraws to a neighbor for routes that were not advertised to the neighbor (1-8P0O9).

  Route filters such as access-profiles, route-maps, or Network Layer Reachability Information (NLRI) filters filter advertisements of routes to BGP peers. BGP no longer withdraws routes from neighbors if the routes were filtered. This provides the following benefits:

  — Reduces peer routers' BGP control processing time

  — Reduces bandwidth overhead over peer session links

  — Improves switch resource utilization by reducing the number of locally originated packets

  The output of the `show bgp neighbor <ip address> transmitted-routes` command now displays the local attributes of the routes that were transmitted to the neighbor, rather than the actual attributes that were transmitted to the neighbor.

This feature is not documented in the *ExtremeWare 7.0.0 Command Reference Guide* or the *ExtremeWare 7.0.0 User Guide*.

- You can now disable auto-polarity detection on the Summit48si (PD2-102329001). The Summit48si automatically detects and corrects the polarity of cables, simplifying installation and maintenance. You can disable this feature using the following command:

  `configure ports <all | portlist> auto-polarity <on | off>`

  The default setting is on. The `show ports {portlist | all } info detail` command displays the autopolarity setting.

  This command is not documented in the *ExtremeWare 7.0.0 Command Reference Guide*.

- You can now create overlapping FDB and ACL IP rules (PD2-63843734).

  IP ACL rules can be implemented either as an ACL rule or an FDB rule. ACL hardware is more flexible and has no restrictions on the kinds of IP rules that can be implemented. ACL rules are implemented on the I/O modules to which they apply. You can assign precedence values to ACL rules. IP rules with a precedence specified are implemented as ACL rules.

  FDB hardware does not limit FDB rules to 255 per module. However, FDB hardware does not support range values for IPSA, L4-DST-PORT, or L4-SRC-PORT. In addition, FDB rules cannot support all IP rule components, such as the ingress port component of IP rules. FDB rules apply to all ingress ports. You cannot assign precedence to an FDB rule.

  ExtremeWare, based on the rule components, decides to implement IP rules either as ACL rules or as FDB rules. Whenever possible, IP rules are implemented as FDB rules because ACL rules are limited.

  When a new IP rule is entered, ExtremeWare checks for overlap with existing IP rules. The new rule is rejected if either it or the overlapping rule does not have a specified precedence. By default, the precedence of FDB rules is higher than that of ACL rules.

  You can now use the new `ignore-overlap` option in the `create access-list` command to ignore overlapping IP rules. IP rules are still implemented as FDB rules if possible. IP rules with a specified precedence are still implemented as ACL rules. FDB rules still take precedence over ACL rules.

  The output of the `show access-list` command indicates rules added with the `ignore-overlap` option.

  This feature is not documented in the *ExtremeWare 7.0.0 Command Reference Guide* or the *ExtremeWare 7.0.0 User Guide*.

- You can now configure a MAC address to be permitted only on a set of ports (PD2-106654174). Secure MAC addresses, if learned, are still aged out like other dynamically learned entries, and can also be cleared. To configure the authorized set of ports on which the MAC address should be permitted, use the following command:

  `create fdbentry secure-mac <mac_address> vlan <vlan name> ports <portlist>`

  To clear all the dynamic, non-permanent blackholed entries that were created due to secure MAC violations, use the following command:

  `clear fdb blackhole`

  To see the number of blackhole entries created due to secure MAC violations, use the following command:

  `show vlan <vlan name> security`

  The output of the `show fdb permanent` command indicates secure MAC addresses.

  A new MIB table was added, and several traps modified so you can configure secure MAC addresses using SNMP.

This feature is not documented in the *ExtremeWare 7.0.0 Command Reference Guide* or the *ExtremeWare 7.0.0 User Guide.*

Features Added or Enhanced in ExtremeWare 7.0.0b61

The following features were added or enhanced in ExtremeWare 7.0.0b61:

- ExtremeWare now supports the following features in a single software image. These features formerly required you to load a dedicated software image on the MSM64i. Features marked with an "*" require a separate software image on the module:
  - MPLS*
  - ATM*
  - PoS*
  - ARM*
  - IS-IS
  - T1*
  - E1*
  - T3*
  - VDSL*
  - 10 Gigabit Ethernet
- Software signatures: each ExtremeWare image now contains a unique signature. The BootROM checks for signature compatibility and denies an incompatible software upgrade. In addition, the software checks both the installed BootROM and software and also denies an incompatible upgrade (PD2-92081601).

  NOTE

  ExtremeWare 6.2.2 build 56 is the first ExtremeWare release to incorporate software signatures. Thus, you must upgrade to ExtremeWare 6.2.2 build 56 before upgrading to later ExtremeWare 6.2.2 builds or ExtremeWare 7.0.

- The error diagnostic messages recorded in the syslog have been modified to provide you with more information on where data corruption is occurring within the switch. The following 5 categories of messages are added:
  - CPU—corrupted packets destined for the CPU
  - DIAG—corrupted packets across the switch backplane
  - EDP—corrupted ExtremeWare Discovery Protocol packets in bound to the CPU
  - EXT—user traffic is corrupted in packet memory in bound to the switching fabric or I/O module
  - INT—user traffic is corrupted in packet memory out bound from the switching fabric or I/O module
- You can now control the number of BGP routes that are deleted and reinstalled with a new gateway using the `config ipfdb route-add [clear-all | clear-subnet]` command (PD2-103735201):
- You can now configure the link detection level using the `config port <portlist> link-detection-level <link detection level>` command (PD2-86873002):

- You can now configure the neutral state timeout value for an ESRP-enabled VLAN using the `config <vlan> esrp esrp-neutral-timeout <neutral-timer(0-512, 0 restores dflt)>` command (PD2-104485403).
- You can now configure reboot loop protection using the `config reboot-loop-protection threshold <time-interval> <count>` command.
- RFC 2439, Section 4 “Stability sensitive suppression of route advertisement”: BGP Route Flap Damping (10223).
- OSPF route priority: you can now configure the same route priority for different route types (1-EIR7Z).
- The `show mpls tls-tunnel` command has been enhanced to show tunnels sorted by VCID (1-DTVM1).
- Selective LPM using an ARM module (1-F3PTC).
- Network Login automatic redirect to login page (1-DNV6U).
- You can now configure Network Login and Vista separately (1-F3PRL).
- VLAN tunneling is now supported over T1 and E1 ports (1-DMXX5).
- Legacy BCP based on RFC 1638 is now supported on T1 and E1 ports (1-DVWBL).
- The maximum number of VLANs supported in a switch has been increased to 4095 from a previous maximum of 3000 (1-F3PTP).
- IGMP Snooping can now be disabled without affecting OSPF and PIM (1-61BM5, 9316).
- If the MAC Address limit is exceeded or the wrong address is learned, ExtremeWare now sends a trap and logs an error in the syslog (1-DVWAP).
- TACACS+: You can now configure the timeout (1-EIYRX)
- When the TACACS+/RADIUS servers are unreachable and the local login occurs, ExtremeWare now displays the message “Could not connect to Authentication server. Logged in with local account”.
- IP FDB aging is implemented to reduce chance of hash collision and optimize FDB space usage as the IP FDB entries are aged (1-B2P35, 1-968RN).
- VDSL ETSI and ANSI standards (multiple rate configurations for upstream and downstream) (1-CHFXT).
- You can now configure T1 and T3 modules with Vista (1-CP11B).
- Summit48si fan speed is now automatically adjusted based on temperature (PD2-67168001).

Supported Hardware

Hardware in the following sections listed in *italics* is new for this release.

ExtremeWare 7.0 supports “*i*” series or “3” series products *only.*

ExtremeWare 7.0 requires BootROM 7.8.

Table 1 lists software filenames for the supported hardware that requires software.

Table 1: Software for supported hardware

Extreme Hardware	ExtremeWare Filename	BootROM Filename/Version
BlackDiamond 6816	v701b11.Gxtr or v701b11.SGxtr	Ngboot7.8.bin/7.8
BlackDiamond 6808	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
BlackDiamond 6804	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
Alpine 3808	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
Alpine 3804	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
Alpine 3802	v701b11.xtr or v701b11.Sxtr/ EW-70-3802.mig	Ngboot7.8.bin/7.8
Summit7i/7iT	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
Summit1i/1iT	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
Summit5i/5iT/5iLX	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
Summit48i	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
Summit48si	v701b11.xtr or v701b11.Sxtr	Ngboot7.8.bin/7.8
ARM module	v701b11.arm	v701b11.nprom/1.18
OC3 PoS module	v701b11.oc3	v701b11.nprom/1.18
OC12 PoS module	v701b11.oc12	v701b11.nprom/1.18
OC3 ATM module	v701b11.atm3	v701b11.nprom/1.18
MPLS module	v701b11.mpls	v701b11.nprom/1.18
T1 module	v701b11.t1	t1boot28.wr/2.8
E1 module	v701b11.e1	e1boot28.wr/2.8
T3 module	v701b11.t3	t3boot28.wr/2.8

NOTE

Please see the "Upgrading to ExtremeWare 7.0.1" chapter for special upgrade instructions.

NOTE

The BlackDiamond 6816 requires its own ExtremeWare image. The image that runs on other BlackDiamond, Alpine, or stackable switches does not support the BlackDiamond 6816.

BlackDiamond Component Support

BlackDiamond components supported with ExtremeWare 7.0.1 include:

Table 2: BlackDiamond component support

BlackDiamond Component
MSM64i
G16X³
G24T³
G12SXi

Table 2: BlackDiamond component support (continued)

BlackDiamond Component
G8Xi
G8Ti
F48Ti
WDMi
F96Ti
F32Fi
10GLRi
MPLS
ARM
P3cMi
P3cSi
P12cMi
P12cSi
A3cMi
A3cSi
DC Power Supply
110 VAC Power Supply
220 VAC Power Supply

NOTE

Do not install mixed versions of the power supplies in the same system. Install power supplies of the same type.

Alpine Component Support

Alpine components supported with ExtremeWare 7.0.1 include:

Table 3: Alpine component support

Alpine Component
SMMi
GM-16X³
GM-16T³
GM-4Si/Xi/Ti
FM-32Ti
FM-24MFi
FM-24Ti
FM-24SFi
GM-WDMi
WM-4T1i

Table 3: Alpine component support (continued)

Alpine Component
WM-4E1i
WM-1T3i
FM-8Vi
AC Power Supply
DC Power Supply

Summit Component Support

Summit components supported with ExtremeWare 7.0.1 include:

Table 4: Summit component support

Summit Module
Summit7i DC Power Supply

GBIC Support

The following table describes how each version of ExtremeWare interprets the media type of the installed GBIC, based on either the Vista web interface, or the `show port config` command. All versions correctly identify Parallel ID GBIC types; however, some versions do not correctly identify the Serial ID GBIC type because the Serial ID GBICs were introduced after the software was released.

Table 5: GBIC support

Software Release	1000BaseSX Parallel ID	1000Base-LX Parallel ID	1000Base-SX Serial ID	1000Base-LX Serial ID	LX70 Serial ID
Release 1.x	SX	LX	Not Supported	Not Supported	Not Supported
Release 2.x	SX	LX	LX	LX	LX
Release 3.x	SX	LX	CX	CX	CX
Release 4.x	SX	LX	SX	LX	LX
Release 6.x	SX	LX	SX	LX	LX70 (6.1.6 and above)
Release 7.x	SX	LX	SX	LX	LX70

Mini-GBIC Support

Extreme products support the Extreme mini-GBIC only. For reliability and stability reasons, third-party mini-GBICs are not supported at this time.

2 Upgrading to ExtremeWare 7.0.1

This chapter contains the following sections:

- "Staying Current" on page 21
- "Upgrading ExtremeWare" on page 21
- "Downgrading Switches" on page 26

CAUTION

You can only load ExtremeWare 7.0 (or later) on a switch running ExtremeWare 6.2.2b56 (or later).

Staying Current

For support purposes, Extreme Networks recommends operating the most current General Deployment (GD) release of ExtremeWare. New releases of ExtremeWare are usually released first as General Availability (GA) releases. A GA release has undergone full regression testing and is supported by your local ExtremeWorks Technical Assistance Center, but should be deployed with the understanding that it is a not a GD release.

Extreme Networks does not recommend that customers perform a network-wide upgrade with any new GA release. As a precaution, you should start with lab testing and edge installations before moving a GA release to the core of networks with over 10,000 nodes.

If you are an Extreme Assist customer, the latest release and release notes are available through the support login portion of the Tech Support web site at http://www.extremenetworks.com/

Upgrading ExtremeWare

You can only load ExtremeWare 7.0 on a switch running ExtremeWare 6.2.2b56 (or later). You can only load ExtremeWare 6.2.2 on a switch running ExtremeWare 6.1.9 (or later). Each of these versions require a different BootROM. Table 6 lists the BootROM required for each version of ExtremeWare.

Table 6: Required BootROM versions

ExtremeWare Version	BootRom Version
ExtremeWare 7.0.0 through ExtremeWare 7.0.1	BootROM 7.8

Table 6: Required BootROM versions

ExtremeWare Version	BootRom Version
ExtremeWare 6.2.2	BootROM 7.6
ExtremeWare 6.1.9 through ExtremeWare 6.2.1	BootROM 7.2
ExtremeWare 6.1 through ExtremeWare 6.1.8	BootROM 6.5

If your switch is running ExtremeWare 6.1.8 (or earlier), you must first upgrade to ExtremeWare 6.1.9, then upgrade to ExtremeWare 6.2.2b56 (or ExtremeWare 6.2.2b68). Following are specific instructions on upgrading to, and downgrading from, ExtremeWare 7.0.0 for Summit, Alpine, and BlackDiamond switches.

Upgrading Switches to ExtremeWare 7.0.0

To install ExtremeWare 7.0.1, you must:

1. Save the configuration to a TFTP server.
2. Upgrade the BootROM to Version 7.6 as described on page 23.
3. Upgrade to ExtremeWare 6.1.9 as described on page 23.
4. Upgrade to ExtremeWare 6.2.2b56 as described on page 23.
5. Upgrade the BootROM to Version 7.8 as described on page 24.
6. Upgrade to ExtremeWare 7.0.1b11 as described on page 24.
7. Upgrade ATM, MPLS, ARM, PoS, T1, E1, or T3 Modules as described on page 25.

If you have already installed ExtremeWare 6.1.9 through ExtremeWare 6.2.2b43, you can skip step 3. If you have already installed ExtremeWare 6.2.2b56 through ExtremeWare 7.0, you can skip steps 2, 3, and 4.

NOTE

The Alpine 3802 requires a different upgrade procedure, described on page 25.

Save the Current Configuration

Before upgrading ExtremeWare, save your configuration using the following steps. This preserves the ability to downgrade should it become necessary.

1. Ensure that the currently used configuration is stored in both the primary and secondary configuration spaces using the `save configuration primary` and `save configuration secondary` commands.
2. Configure the switch to use the primary image and the primary configuration using the `use image primary` and `use configuration primary` commands.
3. Verify that all of the above procedures were completed successfully with the `show switch` command.
4. Upload the configuration of the switch to a TFTP server for safekeeping using the `upload configuration` command.

Upgrade the BootROM to Version 7.6

ExtremeWare 6.1.9 requires BootROM 7.2 (or later). ExtremeWare 6.2.2 requires BootROM 7.6 (or later). Before you upgrade to ExtremeWare 6.1.9, upgrade to BootROM 7.6:

1. Download the BootROM using the `download bootrom [ <host_name> | <ip_addr> ] <ngboot.bin_name>` command.
2. Reboot the switch using the `reboot` command.

Upgrade to ExtremeWare 6.1.9

If you are running ExtremeWare 6.1.8 (or earlier), upgrade to ExtremeWare 6.1.9:

1. TFTP download ExtremeWare 6.1.9 to the primary image space using the `download image primary` command.

   CAUTION

   If you do not upgrade to ExtremeWare 6.1.9 before downloading ExtremeWare 6.2.2, the ExtremeWare 6.2.2 download will fail, and the following message will be printed from the system:

   ```
   ERROR: File too large
   ```

2. Reboot the switch using the `reboot` command. The previous configuration of the switch is preserved.
3. Verify that the correct BootROM and ExtremeWare version are loaded using the `show switch` and `show version` commands.
4. Check the log for configuration errors. Manually enter configurations that did not load.
5. If you configured Random Early Drop Probability in ExtremeWare 6.1.8 (or earlier), re-configure the Random Early Drop Probability using the `configure red drop-probability` command.
6. Save the configuration to the primary space.

Upgrade to ExtremeWare 6.2.2b56

If you are running ExtremeWare 6.1.9 to ExtremeWare 6.2.2b43, upgrade to ExtremeWare 6.2.2b56 (or ExtremeWare 6.2.2b68):

1. TFTP download ExtremeWare 6.2.2b56 (or ExtremeWare 6.2.2b68) to the primary image space using the `download image primary` command.
2. Reboot the switch using the `reboot` command. The previous configuration of the switch is preserved.

   NOTE

   ExtremeWare 6.2.2b56 (and later) stores 75 static log entries. Previous versions stored 100 entries. To accomodate the new entry limit, ExtremeWare 6.2.2b56 clears the static log after your first reboot. To preserve your static log entries, use the `show log` *command and save the output.*

3. Verify that the correct BootROM and ExtremeWare version are loaded using the `show switch` and `show version` commands.
4. TFTP download the saved configuration, and answer `y` at the prompt to reboot the switch.

5 Check the log for configuration errors. Manually enter configurations that did not load.

6 Save the configuration.

Do **not** save to the secondary configuration space until you are certain a downgrade to the previous image is not required.

NOTE

After upgrading from ExtremeWare 6.1.9 to ExtremeWare 6.2.2, the IGMP snooping leave time-out value will be changed from 10 seconds to 0. This results in an IGMP snooping membership entry being removed immediately when an IGMP leave is received from a host.

This is good for an environment where only one host is connected. Use the `configure igmp snooping leave-timeout` *command to change the leave time-out value back to 10 seconds.*

Upgrade the BootROM to Version 7.8

ExtremeWare 7.0.1 requires BootROM 7.8 (or later). Before you upgrade to ExtremeWare 7.0.1, upgrade to BootROM 7.8:

1 Download the BootROM using the `download bootrom [ <host_name> | <ip_addr> ] <ngboot.bin_name>` command.

2 Reboot the switch using the `reboot` command.

Upgrade to ExtremeWare 7.0.1b11

If you are running ExtremeWare 6.2.2b56 (or later), upgrade to ExtremeWare 7.0.1:

1 Upload the configuration to your TFTP server using the `upload configuration` command.

2 Clear your switch using the `unconfigure switch all` command.

3 Answer `y` at the prompt to reboot the switch.

4 TFTP download ExtremeWare 7.0.1 to the primary image space using the `download image primary` command.

5 Reboot the switch using the `reboot` command.

6 Verify that the correct ExtremeWare version is loaded on the switch using the `show switch` command.

7 TFTP download the configuration you saved in Step 1, and enter `y` at the prompt to reboot the switch.

8 Check the log for configuration errors. Manually enter configurations that did not load.

9 Save the new configuration to the primary space.

Do **not** save to the secondary configuration space until you are certain a downgrade to the previous image is not required.

10 If you are upgrading a BlackDiamond switch, synchronize the BootROM, image, and configuration across all installed MSM64i modules using the `synchronize` command. This command reboots the synchronized modules.

You can ignore any diagnostics failure messages generated by the synchronization.

11 Reboot the switch using the `reboot` command.

Upgrade ATM, MPLS, ARM, PoS, T1, E1, or T3 Modules

If you are using a ATM, MPLS, ARM, PoS, T1, E1, or T3 module, upgrade the module to ExtremeWare 7.0.1:

1. TFTP download the latest ExtremeWare version for the module using the `download image slot` command.

   NOTE

   T1, E1, and T3 modules must be using ExtremeWare 6.1.8b79 before upgrading to ExtremeWare 7.0.1.

2. Reboot the module using the `reboot slot` command.

   NOTE

   If you are upgrading multiple modules, skip step 2 until you have upgraded every module, then reboot the switch.

3. Download the BootROM using the `download bootrom slot` command.
4. Reboot the module using the `reboot slot` command.

   NOTE

   If you are upgrading multiple modules, skip step 4, upgrade every module, then reboot the switch.

Upgrading an Alpine 3802 to ExtremeWare 7.0.1

To upgrade an Alpine 3802 to ExtremeWare 7.0.1:

1. Upload the configuration to your TFTP server using the `upload configuration` command.
2. Upgrade to BootROM 7.8 using the `download bootrom` command.
3. Reboot the switch using the `reboot` command.
4. TFTP download ExtremeWare 6.1.8w3.0.1 b79 to the primary image space using the `download image primary` command.
5. Verify that the correct BootROM and ExtremeWare versions are loaded on the switch using the `show switch` and `show version` commands.
6. Clear your switch using the `unconfigure switch all` command (this is required by the beta software).
7. Answer `y` at the prompt to reboot the switch.
8. TFTP download ExtremeWare 7.0.0b46 to the primary image space using the `download image primary` command.
9. Reboot the switch using the `reboot` command.
10. TFTP download the latest ExtremeWare 7.0.1 build to the primary image space using the `download image primary` command.
11. Reboot the switch using the `reboot` command.
12. TFTP download the configuration you saved in Step 1, and enter `y` to reboot the switch.
13. Check the log for configuration errors. Manually enter configurations that did not load.

14 Save the new configuration to the primary space.

Do **not** save to the secondary configuration space until you are certain a downgrade to the previous image is not required.

Downgrading Switches

Assuming that the previous configuration is in the secondary configuration space and the previous image is in the secondary image space:

1 If you saved an ExtremeWare 6.1 (or earlier) configuration during the upgrade process, configure the switch to use that configuration with the `use configuration secondary` command.

If you did not save an earlier configuration, re-configure the switch or download a configuration at the end of this process.

2 If you did not save the earlier ExtremeWare image in the secondary image space, download the image using the `download image secondary` command.

NOTE

If you downgrade to an ExtremeWare version that does not support software signatures (ExtremeWare 6.2.2b56 or later supports software signatures), you must follow the upgrade procedures in the preceding sections to get back to ExtremeWare 7.0.1. You cannot switch between primary and secondary images on the switch unless they both support software signatures.

3 Use the image in the secondary image space with the `use image secondary` command.

4 Verify that the above procedures were completed successfully with the `show switch` command.

5 Clear your switch using the `unconfigure switch all` command (this is required by the beta software).

6 Downgrade to the appropriate BootROM version. The `show version` command displays the BootROM version as "Unknown" when the BootROM is downgraded.

7 Reboot the switch.

NOTE

When downgrading to a previous version of ExtremeWare, ensure that the switch configuration matches that version of ExtremeWare or below. Pointing the configuration to a new version of ExtremeWare and using a previous version of ExtremeWare is not supported. You will get a warning message from the system when attempting to do so.

8 If you did not save an ExtremeWare 6.1 (or earlier) configuration during the upgrade process, re-configure the switch or download a configuration.

3 Supported Limits

This chapter summarizes the supported limits in ExtremeWare.

Supported Limits

The table below summarizes tested metrics for a variety of features. These limits may change but represent the current status. The contents of this table supersede any values mentioned in the *ExtremeWare Software User Guide.*

Table 7: Supported limits

Metric	Description	Limit
Access List rules	Maximum number of Access Lists (best case).	5120
Access List rules—Summit	Maximum number of Access Lists in which all rules utilize all available options (worst case).	255
Access List rules—Alpine	Maximum number of Access Lists in which all rules utilize all available options (worst case).	255
Access List rules—BlackDiamond	Maximum number of Access Lists in which all rules utilize all available options (worst case).	255 per I/O module
Access Profiles	Maximum number of access profiles per switch.	128
Access Profile entries	Maximum number of access profile entries per switch.	256
BGP—Peer Groups	Maximum number of BGP peer groups per switch.	16
BGP—peers	Maximum number of BGP peers per switch.	200
BGP—routes, BlackDiamond, Summit7i, Alpine	Maximum number of routes received and contained in the BGP route table (best case).	1,275,000
BGP—routes, Summit1i, Summit5i, Summit48i, Summit48si	Maximum number of routes received and contained in the BGP route table (best case).	180,000
BGP—NLRI filters	Maximum number of NLRI filters per switch.	128
BGP—NLRI filer add entries	Maximum number of NLRI add entries per switch.	256
BGP—AS-Path filters	Maximum number of AS-Path filters per switch.	128
BGP—AS-Path filter add entries	Maximum number of AS-Path filter add entries per switch.	256

Table 7: Supported limits (continued)

Metric	Description	Limit
BGP—network statements	Maximum number of network statements per switch.	256
BGP—aggregate addresses	Maximum number of aggregate routes that can be originated per switch.	256
Jumbo Frame size	Maximum size supported for Jumbo frames, including the CRC.	9216
EAPS—Domains/switch	Maximum number of EAPS domains.	64
EAPS—Domains/ring	Maximum number of EAPS domains if no switch in the ring is connected to another ring.	64
EAPS—VLAN links	Maximum number of Control or Protected VLANs per EAPS domain.	4093
EAPS—Bridge links	Maximum number of EAPS bridge links per switch.	4096
EAPS—Master nodes	Number of Master nodes per EAPS domain.	1
EAPS—Switches	Maximum number of EAPS switches per ring.	No limit
EMISTP & PVST+ — maximum domains, Summit	Maximum number of EMISTP and PVST+ domains.	128
EMISTP & PVST+ — maximum domains, Alpine	Maximum number of EMISTP and PVST+ domains.	256
EMISTP & PVST+ — maximum domains, BlackDiamond	Maximum number of EMISTP and PVST+ domains.	512
EMISTP & PVST+ — maximum ports	Maximum number of EMISTP and PVST+ ports.	4096
EMISTP & PVST+ — maximum domains per port, Summit	Maximum number of EMISTP and PVST+ domains that can be configured per port.	128
EMISTP & PVST+ — maximum domains per port, Alpine	Maximum number of EMISTP and PVST+ domains that can be configured per port.	256
EMISTP & PVST+ — maximum domains per port, BlackDiamond	Maximum number of EMISTP and PVST+ domains that can be configured per port.	512
ESRP—maximum domains	Maximum number of ESRP domains for a single switch.	64
ESRP—maximum instances	Maximum number of ESRP supported VLANs for a single switch.	64
ESRP—maximum ESRP groups	Maximum number of ESRP groups within a broadcast domain.	4
ESRP—maximum ESRP groups with bi-directional rate shaping	Maximum number of ESRP groups within a broadcast domain when bi-directional rate shaping is enabled.	3
ESRP—maximum VLANs in a single ESRP domain – Summit, Alpine	Maximum number of VLANs that can be joined to a single ESRP instance through an ESRP domain. To obtain higher values see configuration notes.	256 recommended; 3000 max
ESRP—number of VLANs in a single ESRP domain, BlackDiamond	Maximum number of VLANs that can be joined to a single ESRP instance through an ESRP domain. To obtain higher values see configuration notes.	1024 recommended; 3000 max
ESRP—Route-track entries, Summit, Alpine, BlackDiamond	Maximum number of routes that can be tracked for each ESRP domain.	4
ESRP—maximum VLAN tracks	Maximum numbers of VLAN tracks per VLAN.	1

Table 7: Supported limits (continued)

Metric	Description	Limit
FDB—maximum ports for permanent entries	Maximum number of ports supported for permanent FDB entries.	2,000
FDB—maximum L2/L3 entries – BlackDiamond, Summit5i, Summit7i, Alpine 3804, Alpine 3808	Maximum number of MAC addresses/IP host routes for the MSM64i, Alpine 3808, and Summit7i.	262,144
FDB—maximum L2/L3 entries – Summit1i, Summit48i, Summit48si, Alpine 3802	Maximum number of MAC addresses/IP host routes for the Summit1i, Summit5i, and Summit48i.	131,072
Flow Redirection—maximum redirection rules	Maximum number of rules that can point to the same or separate groups of web cache servers.	64 (8 servers is the maximum)
Flow Redirection—maximum enumeration mode entries	Maximum number of active entries for enumeration mode rules. For example, one /16 rule can take all of the available entries.	64,000
Flow Redirection—maximum subnet mode entries	Maximum number of active entries for subnet mode rules. Each mask can have 1 entry.	64
IP ARP entries	Maximum number of IPARP entries.	20,480
IP ARP Static entries	Maximum number of permanent IP static ARP entries supported.	512
IP ARP Static Proxy entries	Maximum number of permanent IP ARP proxy entries.	512
IP Route Sharing Entries (ECMP)—static or OSPF	Maximum number of static or OSPF routes used in route sharing calculations.	12
IP Route Sharing Entries (ECMP)—IS-IS	Maximum number of IS-IS routes used in route sharing calculations.	8
IP Router interfaces	Maximum number of VLANs performing IP routing - excludes SubVLANs.	512
IP Static Routes	Maximum number of permanent IP routes.	1024
IPX Static Routes and Services (RIP and SAP)	Maximum number of static IPX RIP route and IPX SAP entries.	64 for each
IPX dynamic routes and services	Maximum recommended number of dynamically learned IPX RIP routes and SAP entries.	2000 for each
IPX Router interfaces	Maximum number of IPX router interfaces.	256
IPX Access control lists	Maximum number of Access Lists in which all rules utilize all available options.	worst case: 255
IS-IS—maximum routing interfaces	Maximum IS-IS routing interfaces.	255
IS-IS—maximum routes	Maximum IS-IS routes.	25,000
IS-IS—maximum adjacencies	Maximum IS-IS adjacencies per routing interface.	64
IS-IS—maximum domain summary addresses	Maximum IS-IS domain summary addresses.	32
IS-IS—maximum redistributed routes, regular metric	Maximum IS-IS redistributed routes using the regular metric.	20,000
IS-IS—maximum redistributed routes, wide metric	Maximum IS-IS redistributed routes using the wide metric.	30,000
IS-IS—maximum redistributed routes, both metrics	Maximum IS-IS redistributed routes using both metrics.	10,000
Logged Messages	Maximum number of messages logged locally on the system.	1000

Table 7: Supported limits (continued)

Metric	Description	Limit
MAC-based VLANs—MAC addresses	Maximum number of MAC addresses that can be downloaded to the switch when using MAC-based VLANs.	7000
MAC-based security	Maximum number of MAC-based security policies.	1024
Mirroring—mirrored ports	Maximum number of ports that can be mirrored to the mirror port.	8
Mirroring—number of VLANs	Maximum number of VLANs that can be mirrored to the mirror port.	8
NAT—maximum connections	Maximum number of simultaneous connections per switch.	256,000
NAT—maximum rules	Maximum number of rules per switch.	2048
NAT—maximum VLANs	Maximum number of inside or outside VLANs per switch.	The switch's limit
NetFlow—Filters	Maximum number of NetFlow filters in a switch.	128
NetFlow—Groups	Maximum number of NetFlow groups.	32
NetFlow—Hosts	Maximum number of NetFlow hosts.	8/group
OSPF areas	As an ABR, how many OSPF areas are supported within the same switch.	8
OSPF external routes—BlackDiamond, Summit7i, Alpine	Recommended maximum number of external routes contained in an OSPF LSDB without too many other types of OSPF routes.	130,000
OSPF inter- or intra-area routes—BlackDiamond, Summit7i, Alpine	Recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB without too many other types of OSPF routes, with one ABR in OSPF domain.	16,000
OSPF external routes—Summit1i, Summit5i, Summit48i, Summit48si	Recommended maximum number of external routes contained in an OSPF LSDB without too many other types of OSPF routes.	65,000
OSPF inter- or intra-area routes—Summit1i, Summit5i, Summit48i, Summit48si	Recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB without too many other types of OSPF routes, with one ABR in OSPF domain.	8,000
OSPF routers in a single area	Recommended maximum number of routers in a single OSPF area.	200
OSPF interfaces on a single router	Recommended maximum number of OSPF routed interfaces on a switch.	384
OSPF virtual links	Maximum number of OSPF virtual links supported.	32
OSPF adjacencies—Summit1i, Summit5i, Summit48i, Summit48si	Maximum number of OSPF adjacencies on a switch with 128 MB memory.	150
OSPF adjacencies—Summit7i, Alpine, BlackDiamond	Maximum number of OSPF adjacencies on a switch with 256 MB memory.	225
Policy Based Routing	Maximum number of policy based routes that can be stored on a switch.	64
RIP-learned routes	Maximum number of RIP routes supported without aggregation.	8000
RIP interfaces on a single router	Recommended maximum number of RIP routed interfaces on a switch.	384

Table 7: Supported limits (continued)

Metric	Description	Limit
Route Maps	Maximum number of route maps supported on a switch.	128
Route Map Entries	Maximum number of route map entries supported on a switch.	256
Route Map Statements	Maximum number of route map statements supported on a switch.	512
SLB—maximum number of simultaneous sessions	For Transparent and Translational and GoGo modes respectively.	500,000/500,000/ unlimited
SLB—maximum number of VIPs	For Transparent and Translational and GoGo modes respectively.	1000/1000/unlimited
SLB—maximum number of Pools	For Transparent and Translational (does not apply to GoGo mode)	256/256
SLB—maximum number of Nodes per Pool	For Transparent and Translational (does not apply to GoGo mode)	256/256
SLB—maximum number of physical servers per group	Applies to GoGo mode only; a group shares any number of common VIPs.	8
SSH2—number of sessions	Maximum number of simultaneous SSH2 sessions.	8
SNMP—Trap receivers	Maximum number of SNMP trap receiver stations supported.	16
Spanning Tree—maximum STPDs, Summit	Maximum number of Spanning Tree Domains.	128
Spanning Tree—maximum STPDs, Alpine	Maximum number of Spanning Tree Domains.	256
Spanning Tree—maximum STPDs, BlackDiamond	Maximum number of Spanning Tree Domains.	512
Spanning Tree—minimum STPDs	Minimum number of Spanning Tree Domains.	1
Spanning Tree—802.1d domains	Maximum number of 802.1d domains per port.	1
Spanning Tree—number of ports	Maximum number of ports that can participate in a single Spanning Tree Domain.	4096
Spanning Tree—minimum number of ports	Minimum number of ports that can participate in a single Spanning Tree Domain.	1
Spanning Tree—minimum number of domains/port	Minimum number of Spanning Tree Domains that can be configured per port.	1 for default VLAN, 0 for others
Spanning Tree—Spanning Tree modes	Maximum number of Spanning Tree modes per port.	3
Static MAC FDB entries—Summit, Alpine, BlackDiamond	Maximum number of permanent MAC entries configured into the FDB.	1024
Super-VLAN—number of ports & sub-VLANs	Maximum number of ports and sub-VLANs associated with each super-VLAN.	2550
Syslog servers	Maximum number of simultaneous syslog servers that are supported.	4
Telnet—number of sessions	Maximum number of simultaneous Telnet sessions.	8
UDP profiles	Number of profiles that can be created for UDP forwarding.	10
UDP profile entries	Number of entries within a single UDP profile.	16

Table 7: Supported limits (continued)

Metric	Description	Limit
VLANs—Summit, Alpine	Includes all VLANs plus sub VLANs, super VLANs, etc.	4095
VLANs—BlackDiamond 6816 fully populated	Includes all VLANs plus sub VLANs, super VLANs, etc.	681
VLANs—BlackDiamond 6816 with up to 7 I/O modules	Includes all VLANs plus sub VLANs, super VLANs, etc.	1776
VLANs—BlackDiamond	Includes all VLANs plus sub VLANs, super VLANs, etc.	4095
VLANs—maximum active protocol-sensitive filters	The number of simultaneously active protocol filters in the switch.	15
VRRP—maximum VRIDs	Maximum number of unique VRID numbers per switch.	4
VRRP—maximum VRIDs with bi-directional rate shaping	Maximum number of unique VRID numbers per switch when bi-directional rate shaping is enabled.	3
VRRP—maximum VRIDs/switch	Maximum number of VRIDs per switch.	64
VRRP—maximum VRIDs/VLAN	Maximum number of VRIDs per VLAN.	4
VRRP—maximum ping tracks	Maximum number of ping tracks per VLAN.	4
VRRP—maximum iproute tracks	Maximum number of iproute tracks per VLAN.	4
VRRP—maximum VLAN tracks	Maximum number of VLAN tracks per VLAN.	1

4 Clarifications, Known Behaviors, and Resolved Issues

This chapter describes items needing further clarification, behaviors that might not be intuitive, and issues that have been resolved since the last release. Numbers in parentheses are for internal reference and can be ignored.

This chapter contains the following sections:

- "Clarifications and Known Behaviors" on page 33
- "Issues Resolved in ExtremeWare 7.0.1b11" on page 60
- "Issues Resolved in ExtremeWare 7.0.0b68" on page 61
- "Issues Resolved in ExtremeWare 7.0.0b61" on page 62

Clarifications and Known Behaviors

Following are the clarifications and known behaviors in ExtremeWare 7.0.1. For changes made in previous releases, see the release notes specific to the release.

System Related – All Systems

CAUTION

In order for configuration changes to be retained through a switch power cycle or reboot, you must use the `save` *command.*

GBIC Type in the show ports configuration Command Output

ZX GBICs are displayed as LX-70 GBICs in the output of the `show ports configuration` command. This is a display issue only; the GBICs function correctly (PD2-131305301).

Do Not Telnet to Port 80 and Continuously Press Keys

If you telnet to the switch using port 80 and continuously press keys on your keyboard, the switch might eventually reboot (on a BlackDiamond switch, the master MSM might fail over to the slave) (PD2-129688312).

Smart Redundancy Enabled in Saved Configuration

Smart redundancy is always enabled in a saved configuration. To work around this, disable smart redundancy after downloading a configuration (PD2-128133503).

Microsoft Load Balancing

When using Microsoft load balancing, if you replace existing hardware and use the same IP address on the new hardware (thus associating the same IP address with a new MAC address), IP traffic through the IPFDB is not forwarded. To work around this, manually clear the IPFDB (PD2-124851229).

Telnet and the show ports Command

If you telnet to the switch and use the `show ports info detail` command, the line feeds might not be recognized, resulting in output lines overwriting previous lines (PD2-130127501).

The show configuration Output

After using the `unconfigure switch all` command, the `show configuration` output displays the VLAN *default* without any ports assigned. The ports still belong the to VLAN *default*, as the `show vlan` output correctly displays (PD2-128233941).

Configure Slots or VLANs Before Uploading a Configuration

If you do not configure any slots or VLANs, upload the configuration, reboot the switch, and download the configuration, all ports are deleted from the default VLANs (PD2-110787427). The workaround is to configure slots or create a VLAN before you upload the configuration.

LACP not Supported

Contrary to the information in the *ExtremeWare 7.0 Software User Guide* and *ExtremeWare 7.0 Command Reference Guide*, LACP is not supported.

Upgrading to ExtremeWare 7.0 and Bi-Directional Rate Shaping

When you directly upgrade from ExtremeWare 6.2.2 to ExtremeWare 7.0, bi-directional rate shaping does not work if the loopback ports were in autonegotiation mode. This behavior is not displayed by 10/100Base-T or Gigabit fiber ports. A workaround is to remove and re-add the loopback ports to the VLAN (PD2-107820904).

Upgrading to ExtremeWare 7.0 and Debug-Trace

When you directly upgrade from ExtremeWare 6.2.2 to ExtremeWare 7.0, the debug-trace configuration might change. Verify the debug-trace configuration, if any, after upgrading. Use the `show debug-trace` command to display the configuration. You can either re-configure manually, or download the ExtremeWare 6.2.2 configuration instead of doing a direct upgrade (PD2-106733988).

Upgrading to ExtremeWare 7.0 and OSPF

If you upgrade directly from ExtremeWare 6.2.2 to ExtremeWare 7.0, the OSPF metric for 10 Gigabit interfaces is incorrect. A workaround is to manually configure the OSPF metrics, or to upload the configuration before upgrading and then download the ExtremeWare 6.2.2 configuration (PD2-108161623).

Routing Traffic Through the MGMT Port

Routing entries with a next hop in the management interface subnet are not removed from the routing table based on the MGMT port state (PD2-104430127).

Configuring the Timezone

After configuring the timezone, a soft reboot can cause the switch to boot into minimum mode (PD2-109830723).

Blank Space in show port info detail Command Output

The output of the `show port info detail` command contains several blank pages. The output still contains all of the requested information (PD2-107800978).

Using an ExtremeWare 7.0 Configuration with an Earlier Image

If you are using an ExtremeWare 7.0 configuration and attempt to use an earlier image, the switch prompts you for confirmation (because this combination is not recommended). If you answer "n" at the prompt, you receive the following error message:

```
Error: bad image.
```

You can safely ignore this message (PD2-110983501).

Console Response with a Large Number of ARP Entries

Console reponse is slow when the switch is learning 10,000 or more ARP entries. This does not affect performance. Console response returns to normal when the entries are learned (PD2-104103941).

Configuring 1000Base-T Ports for 10,000 Mbps

The switch erroneously allows you to configure a 1000Base-T port to 10,000 Mbps. 1000Base-T ports do not support 10,000 Mbps (PD2-108463706).

The show log chronological Command

When the syslog contains more than 1,000 lines, the `show log chronological` command displays nothing. However, the command `show log` displays correctly (PD2-104062736).

BOOTP-Dependent Routes in Downloaded Configuration not Created

Static and default routes that depend on a BOOTP IP address/subnet are not created when you download a configuration (PD2-86888351).

Enable Flow Statistics Ping-Checking

Flow statistics requires ping-checking to ensure that the flow-collectors are operating. Flow statistics ping-checking is disabled by default. You must enable flow statistics ping-checking to enable flow statistics (PD2-110325062).

UDP Echo Transmit Rate

The UDP Echo utility is designed to verify network connectivity. Transmit rates of 10 pps suffice for this function. UDP Echo rates of 20 pps should be sufficient. Do not set your UDP Echo rate higher than 100 pps, as the switch does not send replies faster than that rate (1-FAO89).

The disable learning Command and Flooding

The disable learning command does not remove the port from the security flood list. Thus, you cannot disable flooding when learning is disabled (PD2-73199618).

Port Mirroring

When a multicast packet egresses from a port, two copies of the packet are sent to the mirror port. This does not affect network traffic in any way, as the duplicate packets are sent only to the mirror port. This does affect accounting and RMON statistics (1-DQK86).

Port mirroring is not supported across BlackDiamond modules (PD2-89313413).

Port mirroring is not supported with CPU-generated traffic (1-64H4J).

Setting Auto-negotiation Off on a Gigabit Port

When connecting to a device that does not support 802.3z auto-negotiation, turn off auto-negotiation for the switch port to which it is connecting. Although a gigabit port only runs at full duplex and at gigabit speed, the command to turn auto-negotiation off must still specify duplex. For example:

```
config port 4 auto off duplex full
```

will turn auto-negotiation off if port 4 is a gigabit port.

Enabled IdleTimeouts and Console Connections

If the IdleTimeout feature is enabled, and a telnet session times out, a subsequent telnet to the switch will be successful but existing direct serial console connections will pause or hang. If the subsequent telnet session is terminated, the console port will resume normal function and subsequent telnet sessions will work correctly (5094).

User Accounts

User account usernames and passwords can have a maximum of 30 characters (PD2-101617708).

TFTP Download of Configuration Files

When using TFTP to download a configuration file and selecting “no” for the switch reboot request, rebooting the switch at a later time will display a message that the configuration file has been corrupted. The user will be prompted to reboot the switch with factory default parameters. If an immediate reboot is performed after the download configuration command, the configuration file will be initiated correctly (12413).

Port Tag Limitation

There is an absolute limit of 3552 port tags available in a system. The usage of these port tags depends on a combination of factors:

- Installed ATM, MPLS, ARM, and PoS modules
- Mirroring
- IPX routing
- Static FDB entries

If the switch reaches the limit of available port tags, the following messages appear in the syslog:

```
<WARN:HW> tNetTask: Reached maximum otp index allocation
<WARN:HW> tBGTask: Reached maximum otp index allocation
```

If this occurs, you must compromise some features (for example, mirroring) in order to expand your use of other functionality. (1-E5U7Y).

BlackDiamond

MPLS Hello Packets

MPLS hello packets are sent every 5 seconds, regardless of the configured value (PD2-131214401).

Slot Failure Messages During a Broadcast Storm

If you have more than 15 Gigabit Ethernet links between two chassis, all in the same VLAN and generating a broadcast storm, the system health check records slot failures in the log. When the broadcast storm stops, the log messages also stop (PD2-117946811).

Hot-Inserting an MSM Disrupts MPLS and ARM Modules

If you hot-insert a second MSM, some IP traffic being forwarded through MPLS and ARM modules is halted. You must reboot the modules to restart the lost traffic flow (PD2-130167901).

No Image Information Reported to SNMP with One MSM

If you only install an MSM in slot B of a BlackDiamond 6804, BlackDiamond 6808, or BlackDiamond 6816, no primary or secondary image information is reported to your SNMP NMS (PD2-129612901).

MPLS and CPU DoS Protect

If you enable CPU DoS protect on an BlackDiamond with an MPLS module, ICMP traffic is blocked. To work around this, disable MPLS before you enable CPU DoS protect (PD2-119097601).

Duplicate Precedence Rules

If you create an ACL rule with the same precedence as an existing rule, an error message warns you of the duplication. However, the rule is still created. You must delete the rule with the duplicate precedence and recreate it with a unique precedence (PD2-116540055).

BlackDiamond 6816 MSM C and D Diagnostics Messages not in Syslog

If you run diagnostics on an MSM in slot C or D of a BlackDiamond 6816, messages are not recorded in the syslog. To view the diagnostics messages, use the `show diagnostics` command (PD2-118049501).

Synchronize a Newly Installed MSM64i

When you add a slave MSM64i, you are not prompted to synchronize. If not synchronized, the slave MSM64i uses its image and the master MSM64i configuration. This image/configuration mismatch will likely cause the switch to operate differently after a failover, thereby defeating the purpose of the dual MSM64i's. Be sure the MSM64i's are synchronized (PD2-101615201).

Disabling CLI Paging from the Slave MSM64i

Enabling or disabling CLI paging from the slave MSM64i has no affect on the master MSM64i paging configuration (PD2-104377501).

Limited Commands Mode and the reboot Command

In limited commands mode, the `reboot` command does not reboot the MSM64i; instead the command causes the MSM64i to fail over (PD2-107053801).

The unconfig switch all Command

If you use the `unconfig switch all` command and immediately use the `config default vlan delete port all` command, the switch reboots (PD2-105474401). To avoid this situation, after you unconfigure the switch, wait for the switch to completely reboot before you delete the ports.

Dynamic Memory Scanning and Mapping Module Support

BlackDiamond I/O module memory scanning and mapping support is listed in Table 8.

Table 8: Memory scanning and mapping support in BlackDiamond modules

Module	Memory Scanning and Mapping
F32Fi	Yes
F48Ti	Yes
F96Ti	Yes
G12SXi	Yes
G8Ti	Yes
G8Xi	Yes
WDMi	Yes
MSM64i	Yes

Extended Diagnostics

The `run diagnostics extended` command can cause the following messages to appear in the log. These messages are expected and indicate that the system is currently busy running the user initiated diagnostics (10800). This does not occur with the `run diagnostics normal` command.

```
<CRIT:SYST> task tdiagTask cpu utilization is 98% PC: 806266e8

<INFO:SYST> task tdiagTask cpu utilization is 98% PC: 806266e8
```

BlackDiamond 6816 MIB Value for Input Power Voltage

On the BlackDiamond 6816, the extremeInputPowerVoltage attribute in extremeSytemCommonInfo is shown as “0” and the extremePowerSupplyInputVoltage in the extremePowerSupplyTable is shown as “unknown.” These values cannot be obtained from the switch (1-841J1).

Backplane Traffic

On the BlackDiamond switch, all backplane traffic is tagged. As a result, for cross-module traffic traversing the switch, dot1P QoS has the highest priority on egress (1-CPL8B).

QoS

If you configure QoS on an untagged ingress port, the dot1p bit of a packet leaving a tagged port on a different module is always replaced, even though dot1p replacement is disabled (1-E2UX2, 1-5I3VA).

Alpine

Alpine 3802 show switch Output Shows Incorrect PSU Placement

The output of the `show switch` command shows PSU A on top and PSU B on the bottom in an Alpine 3802 chassis (PD2-129291301).

System Health Check Events Might Not Be Logged

Messages for system health check events might not be logged due to the software not being intialized (PD2-129795601).

Configuring Two Multilink groups on One T1 or E1 Module

If you configure two multilink groups to use the same T1 or E1 module, multilink throughput is degraded slightly. A workaround is to disable WAN QoS (PD2-117966118).

Limited Commands Mode

When in limited commands mode, the slot status LED remains orange, though the link is taken down (PD2-99107226).

T1 and E1 Error Message

A message similar to the following:

```
12/06/2002 11:58.28 <CRIT:KERN> Restarted fifo on slot 2
```

might appear in the log for T1 and E1 slots during the initialization of the T1 or E1 modules. These messages are not critical and do not affect the operation of the modules (PD2-110059501).

VDSL Modules in a Half-Duplex Link

A VDSL CPE operating in a half-duplex link can lock up when used with a hub and running wire-rate randomized traffic. This is a hardware limitation. A restart of the VDSL port will recover, but if the traffic continues at wire-rate and is randomized, then the problem will reoccur (PD2-71538118).

Summit

Health Check Error Messages

Error messages from the system health check display the incorrect location (PD2-110132842).

Limited Commands Mode

When in limited commands mode, links remain active (PD2-99220424).

Summit48i Redundant PHY

When the primary port of a redundant pair is disabled and the link removed, the LED for that port continues to flash indicating it has a link and is disabled (9239).

Summit48i Single Fiber Signal Loss

The Summit48i is currently not able to detect a single fiber strand signal loss due to the hardware based Auto Negotiation parameters (10995).

SNMP Results for Power Sources

The inputPower MIB is unable to differentiate between 110 VAC and 220 VAC input on the Summit series switches when accessing this MIB attribute through SNMP (10870).

Summit48si MIB value for Input Power Voltage

On the Summit48si, the extremeInputPowerVoltage attribute in extremeSytemCommonInfo is shown as “0” and the extremePowerSupplyInputVoltage in the extremePowerSupplyTable is shown as “unknown.” These values cannot be obtained from the switch (1-841J1).

Command Line Interface (CLI)

Only US Character Set Supported

The CLI supports only the US character set (2-H1OQC).

The show iproute Command

The `show iproute` display has a special flag for routes that are active and in use, these routes are preceded by an “*” in the route table. If there are multiple routes to the same destination network, the “*” will indicate which route is the most preferable route.

The “Use” and “M-Use” fields in the route table indicate the number of times the software routing module is using the route table entry for packet forwarding decisions. The “Use” field indicates a count for unicast routing while the “M-Use” field indicates a count for multicast routing. If the use count is going up in an unexpected manner, this indicates that the software is making route decisions and can be something to investigate further.

Serial and Telnet Configuration

Be sure you have specified VT-100 terminal emulation within the application you are using (2125, 2126).

Be sure to maximize the telnet screen in order for automatically updating screens to display correctly (2380).

Displaying Management Port with show port config

The `show port config` command will only display the "mgmt" port configuration information if the "mgmt" port is explicitly defined in the command - i.e., `show ports mgmt config` (8604).

Auto Negotiation and 1000BaseT Ports

Note that per specification, auto-negotiation cannot be disabled on 1000Base-T ports (8867).

Switching and VLANs

Saving ip-mtu Settings

Dynamic TLS (Martini TLS) checks the MTU received from its peer in order for TLS to come to the established state. It compares against the egress VLAN's IP-MTU. If the egress VLAN does not have an IP address defined, any non-default ip-mtu setting will not be saved through a switch reboot (PD2-64084527).

FDB

Static FDB Entries and Rate-Shaping. If you create a static FDB entry on a port configured for rate-shaping, the static entry incorrectly ages out. Static entries should not age out (PD2-97150551).

FDB Aging Timer. In ExtremeWare 6.2.0, the default value of the FDB aging timer was set to 1800 seconds on a newly configured ExtremeWare 6.2.0 switch. In ExtremeWare 6.2.1 the default value has been changed back to 300 seconds. However, when upgrading from ExtremeWare 6.2.0 to ExtremeWare 6.2.1, the default value will remain and 1800 seconds. For upgrades from ExtremeWare 6.1.9 (or earlier) the default value will remain 300 seconds. The FDB aging time can still be set to all previous values (1-85QD3).

Configure Less Than 400 Ports in a VLAN

If you use the `clear slot` command (which flushes the FDB) when there are 256,000 or more FDB entries, the watchdog timer can cause the switch to reboot. To avoid this, configure less than 400 ports in a VLAN (PD2-90223209).

Cannot Delete "mgmt-1" VLAN

A VLAN created with the name "mgmt-1" cannot be deleted (1-EEUPE).

VLAN priority and STP, EDP

STP and EDP (thus ESRP and EAPS) do not transmit packets in the queue specified by the VLAN priority (1-5HOZ9).

Default Routes or Static Routes

If you define a default or static route, and then delete the VLAN on the subnet associated with the default route, the default route entry remains although it is invalid. You must manually delete the configured route.

If you define multiple default routes, the one with the lowest metric is used. If there are multiple routes with the same lowest metric, the switch picks one of the routes.

Configuring a Protocol Filter with 'ffff'

Creating a protocol filter with LLC of value 0xffff and assigning it to a VLAN will not work, and should not be attempted. Doing so will assign all traffic to the VLANs containing that filter and will prevent VLANs with no protocol filter to function properly. If configured and saved, you must issue an `unconfigure switch all` to restore normal operation (2644, 4935).

Deleting Protocols from a VLAN

Adding a protocol to a VLAN may cause an EPC if the protocol was added to the VLAN, deleted from the VLAN, recreated by the user, and re-added to the VLAN (6128).

MAC Based VLANs and DHCP Relay

MAC based VLAN configurations should not be used in conjunction with DHCP. Currently, a host which enters a MAC-based VLAN will not be able to use DHCP to obtain an IP address.

VLAN to VLAN Access Profiles

VLAN to VLAN access profiles are no longer supported on the BlackDiamond switch in ExtremeWare 6.0 or higher (7022).

Load Sharing

Load Sharing and Software Redundant Ports. If you configure software redundant ports with load sharing, saved configurations do not load properly via TFTP. The configuration file executes the software redundant port configuration before the load sharing configuration, but you must configure load sharing first. To avoid this, edit the configuration file so that the load sharing configuration is executed first (PD2-130597269).

Autonegotiation. Load sharing ports must be configured with autonegotiation set to on. Load sharing ports will not transmit traffic correctly using any other setting (PD2-64617405).

Round Robin Load Sharing. If a port in a round robin load share group is removed, the traffic that was being transmitted on that link will be distributed on only 1 of the other active load share links in the round robin group. The traffic is not distributed evenly between the remaining ports (6977).

Port Based Load Sharing on Summit7i. Port-based load sharing on the Summit7i requires ingress ports to be on the same side of the switch (ports 1 - 4, 9 - 12, 17 - 20, and 25 - 28 on the left, ports 5 - 8, 13 - 16, and 21 - 24 on the right) as the 8 ports in the load share group for all ports in the load share group to transmit/receive traffic (6975).

Alpine and Cross Module Load Sharing. The I/O module configured to contain the "master" port must be physically present in a cross-module load sharing group for the system to pass traffic (8589, PD2-119098401).

Load Sharing and Specific Ports in a Load Share Group. Due to the load sharing algorithm used for round robin load sharing, when using 3, 5, 6 or 7 ports in a load share group packet loss will be observed when sending wire-speed traffic across the load share group. This occurs because some ports will be selected to transmit more packets than other ports resulting in bandwidth over-subscription and subsequent packet loss. This only occurs with round-robin load sharing configurations (10311).

Load Sharing, Software Redundant Ports, and Smart Redundancy. The smart redundancy feature is not supported when using software redundant ports and load sharing (12431).

Disabling Load Sharing if the Master is Down Generates Error. If the load sharing master link goes down, and you disable load sharing, the switch generates a ptag error message (PD2-129379272).

Spanning Tree

Deleting a Port From the STP Domain. If you delete a port from the STP domain and save the configuration, that change is not saved. You must either delete the port again after rebooting or edit the configuration file to delete the port from the configuration (PD2-130809831).

Configuring a VLAN from Vista. If you create an STPD using ExtremeWare 6.1.9 (or earlier), add a VLAN, save the configuration, upgrade to ExtremeWare 6.2.2b68 (or later), and save the configuration, you receive the following error message when you try to modify the VLAN from Vista:

```
ERROR: Cannot assign bridge to stpd! HINT: If a port is part of multiple vlans, the
vlans must be in the same Spanning Tree domain.
```

To work around this problem, make configuration changes from the CLI (PD2-118450190).

STP and VLAN Tagging. VLAN tagging is not supported with 802.1d Spanning Tree (STP) BPDUs. Therefore, all BPDUs in a 802.1d STP domain are untagged. However, Extreme Multiple Instance Spanning Tree (EMISTP) and Per-VLAN Spanning Tree (PVST+) do support VLAN tagging of BPDUs.

EMISTP and Ingress Rate Shaping. If a loop exists in your network, but STP is not enabled but Ingress Rate Shaping is, the switches appear to hang and are rebooted by the watch-dog timer. A similar situation exists if a loop is covered by STP on both sides and is disabled on one side; normally the other switch immediately blocks the right port(s), but when Ingress Rate Shaping is present, both switches appear to hang and are rebooted by the watch-dog timer (1-5E9R1).

Upgrading From an ExtremeWare 6.1.9 (or earlier) Configuration. After downloading an ExtremeWare 6.1.9 (or earlier) configuration to an ExtremeWare 6.2.0 (or later) image, a port belonging to a non-default VLAN will generate the "Stpd s0, Port 1:1 does not exist" error message because that VLAN does not belong to domain s0 by default (1-BMP5D).

MAC Security

The source FDB address configuration will not discard ICMP packets (16340).

Mirroring

Mirroring IP Multicast Traffic. Due to IGMP Snooping capabilities, Multicast traffic may cease to be seen on a “mirror port”. If you issue a “restart” command for the mirror port or remove and reinsert the port connection, multicast traffic will resume for the IGMP Host time out period (260 sec.) (3534).

Mirroring and Flooding. When a mirrored port is configured, the forwarding database for items being mirrored (e.g. ports or VLANs) are automatically cleared from the forwarding database if the link status on the mirror port changes. This will result in some temporary flooding until the normal learning process completes. Thus, removing or inserting a probe device into the mirror port may appear to cause flooding, however this is expected behavior (5128).

QoS

The qosprofile Accepts a Value Greater than 100%

The `maxbw` parameter in the `configure qosprofile` command incorrectly accepts values greater than 100%; however, the maximum bandwidth is still 100% (PD2-123662004).

Re-Ordering Access List Precedence Numbers

When you add a new ACL rule with a precedence number, the switch re-orders existing rules with lower precedence numbers to make room for the new rule. If, during this re-ordering, two rules have a precedence number difference greater than one, the switch generates an error message similar to the following:

```
<WARN:KERN> Access rule does not exist
```

You can safely ignore this error message (1-FAO8M).

Access List FDB Entries not Cleaned Up

If you delete an access list with the “f” flag (flow rule), the associated FDB entries might not be cleared (PD2-110082518).

Access Lists Using the IP Deny Any Rule

When using an access control list with an IP deny any rule, all ICMP traffic will be blocked within a VLAN (Layer 2). If using an access list with an IP deny any rule across VLANs (Layer 3), ICMP traffic will not be blocked.

Access Lists and IP Fragmentation

When using IP fragmentation, since the TCP header is treated as data and only the IP header information is being replicated in each packet, access-lists that apply to that flow will not apply as the TCP/USP port information is not included after the first fragment (for subsequent fragments).

QoS Configuration Bandwidth Parameters

Minimum and maximum percentage parameters for a specific port on the default VLAN will not be saved across reboots. The configuration change will be applied when configured. This issue only occurs on the BlackDiamond (15500).

Access List Precedence Intervals

Access lists with large intervals (greater than 10) between precedence values now perform better. Previously, configuring access lists using large intervals (greater than 10) between precedence values could result in several-minute delays for each add transaction. We still recommend that you configure ACL precedence with an interval value of less than 5 between each rule. This configuration avoids any adverse performance issues such as very long delays between add transactions and loss of access to configuration sessions (1-B6F48, 15717, 15718).

Creating Access Lists from Multiple Sessions

When creating or modifying access control lists, please ensure that no other administrator sessions are attempting to create or modify the system access control lists simultaneously. This may result in data corruption (1-579HD).

QoS and dot1p

If you configure VLAN QoS to a higher precedence than dot1p QoS using QoStype priority, egress traffic will go out through Q0 (1-CH3MD).

5,120 Access Lists and SNMP

Although you can configure up to 5,120 ACLs, SNMP only recognizes 1,280. Deleting an ACL that is not recognized by SNMP generates the following error (PD2-64880917):

```
<WARN:SNMP> SNMP IPQOS Could not find entry instance 5083 to delete
```

Monitoring QoS and the show port qos Command

When monitoring QoS, do not use the `show port qos` and `enable qosmonitor` commands on the same port at the same time. These commands in conjunction lock the console session. However, the syslog does capture the output (PD2-64202681, PD2-80836531).

Ingress QoS

Ingress QoS Not Supported on Other Modules

Ingress QoS is only supported on "3" series modules. Though you can configure ingress QoS on other modules, the feature is not supported and the configuration has no effect (PD2-129625008).

The show ports ingress stats Command Truncates

The `Tx Xoff` column in the `show ports ingress stats` command output truncates values to seven characters (PD2-130148001).

Bi-Directional Rate Shaping

Locking and Unlocking Learning

If you configure a rate shaping port to lock learning and unlock learning, the loopback FDB is not flushed. This causes traffic destined for the port to be flooded. You must manually flush the FDB using the `clear fdb` command (PD2-124568416).

Loopback Port Must be on Same Module

The loopback port must be on the same module as the rate shaped ports. Though you can configure a loopback port on another module, this is still not a supported configuration (PD2-124299901).

1000Base-T Ports as Loopback Ports

If the loopback port for bi-directional rate shaping configurations is configured on 1000Base-T ports, the speed of that port cannot be changed from 1000 Mbps to 100 Mbps as the bandwidth settings will not be accurate when configured in 100 Mbps mode.

Changing the Configuration of a Loopback Port

If you change the configuration (speed, duplex setting, etc.) of a loopback port, you must either save the configuration and reboot the switch, or delete the port from the VLAN and add it back (PD2-127582534).

EAPS

WAN Modules Not Currently Supported with EAPS

Do not use WAN modules with EAPS (PD2-120015201).

Do Not Configure a Hello Time of 0

Though the minimum hello time is 1, the switch accepts a hello time of 0. Do not configure the hello time to 0, as this effectively disables EAPS (PD2-119139425).

A Large EAPS Configuration with a Link Transition

If you configure a single EAPS ring with 64 domains and more than 3,000 VLANs, a link transition could cause a 300 second traffic outage. To work around this, delete the old FDB entries using the `clear fdb` command (PD2-119139401).

Changing the Protected VLAN Tag

Do not change the protected VLAN tag if EAPS is configured and enabled. Doing so can create a loop in the network. First disable EAPS, then make changes (PD2-121610287).

EAPS Performance Statistics

Table 9 lists the EAPS performance statistics for a single EAPS domain with the default filter.

Table 9: EAPS performance statistics with the default filter

Protected VLANs	Link Down Convergence Upstream (ms)	Link Down Convergence Downstream (ms)	Link Up Convergence Upstream (ms)	Link Up Convergence Downstream (ms)
1	219	219	100	101
500	220	219	128	126
1,000	220	220	158	150
4,000	262	266	289	244

Table 10 lists the EAPS performance statistics for a single EAPS domain with no filters.

Table 10: EAPS performance statistics with no filters

Protected VLANs	Link Down Convergence Upstream (ms)	Link Down Convergence Downstream (ms)	Link Up Convergence Upstream (ms)	Link Up Convergence Downstream (ms)
1	114	114	100	101
500	114	114	129	127
1,000	115	115	158	150
4,000	165	170	340	295

Table 11 lists the EAPS performance statistics for a single EAPS domain with a single protected VLAN and varying FDB sizes.

Table 11: EAPS performance statistics with varying FDB sizes

FDB Entries	Link Down Convergence Upstream (ms)	Link Down Convergence Downstream (ms)	Link Up Convergence Upstream (ms)	Link Up Convergence Downstream (ms)
2,000	132	127	117	120
10,000	190	162	176	176
50,000	478	341	472	476
100,000	829	554	828	688

EAPS and STP or EMISTP

If you have an EAPS master domain and an EAPS transit domain on a single switch, only add the STP or EMISTP VLAN to the EAPS master domain. On switches running only an EAPS transit domain, add the STP or EMISTP VLAN to both EAPS domains.

If you configure two different EAPS master domains on the same switch, use two separate STP or EMISTP VLANs and two separate STP or EMISTP domains (PD2-72446883).

EAPS Secondary Port Recovery

The EAPS secondary port does not recover if the following events occur in the following order (1-FY31X):

1 The EAPS ring fails, due to a Hello timeout or a link failure.

2 The EAPS master node secondary port fails or is disabled.

3 The EAPS master node secondary port recovers or is re-enabled. The port incorrectly blocks incoming traffic even though it is enabled.

ESRP and EAPS Secondary Port

Configuring ESRP Host Attach on an EAPS secondary port causes a broadcast storm (1-B1O4L).

Incorrect show vlan Output

The `show vlan` output incorrectly lists the EAPS secondary port as active with an asterisk (*). The number of active ports is correctly displayed (PD2-59142420).

ESRP

Configure a Neighbor Timeout Less than 6 Times Hello Timer

If you configure the neighbor timeout to greater than six times the hello timer, and the link between the master and the slave goes down, the slave might not immediately flush the FDB table. To avoid this, configure a neighbor timeout less than six times the hello timer. To correct this situation, manually clear the FDB (PD2-124371801).

Transition Incorrectly Logged

If you change the priority of the ESRP master to 255, in rare situations it might change to slave, transition back to master, then finally transition to slave (PD2-129379243).

Dual Master Recovery Not Logged

When two switches recover from a dual-master situation, in rare situations the new master might not log the state change (PD2-111406501).

A Flapping Redundant Link Might Cause ESRP to Fail Over

A flapping redundant link might cause the port counter to increase its count on the neigbor’s side, increasing the neighbor's port count. This could cause an ESRP state transition. To avoid this, disable smart redundancy (PD2-111264407).

ESRP and Ingress Rate Shaping

Do not use ingress rate shaping on an ESRP-enabled port (PD2-107800933).

ESRP and Protocol-Based VLANs

ESRP-aware switches cannot connect to an ESRP switch through a port configured for a protocol-sensitive VLAN using untagged traffic (PD2-99007701).

ESRP and Load Sharing

If you enable load sharing on ports that belong to more than 200 VLANs, the switch reboots. To avoid this, first enable load sharing, then add the ports to the VLANs (PD2-99259801).

When using load sharing with the ESRP host attach or don't count features, configure *all* ports in the same load-sharing group as host attach ports or don’t-count ports (PD2-97342427, PD2-106782876).

Hot-Swapping a Module with 5,000 ACLs

Hot-swapping a module on a switch that has 5,000 or more ACLs configured can cause an ESRP state change (PD2-107800998, PD2-103938301). To avoid the state change, configure the neighbor timeout value to 12 seconds.

Traffic Convergence Time

Traffic convergence after a link failure can take as long as 5 seconds with 2,000 VLANs and 256,000 FDB entries. This delay can cause ESRP state changes as traffic converges (PD2-89915300).

ESRP PDUs on Ports

ESRP PDUs received on ports that do not belong to any VLAN are processed as valid ESRP PDUs and can trigger state changes (PD2-89481346). To avoid this, assign all ports to valid VLANs.

Multiple ESRP VLANs

If multiple ESRP VLANs share a host port, each VLAN must be in a different ESRP group.

ESRP Interoperability

We recommend that all switches participating directly in ESRP be running the same version of ExtremeWare. If you must mix ExtremeWare versions, do not use any of the ESRP features new to this release.

Mixing Clients and Routers on an ESRP-Enabled VLAN

Typically, ESRP is not be enabled on a VLAN that is also expected to exchange routes with other non-ESRP routers (e.g.: routers using RIP or OSPF). ESRP is intended and designed as a Layer 2 or Layer 3 redundancy method for clients with a single default route. ESRP’s fail-over operation may interfere with normal routing protocol communication if an ESRP-enabled VLAN contains other routers not using ESRP (4874).

ESRP and Bi-Directional Rate Shaping

When a single ESRP VLAN is configured with bi-directional rate shaping ports and no direct physical connection to the 2nd ESRP router, the ESRP slave router flips back and forth to Master state. If a second rate-shaped VLAN or a direct link between the 2 ESRP routers exists, this will not occur (10739).

When ESRP and bi-directional rate shaping are configured simultaneously on the same switch, rate shaping traffic to the ESRP MAC address will not take effect until the switch is rebooted (13583).

VRRP

The show tech-support Command Through Telnet

In a configuration with more than 20 VLANs, if you use the `show tech-support` command on the backup switch through a telnet connection, the backup transitions to master and back. To avoid this, use the `show tech-support` command only through a direct console connection (PD2-128764506).

Increase Advertisement Interval When CPU is Busy

An extremely busy CPU can create a short dual master situation. To avoid this, increase the advertisement interval (PD2-130779223).

Backup Transition Creates Duplicate Packets

A VRRP transition from backup to master might cause duplicate packets to be transmitted for a short period of time (PD2-129379226).

Changing the Advertisement Interval

If you configure a new advertisement interval and then reconfigure the interval back to the default, VRRP elects a new master but keeps the existing master, resulting in two master VRRP VLANs. To avoid this, disable and re-enable VRRP (PD2-127681301).

Changing the Priority

If you configure the VRRP master priority to 0 (releasing it as the virtual router) and then configure the priority to 255, the master is not released even though a new master is elected. This results in two VRRP masters. To avoid this, disable and re-enable VRRP (PD2-127681312).

The track-diagnostic and track-environment Features Not Supported

The track-diagnostic and track-environment features are not currently supported with VRRP (PD2-127681344).

IP Unicast Routing

Deleting a Static Entry Using SNMP

If you delete a static IPARP entry using SNMP, the line in the configuration creating that entry is not deleted. Thus, if you reboot, the static entry is again created. To work around this, either edit the configuration or delete static IPARP entries through a direct connection to the switch (PD2-130505418).

The show iproute Output

The output of the show iproute command displays only the first 8 characters of the VLAN name (PD2-128392829).

Traffic Crosses Layer 3 Boundary

If ingress and egress VLANs do not share a port, layer 3 traffic with a broadcast MAC and unicast IP address is incorrectly forwarded to the default route across a layer 3 boundary (PD2-119375325).

Moving a sub-VLAN Client

When a client is moved from one sub-VLAN to another, the client may not be able to ping or communicate through the super-VLAN until the client has cleared its IP ARP cache for the default router or the switch has that IP ARP cache entry cleared (4977).

No Static ARP Entries

The use of Static ARP entries associated with superVLANs or sub-VLANs is not supported in this release (5106).

VLAN Aggregation and ESRP

A sub-VLAN should not be configured to run ESRP. The system will allow you to enable ESRP on a VLAN and then designate the VLAN as a sub-VLAN, but this is not a supported configuration (5193).

ARP Entry Age

The age of ARP entries changes to a large value when system time is changed (1-E7FIV).

Multinetting and Client Default Gateways

It is critical that clients attached to multinetted segments have their default gateways correspond to the same subnet as their IP addresses and that subnet masks be configured correctly. Not doing so will result in slow performance of the switch (4938).

Multinetting and the Show VLAN Stats Command

The `show vlan stats <vlan_name>` command is not supported on multinetted VLANs.

Multinetting and VRRP

Multinetting is not supported with VRRP.

RIP Routing

RIP V2 Authentication

The authentication feature of RIPv2 is not supported.

RIP in Conjunction with other Routing Protocols

It is recommended that RIP be enabled only on routers running with less than 10,000 routes from other routing protocols, such as BGP or OSPF.

OSPF

Default Route Entries in the IP FDB

After a link transition, entries created by the OSPF originated default route are still in the IP FDB (PD2-109830730, PD2-109830723).

Disable OSPF Before Adding or Removing External Area Filters

If you configure an OSPF area external filter on an ABR, and the filter is set to exclude routes that have already been learned, an OSPF failure occurs. A workaround is to disable OSPF before adding or removing OSPF external area filters (PD2-105170634).

BGP

Multi Exist Descriminator Not Compared

If a route is received from the same AS via EBGP and the IBGP peer, the switch does not compare the multi exist descriminator. To avoid this, use the `enable bgp always-compare-med` command (PD2-126767407).

Route Dropped if Switch's AS is First AS in Path

If the switch receives a route from an IBGP peer and the first AS number in the AS path sequence is the switch's own AS number, the route is dropped as a loop. To avoid this, do not prepend the switch's AS number to the AS path (PD2-126767401).

BGP Set Community Inadvertantly Advertised

The BGP Set Community NO_EXPORT_SUBCONFED is inadvertantly advertised to EBGP peers (PD2-120403214).

Do Not Use configure access-profile Command to Set Community

Do not use the `configure access-profile add` command to set the BGP community , as the command does not correctly set the value. Use the `configure route-map add` command instead (PD2-129638011).

Best Routes

If a new best route comes from an I-BGP peer, an older best route that comes from E-BGP won't be withdrawn (PD2-108750310).

BGP Loops

If a switch detects a BGP route loop, it tears down the link to the neighbor that forwarded the route. To avoid this, disable and re-enable BGP (PD2-99209507).

Redistributing BGP Routes to OSPF

Redistributing 70,000 BGP routes into OSPF depletes the system resources. You must reboot the system (PD2-74932501).

Removed encrypted Option from enable bgp neighbor password Command

The `enable bgp neighbor password` command no longer has the option to encrypt the password (PD2-101778801).

IP Multicast Routing

Use the always Parameter to Guarantee Advertisement

The `enable rip originate-default` command does not always advertise the default RIP route to peers. To guarantee that the default RIP route is advertised, use the `always` parameter (PD2-124368763).

Cisco Interoperation

For proper Cisco interoperation, use Cisco IOS version 11.3 or better, which supports PIM 2.0. Cisco customer support also recommends using PIM in favor of DVMRP whenever possible on Cisco routers (4669).

Traffic Rate Exceeding Last Hop Threshold

When the traffic rate exceeds the configured last hop threshold, the last hop does not initialize; but if the sending traffic rate is set to 50 Kbps, it switches to STP correctly (1-57NMY).

IPX Routing

Tuning

In larger environments, it is helpful to increase the IPX SAP and IPX RIP update intervals to reduce CPU load (e.g. from default of 60 to 120 seconds).

To increase route stability, you may wish to increase the hold multiplier (default is 3 for 180 seconds). To modify these parameters use the following CLI commands: (4859).

```
config ipxrip <vlan name> update-interval <time> hold-multiplier <number>

config ipxsap <vlan name> update-interval <time> hold-multiplier <number>
```

IPX and Round-Robin Loadsharing

Due to packet sequencing problems, it is not recommended that IPX loadsharing run in conjunction with the round-robin loadsharing algorithm (8733, 9467).

IPX Performance Testing Using Traffic Generators

When using traffic generation equipment to test the wire-speed capability of IPX routing, if entries are allowed to age out with the ports remaining active, those entries cannot be re-learned on that port and will not be forwarded at wire-speed. Restarting the port or clearing the FDB will not address this issue.

In a “real-world” IPX environment, clients and servers generally do not lose communication with the directly attached switch for the FDB entries to age out (9338).

IPX and Bi-Directional Rate Shaping

Bi-directional Rate Shaping is not supported in conjunction with IPX traffic (9226, 9153).

Security and Access Policies

Simulated Mode Creates ACL

When you enable the CPU-DoS-Protect feature in simulated mode, an ACL is still created when a DoS attack is simulated and traffic is blocked (PD2-129163414).

Network Login Design Guidelines and Limitations

Following are Network Login design guidelines and limitations (PD2-130051101):

- All client MACs on an authenticated port will have network access. You cannot authenticate on a per-MAC basis, only per-port.
- All client MACs on an unauthenticated port will see broadcast and multicast traffic.
- Network Login must be disabled on a port before that port can be deleted from a VLAN.
- Campus Mode login will not show the original VLAN to which the port was connected to once the port transition to destination VLAN takes place.
- A Network Login VLAN port should be an untagged Ethernet port and should not be a part of following protocols:
  - ESRP
  - STP
  - VLAN aggregation
  - Load-sharing
- Enabling any of these protocols on Network Login ports will take higher precedence. This may result in a port transitioning from a blocked state to a forwarding state.
- Network Login is not supported for T1, ATM, PoS and MPLS TLS interfaces.
- MSM-failover will clear Network Login state information.

Configure RADIUS with Existing VLAN for Network Login

If you configure your RADIUS server with a VLAN that does not exist on the switch, you cannot log in with Network Login. You must either create the VLAN on the switch or correct the RADIUS configuration. After you correct the configuration, clear the session associated with the failed login before you log in again (PD2-101984392).

RADIUS and the BlackDiamond

When RADIUS authentication is configured on a BlackDiamond switch, upon reboot, you will see the following message indicating that the system is initializing before authentication messages will be transmitted to the configured RADIUS server(s) (7046):

```
“Warning: Radius is going to take one minute to initialize.”
```

RADIUS and Telnet

If one of the following two situations occurs:

1. You have a single RADIUS server configured with a RADIUS timeout value of 10 seconds or more
2. Both primary and secondary RADIUS servers lose their connections and the configured RADIUS timeout value is 5 seconds or more

The switch might not be able to fail over to the local user authentication for telnet sessions. If this happens, the switch cannot be accessed via telnet. This does not occur with the default RADIUS timeout configuration of 3 seconds, or when using alternate session types such as console, SSH, or Vista management (PD2-109828821).

TACACS+ and RADIUS

If TACACS or RADIUS is enabled, but access to the TACACS/RADIUS primary and secondary server fails, the switch uses its local database for authentication.

Network Login and Saving the Configuration

If you save the configuration on a switch while there are open authenticated Network Login sessions, all those sessions will become unauthenticated. This occurs to prevent the authenticated ports from being permanently saved in the authenticated VLAN (1-981ML).

The show netlogin Command Output

If you remove a module with configured Network Login ports and reboot the switch, the output of the `show netlogin` command incorrectly omits the configured ports. Network Login remains enabled on the configured ports and operates correctly if you reinstall the module (PD2-92593101).

Flow Redirection

Enumeration Mode Redirects ICMP Packets

When you create a flow redirection rule for source address based on a subnet mask of /24, enumeration mode is selected, and all ICMP packets are redirected to the next hop. To work around this, use a subnet mask of /16 (PD2-118471863).

Cache Servers Set To “Down” Under Sustained High Traffic Loads

Under very high sustained loads flow redirection might fail and set a cache server to the “down” state and then bring it back up. This only occurs during high loads for a duration of more than 2 minutes. The server will come back up immediately; however, during that time connections that were established might be dropped due to a flushing of the associated IP forwarding database entries. A “down” state is depicted in the log with the following message:

```
09/01/2000 10:51.56 <INFO:IPRT> redirect next hop test <ip_addr> changed to down
```

Health Checking Cannot be Disabled

Flow redirection health checking of the next hop address is turned on by default and cannot be disabled.

NAT

If you change the name of a VLAN that is part of your NAT configuration, the NAT rule configuation is not updated. NAT rule matching continues to operate correctly, but if you save or upload the configuration, the rule is saved or uploaded incorrectly (PD2-82963707).

Vista

VLAN Ports Tagging Information Incorrect

In the Virtual LAN Configuration screen, the information for VLAN ports displays incorrect tagging information (PD2-130140999).

Blackhole Flag Missing

The blackhole flag is missing from the FDB statistics screen (PD2-129387401).

Multicast Address Display

If you configure a routing protocol on multiple interfaces, the Vista statistics page displays the wrong Locally Registered Multicast Address (PD2-105094265).

Configuration Statistics PSU Display

The Vista configuration statistics switch display for the BlackDiamond 6808 shows four power supplies when only two are installed (1-D3RSP).

Closing Internet Explorer 4.0

IE 4.0 caches user login information. In some environments, this can be a security issue. As a work-around, it is best to close the browser after logging out of the switch (1873, 1994).

Vista and RADIUS

If a switch configured for RADIUS authentication loses communication with the RADIUS server, subsequent attempts to use the Vista Web application will take a very long time (measured in several minutes). When communication is lost with the RADIUS server, Vista will continue to make authentication requests for every page. Each request has a timer that will expire leading to an excessive amount of time to bring up each page (5144).

Configuration Options with Large Number of Interfaces

When selecting a configuration applet with a large number of configured interfaces, the traversal of the VLAN interfaces by Vista can cause a Watchdog reset due to the task utilization of Vista during the interface data collection. It is recommended that Vista not be used for configurations with Watchdog enabled where the Vista Configuration applet is used with a large number of VLAN interfaces.

SNMP

Modular Switch get Error

A get request from an NMS to a modular switch for the ifMau<object> on the management port returns a “no such instance” error (PD2-124250702).

SNMP v1 Traps

SNMP v1 traps for link up and link down are not supported. ExtremeWare uses SNMP v2 traps (PD2-110113025).

SNMP and ACLs

Polling the ACL table with a network manager can cause high CPU utilization. For example, with 1,000 ACLs, CPU utilization could be as high as 95%, which could make the console unresponsive (PD2-57475201).

Adding or Deleting a Trapreceiver

Adding or deleting a trapreceiver does not detect the correct community string (1-9I5LD).

Incrementing the intfIf Value

With a getnext or bulkget on a non-existent ifTable object ID, the intf returns next OID value instead of incrementing the intfIf (2-H1OOF).

WinSCP2 Not Supported

The application WinSCP2.exe is not supported. Using WinSCP2 does not cause any problems (1-A5C6C).

SNMP ifAdminStatus MIB Value

The SNMP ifAdminStatus MIB value is not saved after a reboot. Ports set to down in the SNMP ifAdminStatus MIB come back up after rebooting. However, if you save the configuration using the CLI or SNMP after changing the port status to down in the ifAdminStatus MIB, the change is saved after a reboot (2-GOQMD).

Trap Receivers as Broadcast Entry

Although it is possible to enter a broadcast or IP multicast address as an SNMP trap receiver, it will not function (2545).

Bridge MIB Attributes

The IEEE Bridge MIB dot1dTpPortEntry PortInDiscards and dot1dBasePortEntry counters do not increment (4937).

SNMP Time-out Setting

SNMP management stations may need to set the SNMP time-out value to 10 seconds as some large configuration operations take longer to perform (7151).

In addition, when using SNMP tools that use the bulk get request function as opposed to generic get next requests, the MIB walk can time out and subsequently fail with the default time-out setting. It is suggested that the default time-out value be increased from 5 seconds to 60 seconds to decrease the frequency of such time-outs when the get bulk request contains a large number of entries (9592).

SNMP Access Profile

The access profile for SNMP read-only or SNMP read-write can be used for permit-mode only, deny-mode is not operational (7153).

SNMP and Auto-negotiation Settings

For 100/1000Base-TX ports, the ifMauAutoNegAdminStatus can only be disabled if the ifMauDefaultType is set to a speed of 100Mbps. For 10/100Base-TX ports, the user must first set the value of ifMauDefaultType to the correct setting before disabling the ifMauAutoNegAdminStatus (9416).

SNMP and the BGP MIB

When exercising the route table in the BGP MIB, high SNMP utilization messages might be sent to the syslog (11718). This access to the MIB has no adverse effects to any protocol stability (i.e., ESRP, OSPF, BGP).

SNMP and the FDB MIB

When exercising the route table in the FDB MIB with dot1dTpFdbTable enabled, high CPU utilization messages might be displayed in the syslog (PD2-102926801). This occurs when there is a large number of FDB entries and has no adverse affects on protocol stability.

Extreme Fan Traps

The `extremeFanOK` and `extremeFanFailed` traps will contain the `extremeFanNumber` indicating which fan has failed (1-7J571).

Extreme Power Supply Traps

A new object was added “extremePowerSupplyNumber” to the power supply traps. The two RPS traps will no longer be sent out. Instead the extremePowerSupplyGood and extremePowerSupplyFail traps will contain the power supply number indicating which power supply has failed (1-7J56T).

DHCP

The DHCP server is not supported as a standalone feature. It is used as part of the Network Login feature only (1-8SAI6).

Some of the counters for DHCP/BOOTP statistics do not display the correct value. As a result, DHCPRelay statistics are not correctly reported in the IPStats (PD2-73587422).

Diagnostics and Troubleshooting

The show diagnostics backplane-utilization Command Available

The `show diagnostics backplane-utilization` command is not supported on Alpine or Summit switches. Though the command is available, there are no backplane utilization diagnostics available for Alpine or Summit switches (PD2-130597218).

Spurious Message When system-down is Configured

If you configure the system health check alarm level for system-down and a fault is detected, the switch is turned off but continuously logs the message "Card in slot N is off line." You can ignore this message (PD2-129386201).

The use configuration Command

When the switch is in minimum mode, the `use configuration` command has no effect on the backup MSM (PD2-129133801).

Output of the show diagnostics Command

The output of the `show diagnostics` command for the CPU system might display negative numbers, and the totals might not add up properly (PD2-128460401).

Configure Auto-Recovery to online or Alarm-Level to traps

If you configure the system health check auto-recovery to `offline`, save the configuration, and configure the alarm-level to `log`, a health check brings the module or switch offline regardless of how many errors the health check detects. To avoid this, either configure auto-recovery to `online`, or configure alarm-level to `traps` (PD2-124368101).

Error Count Not Accurate

If the switch is flooded with heavy traffic for more than 10 minutes, the `CPU System` field in the `show diagnostics` output is not accurate. The display reports up to 20 more errors (PD2-122738701).

Configuring Diagnostics Mode Off

If you configure diagnostics mode OFF, and then execute the `unconfigure switch all` command, when the switch returns to active state the diagnostics mode is still set to OFF. The default diagnostics mode should be `fastpost`. To verify which diagnostics mode is set for the switch, use the `show switch` command (1-97NL1).

Disable Remote Syslog Before Enabling IPARP Debug-Tracing

With remote syslog enabled, if you configure the IPARP debug-trace to level 2 or higher, the switch hangs and is rebooted by the watchdog timer. To avoid this, disable the remote syslog prior to configuring the debug-trace (PD2-110983505).

Rebooting Using SNMP or RMONII With Reboot Loop Protection

If you use SNMP or RMONII to issue the `reboot` command, and reboot loop protection is configured with a threshold of 1, the switch will reboot into minimal mode (PD2-111307101).

Configuring a New Threshold for Reboot Loop Protection

When a new threshold is configured for reboot loop protection, the time stamp is not cleared, and the reboot-threshold can be violated (PD2-109830745).

The card-down Option

In a fully redundant configuration, if you configure the `card-down` option in the `configure sys-health-check` command and checksum errors are detected, the MSM is not taken offline as expected. To work around this, use the `configure sys-health-check auto recovery 3 offline` command (PD2-105991401).

Do Not Use a Count of One for Reboot Loop Protection

If you configure a large threshold, do not configure a count of one. If you reboot the switch manually, which resets the timer, and the time to reboot falls within the threshold you have configured, the switch detects the reboot and enters minimal mode (PD2-111222216, PD2-111201401).

Documentation

The *ExtremeWare 7.0.0 Command Reference Guide* incorrectly states that flow statistics ping-checking is enabled by default. In ExtremeWare 7.0.0, flow statistics ping-checking is disabled by default.

The *ExtremeWare 7.0.0 Software User Guide* incorrectly states that indirect LSPs are supported on IS-IS networks. Indirect LSPs are supported on OSPF networks only.

The *ExtremeWare 7.0.0 Command Reference Guide* incorrectly states that loopback detection is enabled by default. By default, loopback detection is disabled.

Issues Resolved in ExtremeWare 7.0.1b11

The following issues were resolved in ExtremeWare 7.0.1b11. Numbers in parentheses are for internal use and can be ignored.

General

The `show ports configuration` command now correctly displays loopback status (PD2-121610228).

BlackDiamond

IP forwarding from an MPLS LSP to a TLS tunnel is now supported (PD2-118230301).

TLS VLANs no longer block IGMP joins, so routers downstream from the BlackDiamond can now join multicast groups (PD2-117499435).

Multicast traffic is now routed correctly when the T1 module is configured for IPCP (PD2-121607310).

If you configure more than 10 ATM PVCs on a switch and upload the configuration, the switch no longer crashes (PD2-120100801).

Summit

If you download a configuration, autopolarity detection is no longer automatically enabled on the Summit48si (PD2-118279201).

The `configure ports auto-polarity` command is no longer available on platforms other than the Summit48si (PD2-118503001).

Issues Resolved in ExtremeWare 7.0.0b68

The following issues were resolved in ExtremeWare 7.0.0b68. Numbers in parentheses are for internal use and can be ignored.

BlackDiamond

On the BlackDiamond 6816, the syslog no longer reports memory mapping failures on the MSM64i in slot B as being on the MSM64i in slot C (PD2-112505701).

Alpine

When using T1 modules, if all ports are tagged in the same VLAN and traffic stops completely, the T1 FDB entries now age out correctly (PD2-117286601).

ESRP

You can no longer configure a failover priority greater than the upper limit, which is 255. Values greater than 255 generate error messages (PD2-97286301).

You can no longer configure the ESRP tracking failover priority to a value greater than or equal to the ESRP VLAN priority (PD2-81790880).

You are no longer required to disable load sharing before adding or deleting ESRP restart ports that are also in the load sharing group (PD2-110113235).

The `show esrp` command output now contains VLANs with both IPX and IP enabled (PD2-102292601).

You are no longer able to create more than 64 ESRP domains (PD2-97286303).

Spanning Tree

After an STP recovery, an IGMP snooping disabled VLAN now correctly forwards multicast streams. (PD2-109828901).

QoS

The ACL FDB is now completely cleared when you delete flow redirect rules (1-EQRVD, PD2-110802201).

IS-IS

If you configure IS-IS authentication on a PoS interface, the configuration is now saved after a reboot or configuration download (PD2-108735813).

If you configure authentication before enabling IS-IS on a VLAN, the VLAN interface is now correctly authenticated (PD2-108735827).

Issues Resolved in ExtremeWare 7.0.0b61

The following issues were resolved in ExtremeWare 7.0.0b61. Numbers in parentheses are for internal use and can be ignored.

General

If you unconfigure a software redundant port, the redundant port no longer remains down (PD2-105118423).

Link transitions no longer generate additional checksum errors when checksum errors have already been recorded (PD2-104485401).

IPX now operates correctly after upgrading from ExtremeWare 6.2.1 to ExtremeWare 6.2.2 (PD2-83558641).

The output from the `show port configuration` command now displays flow control autonegotiation status correctly (1-5VKAH).

Broadcast storms no longer lock up the management port (PD2-71281360).

Name completion now works correctly with the `show iparp` and `show ipfdb` commands (PD2-97719004).

Enabled jumbo frame ports are no longer displayed twice in the output of the `show config` command (PD2-94298501).

Upgrading from ExtremeWare 6.1.9 to ExtremeWare 6.2.2 no longer corrupts permanent FDB entries (PD2-98730521).

False checksum error messages due to a software error when checking for a return code are no longer displayed to the system log (PD2-108354018).

If you create a tagged VLAN and port, delete the VLAN, and recreate an untagged VLAN and port, broadcast traffic is now forwarded correctly (PD2-100769947).

The maximum IP-MTU size is now 9194 to fit within the maximum jumbo frame size of 9216 (PD2-105579101).

When a ping is redirected, the statistics for the last packet received are no longer reported as lost (5170).

BlackDiamond

The `synchronize` command no longer causes the BlackDiamond to crash (PD2-82006728).

Multicast packets are now forwarded correctly after you disable IGMP snooping (1-EUCKP).

Extended diagnostics on a P12cSi or P12cMi module cause the BlackDiamond to crash. Normal diagnostics work correctly (PD2-83547401).

ESRP status change no longer cause the OSPF neighbor to crash with the following error message (PD2-81648327):

```
"Error: Insertion FAILED Neighbor already exists."
```

ExtremeWare now correctly checks fan tray status (PD2-63609301).

The ARM no longer crashes on a fully loaded BlackDiamond with a heavy traffic load (PD2-95645103).

If you remove a power supply or have an empty power supply bay, ESRP now changes the priority of the ESRP VLAN to the failover setting (PD2-86682767).

The F96Ti and G12SXi modules no longer lose the connection to the MSM64i (PD2-93060104).

If you have 3,000 VLANs configured, the `clear slot` and `unconfig slot` commands no longer cause the switch to see the slot as mismatched and unconfigured (PD2-90223205).

Summit

The FDB now correctly recognizes ESRP-aware Summi48i redundant ports after a link transition, and the ESRP state is now also correctly recognized. In addition, connections to the redundant ports no longer show as simultaneously active after one of the redundant links transitions (PD2-89481383).

When using 802.1Q tagged, odd size packets on 10/100 Mbps links no longer cause the Summit48i to drop packets (1-EGCA8).

Enabling the master load-sharing port can no longer cause the redundant port on the Summit48i to transition (PD2-105853230).

IP Multicast Routing

ExtremeWare no longer continues to send a prune to a pruned interface even if a new sender is introduced to the pruned interface, which was causing multicast streams to intermittently stop (PD2-89479820).

An ARP request with 0 in the protocol address is no longer reported as a duplicate address if the switch has the default VLAN with no IP address assigned (PD2-70889799).

RIP Routing

After a gateway is changed, affected RIP routes are now considered the preferred routes and the routes are advertised to RIP peers (PD2-97347827).

EAPS

EAPS and Spanning Tree superloop configurations no longer result in blackholes due to incorrect FDB table entries (PD2-89367538). To correct this problem, you must upgrade all switches to ExtremeWare 7.0 (or later).

The EAPS master now sends an IGMP query upon link failure, if you configure the EAPS master interface with an IP address on the same protected VLAN as the transit node (PD2-100679119, PD2-91657073, PD2-112059601).

A rapid link transition up, down, then up again on an EAPS transit switch no longer causes traffic disruption (PD2-102929324).

FDB

IP FDB aging is now disabled by default (PD2-93733656).

IP FDB entries are now correctly cleared when a new indirect LSP comes up (PD2-97156301).

Aging 250,000 or more IP FDB entries no longer consistently uses over 50% of the CPU (PD2-93733654).

If 256,000 or more FDB entries are injected into the switch and then you create a new VLAN, that VLAN might now correctly has an FDB entry (PD2-97137701).

Using the `clear FDB` command on switch that has over 3,000 VLANs assigned to all ports no longer triggers a watchdog reboot if 50,000 or more FDB entries are present on one of the VLANs (PD2-69816950).

BGP

When exporting static and direct routes, the BGP origin is now correctly incomplete (PD2-92668201).

If the number of routes to be processed for next hop changes is a multiple of 5000 and no other BGP activity causes BGP to release the CPU, next hop change processing no longer re-starts after completion. (PD2-95403801).

OSPF

OSPF now recalculates the cost of external routes correctly when a redistribution from static to OSPF is deleted and re-inserted (PD2-64596001).

In an environment with a large number of IPFDB entries associated with OSPF routes/default route, an OSPF change that caused an SPF calculation no longer consumes significant CPU cycles, which caused the system to become unresponsive for several minutes (PD2-99111708).

Spanning Tree

When you configure STP with 4,000 VLANs, an STP topology change no longer triggers a watchdog reboot due to the amount of time required to complete aging of the FDB entries for the VLANs (PD2-80183102).

ESRP

A link transition on a redundant switch no longer causes a momentary ESRP dual master situation (PD2-95068802).

Failover priority 0 now operates correctly (PD2-68325201).

The slave ESRP VLAN no longer forwards traffic when the slave link transitions without an ESRP state change (PD2-89481303).

If you create, delete, or modify a VLAN tag when there are 256,000 MAC address, you no longer receive the following error message (PD2-90054207, PD2-90223201):

```
updateEdpFilter401: Unable to locate EDP MAC (VID=0xffd)
```

VLANs

You are no longer required to set the CPU-transmit-priority level to "normal" to configure more than 1024 VLANs (7120, 8908).

IS-IS

The `show isis` command now operates correctly with user accounts (PD2-92240902).

The `show configuration` command now correctly displays layer 2 VLAN status after you add a VLAN interface to `level-2` (PD2-94272213).

NetFlow

If a flow record filter is configured on one port with type "match-all-flows" you can now configure the same flow filter on other ports (1-7G1D8).

SNMP

The ExtremeWare 7.0.0 MIB now compiles correctly in 3rd party applications (PD2-97119501).

The `disable snmp trap port-up-down all` command now operates correctly (PD2-96936507).

Entries in the alarmTable related to SMON, extremeRtStats, and extremeVlanL2Stats are lost after reboot and no longer create spurious corrupt entries (PD2-91569801).