ExtremeWare Release Note

Software Version 6.2.2 b18

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

Published: June 2002
Part Number: 120156-00 Rev 01

Author: Rich Small

Editor: Rich Small

Production: Rich Small

Special Thanks: Affan Zaheer, Connie Knighton, Daniel Chew

Contents

1 Overview

These Release Notes document ExtremeWare 6.2.2b18. ExtremeWare 6.2.2 introduces new hardware products and software features.

This chapter contains the following sections:

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 6.2

Following are descriptions of features introduced or enhanced in ExtremeWare 6.2.0 and subsequent releases. These features are documented in detail in the *ExtremeWare Software User Guide* or the *ExtremeWare Software Command Reference Guide* for the relevant software version, unless otherwise noted.

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

Features Added or Enhanced in ExtremeWare 6.2.2

- Dynamic Memory Scanning and Memory Mapping: These features now work on the following products, in addition to the BlackDiamond (1-B98W9):
  - Alpine 3808
  - Alpine 3804
  - Summit1i
  - Summit5i
  - Summit7i
  - Summit48i
  - Summit48si

NOTE

The `config sys-health-check auto-recovery` *command does not support the* `number-of-tries` *option on Summit switches.*

- MSM64i Failover: Failover times for the MSM64i have been improved. Boot time is up to 37% faster, configuration saves are up to 35% faster, and software packet forwarding is up to 65% faster (1-B2G7X). The actual improvement is dependent upon the type of traffic, your specific configuration, and other issues.
- You can now force the master MSM64i to immediately fail over to the slave MSM64i with the `run msm-failover` command.
- Improved BlackDiamond 6816 POST and Diagnostics: The BlackDiamond 6816 POST speed, error output, and diagnostics now match those of the BlackDiamond 6808 (1-AJTM7).
- System Memory Dump: You can now download the entire contents of memory through the Ethernet management port (1-B98W1). This feature is for troubleshooting, and should not be used without assistance from TAC. The following command transfers the dump. If you do not specify an IP address, the configured system-dump server IP address is used:

  ```
  upload system-dump [<ip address>]
  ```

  This command specifies the IP address to which to transfer a dump if the `system-dump` option is specified in the configuration. This address is also used if no address is provided in the `upload system-dump` command. The default is 0 or "no IP".

  ```
  config system-dump server <ip address>
  ```

  The following command sets an optional timeout for the dump transfer. The default is 0. The minimum non-zero value is 120 seconds. The minimum recommended value is 480 seconds.

  ```
  config system-dump timeout <seconds>
  ```

  The following command returns the system-dump configuration to the defaults.

  ```
  unconfig system-dump
  ```

  The following command displays the system-dump server IP and dump-timeout.

  ```
  show system-dump
  ```

  The sys-recovery-level command has a new option: `system-dump`. The `system-dump` option specifies a memory dump if a task generates a software exception. The four options specify the action taken when the dump transfer is complete. The actions occur whether or not the dump was successful. The `maintenance-mode` option leaves the switch in whatever state it was in before the dump.

  ```
  config sys-recovery-level [all | critical] system-dump [maintenance-mode |
  msm-failover | reboot | shutdown]
  ```

  These commands are not documented in the *ExtremeWare 6.2.2 Command Reference Guide*.
- System Odometer Enhancement: The system odometer has been enhanced to record the date the unit was first installed (1-7KEEX).
- ARP response time is now faster (1-631SX).
- BGP Enhancements:
  - You can now limit the number of IP prefixes from a BGP neighbor (1-5LQI3).
  - The `show bgp neighbor detail` output was enhanced to provide information on the maximum prefix limit.
  - The Type of Service (TOS) field in the IP header is now set to "Internetwork Control" (1-9GPU2). This gives the IP packet preferential treatment at the intermediate nodes and at the destination.

— The BGP task has been restricted to 75% utilization of the CPU.

— TCP can now ignore the sequence number if it is lower than the expected receive sequence number, but process the rest of the fields in the ACK packet (1-8406T).

— If the reachability to a BGP next hop changes, only associated routes are reprocessed. If only the immediate gateway changes, no BGP updates are sent to the peers. Last, IGP next hop changes are processed before the BGP next hops (1-9GM22).

— BGP error processing is now delayed by 5 seconds, to allow IGP recalculation (1-AX8MX).

— The Input Policy can no longer be re-applied to an already-modified Path Attribute, thus protecting the Path Attributes (1-96X1V).

— The `config debug trace bgp-neighbor` command is now supported and accepts IP addresses (1-9FBFL).

— The `clear bgp neighbor counters` command now clears FsmTransitions (1-BZC35).

- EAPS Enhancements:

  — You can no longer modify the configuration of a SuperBridge, Subbridge, or IP Range control VLAN (1-A7RPG).

  This limitation is not documented in the *ExtremeWare 6.2.2 Software User Guide*.

  — You can no longer modify the configuration of any control VLAN if the domain is active (1-A7RPG).

  This limitation is not documented in the *ExtremeWare 6.2.2 Software User Guide*.

  — You can no longer delete a domain if that domain is active (1-A7RPG).

  This limitation is not documented in the *ExtremeWare 6.2.2 Software User Guide*.

  — Enhanced `show eaps` command: This command now provides a `summary` option, as follows (1-8330Z):

  ```
  show eaps [detail | summary]
  ```

- Multicast Enhancements:

  — Configure PIM Register Interval: You can now configure the PIM register interval threshold and probe interval using the `config pim register-suppress-interval <time> register-probe-interval <time>` command (1-8BB01, 1-885VT).

  — PIM Scalability: You can now configure the register interval for initial registers using the `config register-rate-limit-interval` command (1-8BB01, 1-993WD).

  This command is not documented in the *ExtremeWare 6.2.2 Command Reference Guide*.

  — Configurable SPT Threshold: You can now configure the SPT threshold using the `config pim spt-threshold <last hop router threshold>` command without configuring the RP threshold. In addition, a cache of the Source Specific Tree is created on the RP upon receipt of the register packet even though no cache of Shared Tree exists. Thus, new receivers need not wait for the next cycle of the null-register to receive the desired multicast stream (1-9157U).

- NAT Enhancement: NAT has been improved to increase performance by up to 25%.

- SNMP Enhancements:

  — You can now enable and disable traps by port, using the `enable/disable snmp traps { port-up-down ports [<portList> | all] }` command (1-6DDDC).

  — Enhanced `show management` output now includes trap information per port.

  — Trap receivers now support an enhanced mode, which accepts traps containing extra varbinds at the end. Only standard traps are sent to a standard mode trap receiver and only enhanced traps

are sent to an enhanced mode trap receiver. In addition, a standard mode trap receiver is only sent standard traps from domain “s0”. For other domains, no traps are sent (1-61BMV, 1-BPZO1, 1-BP5AN).

— The “ifDescr” and “ifAlias” varbinds were added to the proprietary EDP traps (1-60FXQ).

— You can now enable loopback mode on a VLAN using SNMP (13029).

— You can now configure the CPU transmit priority using SNMP (1-5E4YT).

— You can now manage multiple Spanning Tree domains using SNMP (1-60FXK).

— The BlackDiamond 6804 is now supported through SNMP (1-DDMYT).

— The show management command output was reorganized for better clarity (1-CZWD9).

- The number of OSPF Equal Cost Multiple Paths was increased from 8 to 12 (1-9G4OY).
- The syntax of the `config ospf timer <transmit_delay>` command is now `config ospf timer <transit delay>` (1-B1K6X).
- The VLAN name in the `show vlan` command output was expanded to display up to a maximum of 18 characters (1-99BX9).
- You can now overwrite 802.1p priority based on VLAN using the `create fdb any-mac vlan <name> dynamic ingress-qos <qosprofile>` command (1-968RZ).

  This command is not documented in the *ExtremeWare 6.2.2 Command Reference Guide*.
- The `show ports info` command output was enhances to include “enabled” and “disabled” flags for each port (1-8CEA7).
- You can now display the loadsharing groups using the `show ports sharing` command (1-8P8B9).
- You can now specify a default route as a blackhole using the `config iproute add blackhole [default | ipaddress <ip address>]` and `config iproute delete blackhole [default | ipaddress <ip address>]` commands (1-99CRT).
- Up to 1024 locked-down MAC addresses are now saved through reboots and downloaded configurations (1-ADCFL).
- Network Login users now have a Vendor-Specific Attribute to limit their access to the switch (1-98DW9).
- You can now configure a banner for Network Login users using the `config banner netlogin` command (2-GQ1XR).
- You can now configure a domain suffix list up to six items long for the ExtremeWare DNS client (1-69RS9).
- ExtremeWare now supports protection against the “jolt,” “opentear,” “raped,” “boink,” and “winfreeze” DoS attacks (1-AY26W).
- If CLI paging is disabled, you can now press [q] and [Enter] to force the output to stop (1-90L0R, 1-7H2KP).
- You can now display the log in ascending chronological order using the `show log chronological` command (1-6AWX6).
- The maximum number of Access List entries is now 5120 (1-BPLJP).
- When ESRP restarts a port, it sends a message to the syslog (1-BK08Q).
- You can now configure the ESRP timeout using the `config vlan esrp timer` command (1-BKEQP).
- Permanent FDB entry lookup is now faster (1-EB4IT).
- The software redundant port feature now fully supports IPX (1-A5EMP).

- The show iproute command output now includes a time stamp (1-60FYA).
- Spanning Tree now detects loopback ports, blocks the ports, and recalculates (1-D3TSX, 1-DWC1D).
- The auto-negotiation status of Gigabit Ethernet ports is now contained in the output of the `show configuration` command (1-962XM).
- The `mgmt` option was removed from the `restart port` command (1-9KMSI).
- The `show tech-support` command now captures the `show ipmc fdb` command output (1-A85H1).
- The show management command output now includes a reminder that disabling a web access configuration requires a reboot (1-BTWKQ, 1-8PJVZ).
- The `show vlan` command output now indicates active and inactive ports (1-E3MRP).

Features Added or Enhanced in ExtremeWare 6.2.1

- Show ESRP Aware Command: To display ESRP aware information, use the following command:

  ```
  show esrp-aware vlan <vlan name>
  ```

  If you do not specify a VLAN, ESRP aware information for all VLANs is displayed. The display includes the group number, MAC address for the master of the group, and age of the information.

  This command is not documented in the *ExtremeWare 6.2.1 Command Reference Guide*.

- UPD Echo Server Support: You can now use UDP Echo packets to measure the transit time for data between the transmitting and receiving end. To enable UDP echo server support, use the following command:

  ```
  enable udp-echo-server
  ```

  To disable UDP echo server support, use the following command:

  ```
  disable udp-echo-server
  ```

  These commands are not documented in the *ExtremeWare 6.2.1 Command Reference Guide*.

- Show DHCP Commands: Two new commands display DHCP information. To display the IP address, MAC address, and time assigned to each end device, use the following command:

  ```
  show vlan dhcp-address-allocation vlan <vlan name>
  ```

  To display the DHCP configuration, including the DHCP range, DHCP lease timer, network login lease timer, and DHCP-enabled ports, use the following command:

  ```
  show vlan dhcp-config vlan <vlan name>
  ```

  These commands are not documented in the *ExtremeWare 6.2.1 Command Reference Guide*.

- Odometer: An Odometer feature was added to keep track of how long each individual component in the whole switch has been functioning since it is manufactured. This odometer counter will be kept in the EEPROM of each monitored component. This means that even when the component is plugged into different chassis, the odometer counter will be available in the new switch chassis. The following components are monitored by the odometer:

  — For the BlackDiamond — MSM64i and I/O modules

  — For the Alpine — SMMi, I/O slots, and power supplies

  — For stackable switches — the CPU

  The odometer "readings" can be viewed using the `show odometers` command.

  This command is not documented in the *ExtremeWare 6.2.1 Command Reference Guide*.

- Show Flow-Redirect Command Enhancement: The output of the `show flow-redirect` command is clarified to better indicate configured protocols.

- Web Cache Redirection (WCR) Modifications: Changes were made to support the option of specifying the L4 source IP port in place of the L4 destination IP port.

  WCR was enhanced to load balance among the next hops based on the IP source addresses of the flow redirect rule.

  In addition, configurable health check support has been added to control the frequency and timeout values of the `ping-check`, `tcp-port-check`, and `service-check` health checks.
- BGP update message transmission has been optimized for faster synchronization and convergence.
- A number of enhancements were made in ExtremeWare diagnostics. These include improved speed (decreased boot time) for the Fastpost diagnostic, improved error output, and adding some minor packet memory testing to the Normal level diagnostic.
- Dynamic Memory Scanning and Memory Mapping (BlackDiamond)
- Disable BlackDiamond and Alpine slots
- EAPS enhancements: The EAPS functionality was enhanced in two ways in ExtremeWare 6.2.1:
  - EAPS-Transit switches and ESRP-aware switches now use different hardware queues for broadcasting Control Messages and for forwarding generic broadcast traffic.
  - Sanity checking is now done to ensure that the configuration of an EAPS domain and the configuration of the Control VLAN are correct before attempting to start the EAPS domain.
- Extreme Standby Router Protocol (ESRP) enhancements:
  - Domain Member and Sub-VLAN support on Host Attached ports.
  - Two new ESRP election algorithms are provided to ensure no fail back if the original Master recovers.
- SCP2/SSH2 client
- MAC Address Lock-down
- Port Frame/Packet flooding
- Worldwide Daylight Savings Time support
- An option was added to the `show log` command to display log entries in chronological order.
- The `show iproute` command was enhanced to display only the routes selected.
- VLAN names can be specified using the tab key for command completion.
- An `all` option was added to the `delete access-list` command to allow deletion of all access lists.
- The `show iparp` command was enhanced to accept a MAC address as an option.
- The number of primary and secondary DNS name servers that can be configured was increased to 8.
- You can now modify health check timeout and frequency for the flow redirection health checks.

Features Added or Enhanced in ExtremeWare 6.2.0

- Network Address Translation (NAT)
- MAC Address Security
- NetFlow
- Virtual Router Redundancy Protocol (VRRP)
- Extreme Multiple Instance Spanning Tree Protocol (EMISTP)
- Network Login

- Ethernet Automatic Protection Switching (EAPS)
- Software Controlled Redundant Port
- BGP Enhancements: A number of enhancements were made to BGP support. They include the following:
  - Support for displaying the BGP community in either decimal number format or <AS number>:<community number> format.
  - Support for adding or deleting a particular value to the MED attribute in the BGP Route that has been received, and support for removing the MED attribute completely.
  - Support for arithmetic determination of a MED value from IGP metric to route's NextHop.
  - The ability to remove Private AS numbers from the AS Path attribute associated with the routes in the updates sent to EBGP neighbors.

Supported Hardware

This release of ExtremeWare 6.2.2 is designed to support products using the "*i*" chipset *only.*

ExtremeWare 6.1.9 and 6.2.0 or above requires version 7.2 BootROM. Note that BootROM 7.2 is not backward compatible with previous versions of ExtremeWare 6.x.

ExtremeWare 6.1 or later requires that the BlackDiamond switch use only the MSM64i in slots marked "A", "B", "C", and "D". It is not possible to use MSM32 modules with ExtremeWare 6.x or higher.

This release supports the following hardware in addition to the hardware mentioned in the User Guides (support for hardware listed in *italics* is new for this release):

Table 1: Supported hardware

Extreme Switch Platform	ExtremeWare Filename/Version	BootRom Filename/Version
BlackDiamond 6804	v622b18 .xtr or v622b18_ssh.xtr/v6.2.2b18	Ngboot7.6.bin/v7.6
BlackDiamond 6816	v621b27_6816.xtr or v621b27_6816_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2
BlackDiamond switch using MSM64i MSMs	v621b27.xtr or v621b27_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2
Alpine 3808	v621b27.xtr or v621b27_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2
Alpine 3804	v621b27.xtr or v621b27_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2
Summit 7i/7iT	v621b27.xtr or v621b27_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2
Summit 1i/1iT	v621b27.xtr or v621b27_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2
Summit 5i/5iT/5iLX	v621b27.xtr or v621b27_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2
Summit 48i	v621b27.xtr or v621b27_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2
Summit 48Si	v621b27.xtr or v621b27_ssh.xtr/v6.2.1b27	Ngboot72.bin/ v7.2

NOTE

Please see the "Upgrading to ExtremeWare 6.2" chapter for special upgrade instructions when upgrading from 6.1.8b13 or below.

NOTE

The BlackDiamond 6816 requires its own ExtremeWare image (v621b20_6816.xtr or v621b20_6816_ssh.xtr). The image that runs on other BlackDiamond, Alpine, or stackable switches does not support the BlackDiamond 6816.

BlackDiamond Module Support

BlackDiamond modules supported with ExtremeWare 6.1.5 and above and the MSM64i include:

Table 2: BlackDiamond module support

BlackDiamond Module	ExtremeWare 6.1.5 and above Support	Uses "*i*" Chipset
MSM64i	Yes	Yes
G12SXi	Yes	Yes
G8Xi	Yes	Yes
G8Ti	Yes	Yes
F48Ti	Yes	Yes
WDMi	Yes	Yes
F96Ti	Yes (EW 6.1.8b12 or above)	Yes
F32Fi	Yes (EW 6.1.8b13 or above)	Yes
F32T	Yes	No
F32F	Yes	No
G4SX - G4LX	Yes	No
G6SX - G6LX	Yes	No
DC Power Supply	Yes	N/A
110V AC Power Supply	Yes	N/A

NOTE

Mixed versions of the power supplies should not be installed in the same system. Both power supplies should be of the same type.

Alpine Module Support

Alpine modules for the Alpine 3808 or 3804 Chassis supported with ExtremeWare 6.1.5 and above include:

Table 3: Alpine module support

Alpine Module	ExtremeWare 6.1.5 and above Support	Uses "*i*" Chipset
SMMi	Basic or Advanced license	N/A
GM-4Si/Xi/Ti	Yes	Yes
FM-32Ti	Yes	Yes
FM-24MFi	Yes	Yes
FM-24Ti	Yes (EW 6.1.7 or above)	Yes
FM-24SFi	Yes (EW 6.1.7 or above)	Yes
GM-WDMi	Yes (EW 6.1.8 or above)	Yes
DC Power Supply	Yes	N/A

Summit Component Support

Summit components supported with ExtremeWare 6.1.5 and above include:

Table 4: Summit component support

Summit Module	ExtremeWare 6.1.5 and above Support	Uses "*i*" Chipset
Summit 7i DC Power Supply	Yes	N/A

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)

Summit48si Mini-GBIC Support

The Summit48si supports the Extreme mini-GBIC only. For reliability and stability reasons, third-party mini-GBICs are not supported at this time.

Features Unique to the "*i*" Chipset

The following list summarizes the feature areas specific to the "*i*" chipset products. Unless noted otherwise, both I/O module and MSM must make use of the "*i*" chipset to make use of the features listed below.

- QoS and Access Policies — Complete use of IP Access Lists; support for IP DiffServ; support for eight QoS queues per port, instead of four; support for Random Early Detection.
- Bridging/Switching — Support for jumbo frames; support for address and round-robin-based load-sharing algorithms and non-contiguous load-sharing port groups.
- Routing — Wire-speed IPX routing (products without the "*i*" chipset support IPX routing, but not at wire-speed). Support for BGP4 (though it is not necessary to have "*i*"-based I/O modules to support BGP4 on the BlackDiamond). Policy-based Routing.
- Server Load Balancing — Support for all Server Load Balancing functions.
- Web Cache Redirection — Support for all WCR functions.
- QoS Bi-directional Rate Shaping — Ability to perform Policy-based QoS for a VLAN's traffic both into and out of the switch.
- Extreme Standby Router Protocol (ESRP) options — Support for ESRP Groups, ESRP Domains and ESRP Host Attach.
- Traffic statistics on a per VLAN basis.
- Subnet directed broadcast packet forwarding improvements.
- System health-checker on the BlackDiamond.
- Extreme Multiple Instance Spanning Tree Protocol (EMISTP) — An extension of STP that allows a physical port to belong to multiple STPDs by assigning the port to multiple VLANs.
- Extreme Automatic Protection Switching (EAPS) — Support for fast protection switching to layer 2 switches interconnected in an Ethernet ring topology.
- Virtual Redundant Router Protocol (VRRP) — Support for the proposed IETF standard protocol that allows multiple switches to provide redundant routing services to users
- Network Address Translation (NAT) — Ability to convert a set of IP addresses, typically private IP addresses, to another set of IP addresses, typically public Internet IP addresses.
- Network login — Ability to control the admission of user packets into a network by giving addresses only to users that have been properly authenticated.

2 Upgrading to ExtremeWare 6.2

This chapter contains the following sections:

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 flash 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 the BootROM

This release is supplied with a new BootROM image: BootROM 7.6. The new BootROM is critical to ExtremeWare 6.2.2 and is not backward compatible with ExtremeWare 6.1.8 or earlier. Upgrade the BootROM upgrade *before* upgrading ExtremeWare using the following command:

```
download bootrom [ <host_name> | <ip_addr> ] <ngboot.bin_name>
```

To downgrade to an earlier version of ExtremeWare, perform a BootROM downgrade *before* downgrading ExtremeWare. To downgrade BootROM, use the following command:

```
download bootrom [ <host_name> | <ip_addr> ] <ngboot.bin_name>
```

Once the BootROM downgrade is complete, you can reboot the system with the previously loaded ExtremeWare version.

Upgrading ExtremeWare

Below are instructions specific to upgrading to, and downgrading from, ExtremeWare 6.2 for Summit, Alpine, and BlackDiamond switches.

NOTE

You must upgrade to BootROM 7.2 to run ExtremeWare 6.1.9 or later. Also note that you must downgrade to BootROM 6.5 to run ExtremeWare 6.1.8 or earlier. See below for instructions on BootROM upgrades.

Upgrading Switches

ExtremeWare 6.2 can read a stored configuration saved by ExtremeWare 6.x. The procedures outlined below will preserve 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 primary` and `save secondary` commands.
2. Configure the switch to use the primary image and the primary configuration using the `use image primary` and `use config 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 config <ipaddress> <filename>` command.
5. If you are not already running BootROM 7.6, TFTP download BootROM 7.6 to the switch. An example command is "`download bootrom <ipaddress> ngboot7x.bin`". Reboot the switch to come up with BootROM 7.6.

Upgrading from ExtremeWare 6.1.8 or Earlier

If you are running a version of ExtremeWare **prior** to **6.1.9b11**:

1. TFTP download a version of ExtremeWare 6.1.9 to the primary image space. An example command is "`download image <ipaddress> 619b22.xtr primary`".

   CAUTION

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

   ```
   ERROR: File too large
   ```

2. Reboot the switch. The previous configuration of the switch will be preserved going from the previous version of ExtremeWare to ExtremeWare 6.1.9. Verify that the switch is operating as expected. Save the configuration to the desired configuration location.

   NOTE

   If you have configured "Random Early Drop Probability" in the 6.1.8 or below configuration, after upgrading to 6.1.9, please re-configure the "Random Early Drop Probability" to the desired value

using "config red drop-probability" command and save the configuration to the desired location prior to upgrading to 6.2.2.

3 Follow the instructions for upgrading from ExtremeWare v 6.1.9 in the next section.

Upgrading from ExtremeWare 6.1.9

If you are running ExtremeWare version 6.1.9:

1 TFTP download ExtremeWare 6.2.2 to the primary image space using the `download image <ipaddress> v622b18.xtr primary` command.

2 Reboot the switch. The previous configuration of the switch will be preserved. Verify that the switch is operating as expected. After verification, you can configure features specific to the current version of ExtremeWare.

3 Save the configuration to the primary space.

NOTE

After upgrading from 6.1.9 to 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. To change the leave time-out value back to 10 seconds, use the following command:

```
configure igmp snooping leave-timeout 10000
```

Upgrading from ExtremeWare 6.2.0b60 or Later

If you are running ExtremeWare version 6.2.0b60:

1 Upload the configuration to your TFTP server.

2 TFTP download ExtremeWare 6.2.2 to the primary image space using the `download image <ipaddress> 62xby.xtr primary` command.

3 Reboot the switch.

4 TFTP download the configuration you saved in Step 1, and enter "Y" to reboot the switch.

Verify that the switch is operating as expected. After verification, you can configure features specific to the current version of ExtremeWare.

5 Save the new 6.2.2 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

It is assumed that you have followed the upgrade instructions correctly and that the desired previous configuration has been preserved in the secondary configuration space.

1 If, as per upgrade instructions, the secondary configuration was saved while using a 6.1 or previous 6.1 image, configure the switch to use the secondary configuration with the `use config secondary` command. If there is no stored configuration saved for that version of ExtremeWare,

you will need to re-configure or download the correct configuration file to the switch when running the desired image.

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

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

4 Downgrade the BootROM version to 6.5 if you are pointing the image back to a version of ExtremeWare previous to 6.1.9. The `show version` command displays the BootROM version as "Unknown" when the BootROM is downgraded.

5 Reboot the switch. If you have followed upgrade instructions, your original configuration should be in place. If you did not have the correct configuration downloaded, you may provide a minimal configuration for the switch through CLI sufficient to TFTP download the configuration file generated during the upgrade procedure. If you do not have the configuration file, re-configure the switch manually.

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.

Installing ExtremeWare 6.2.2 on a BlackDiamond 6816

To install the BlackDiamond 6816 version of ExtremeWare 6.2.2 on a BlackDiamond 6816 for the first time, you must follow the procedure that follows. These steps must be followed even if you already have ExtremeWare 6816b6 installed.

1 Remove all modules (MSM64i and I/O) except the MSM64i in slot A.

NOTE

Make sure you have no configurations saved in primary or secondary.

2 Download BootROM 7.6 using the `download bootrom` command.

3 Enter "Y" to complete the upgrade.

4 Reboot the switch using the `reboot` command.

5 Download ExtremeWare 6.1.9b11 or 6.1.9b22 using the `download image` command.

NOTE

You only need to load code into the primary image.

6 Confirm the installation using the `show version` and `show switch` commands.

NOTE

Make sure you are booting to your primary image. Otherwise, configure the switch to boot from the primary image with the `use image primary` *command.*

7 Reboot the switch using the `reboot` command.

8 Download ExtremeWare 6.2.2 using the `download image` command.

NOTE

Install code image into both primary and secondary.

9 Confirm the installation using the `show version` and `show switch` commands.

10 Clear the log using the `clear log static` command.

You must perform the same steps for each MSM64i.

Repopulate the Chassis

To repopulate the chassis after you have installed the BlackDiamond 6816 version of ExtremeWare 6.2.2 on each MSM64i, perform the following steps:

1 Power down the chassis.

2 Install MSM64i modules in slots A - D.

3 Install all I/O modules.

4 Power up the chassis.

5 Confirm that each MSM64i is running the correct version of ExtremeWare using the `show switch` command.

6 Check the log using the `show log` command.

7 If you have critical or major errors, save them into a text file and contact Extreme Technical Support.

Error Message

If you install an MSM64i with a BlackDiamond 6816 image onto a BlackDiamond 6808 chassis, you might get an error message in the syslog indicating the image is not supported, as indicated by the MGMT LED. The message is part of the 6816/6808 download protection. You still have minimal system functionality available to download the proper image. To reset the LED, load the same image in both image spaces and synchronize the switch using the `synchronize` command.

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 6: 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, and Alpine	Maximum number of routes contained in the BGP route table (best case).	400,000
BGP—routes, Summit1i, Summit5i, and Summit48i	Maximum number of routes contained in the BGP route table (best case).	200,000
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 entried per switch.	256
BGP—network statements	Maximum number of network statements per switch.	256

Table 6: Supported limits (continued)

Metric	Description	Limit
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.	3000
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
ESRP—Max instances	Maximum number of ESRP supported VLANs for a single switch.	64
ESRP—Max ESRP groups	Maximum number of ESRP groups within a broadcast domain.	4
ESRP—Max VLANs in a single ESRP domain – Summit “i” series and 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 “i” series and Alpine	Maximum number of routes that can be tracked by ESRP.	4
ESRP—Route-track entries, BlackDiamond	Maximum number of routes that can be tracked by ESRP.	4
ESRP—Max VLAN tracks	Maximum numbers of VLAN tracks per VLAN.	1
FDB—Max L2/L3 entries – MSM64i with “i” series I/O modules, Summit7i, Alpine 3808/3804	Maximum number of MAC addresses/IP host routes for the MSM64i, Alpine 3808, and Summit7i.	256,000
FDB—Max L2/L3 entries – Summit1i, Summit5i, and Summit48i	Maximum number of MAC addresses/IP host routes for the Summit1i, Summit5i, and Summit48i.	128,000
FDB—Max L2/L3 entries for non-“i” series BlackDiamond I/O modules.	Maximum number of MAC addresses/IP host routes for the G4X, G6X, F32T, and F32F.	32,000
Flow Redirection—Max 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—Max 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—Max 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	Maximum number of IP routes used in route sharing calculations. This includes static routes and OSPF ECMP.	12

Table 6: Supported limits (continued)

Metric	Description	Limit
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
Logged Messages	Maximum number of messages logged locally on the system.	1000
MAC-based VLANs—MAC addresses	Maximum number of MAC addresses that can be downloaded to the switch when using MAC-based VLANs.	7000
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, and Alpine 3808/3804	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, and Alpine 3808/3804	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, and Summit48i	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, and Summit48i	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.	40
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

Table 6: Supported limits (continued)

Metric	Description	Limit
OSPF adjacencies—Summit1i, Summit5i, Summit48i, Summit48si	Maximum number of OSPF adjacencies on a switch with 128 MB memory.	200
OSPF adjacencies—Summit7i, Alpine, BlackDiamond	Maximum number of OSPF adjacencies on a switch with 256 MB memory.	300
Policy Based Routing	Maximum number of policy based routes that can be stored on a switch.	64
Protocol-sensitive VLANs—active protocol filters	The number if simultaneously active protocol filters in the switch.	15 for "*i*" series switch products; 7 otherwise
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
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—Max number of simultaneous sessions	For Transparent and Translational and GoGo modes respectively.	500,000/500,000/ unlimited
SLB—Max number of VIPs	For Transparent and Translational and GoGo modes respectively.	1000/1000/unlimited
SLB—Max number of Pools	For Transparent and Translational (does not apply to GoGo mode)	256/256
SLB—Max number of Nodes per Pool	For Transparent and Translational (does not apply to GoGo mode)	256/256
SLB—Max 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—Max STPDs, Summit	Maximum number of Spanning Tree Domains.	128
Spanning Tree—Max STPDs, Alpine	Maximum number of Spanning Tree Domains.	256
Spanning Tree—Max STPDs, BlackDiamond	Maximum number of Spanning Tree Domains.	512
Spanning Tree—Minimum STPDs	Minimum number of Spanning Tree Domains.	1
Spanning Tree—Max 802.1d domains	Maximum number of 802.1d domains per port.	1
Spanning Tree—Max 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

Table 6: Supported limits (continued)

Metric	Description	Limit
EMISTP & PVST+ — Max domains, Summit	Maximum number of EMISTP and PVST+ domains.	128
EMISTP & PVST+ — Max domains, Alpine	Maximum number of EMISTP and PVST+ domains.	256
EMISTP & PVST+ — Max domains, BlackDiamond	Maximum number of EMISTP and PVST+ domains.	512
EMISTP & PVST+ — Max ports	Maximum number of EMISTP and PVST+ ports.	4096
EMISTP & PVST+ — Max domains per port, Summit	Maximum number of EMISTP and PVST+ domains that can be configured per port.	128
EMISTP & PVST+ — Max domains per port, Alpine	Maximum number of EMISTP and PVST+ domains that can be configured per port.	256
EMISTP & PVST+ — Max domains per port, BlackDiamond	Maximum number of EMISTP and PVST+ domains that can be configured per port.	512
Static MAC FDB entries—Summit “i” series, Alpine, and BlackDiamond	Maximum number of permanent MAC entries configured into the FDB.	1024
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
VLANs—Summit “i”-series and Alpine	Includes all VLANs plus sub VLANs, super VLANs, etc.	3000
VLANs—BlackDiamond 6816	Includes all VLANs plus sub VLANs, super VLANs, etc.	600
VLANs—BlackDiamond	Includes all VLANs plus sub VLANs, super VLANs, etc.	3000 in an all “i”-series system. 1024 in a mixed “i”-series/non “i”-series system
VRRP—Maximum VRIDs	Maximum numbers of unique VRID numbers per switch.	4
VRRP—Max VRIDs/switch	Maximum numbers of VRIDs per switch.	64
VRRP—Max VRIDs/VLAN	Maximum numbers of VRIDs per VLAN.	4
VRRP—Max ping tracks	Maximum numbers of ping tracks per VLAN.	4
VRRP—Max iproute tracks	Maximum numbers of iproute tracks per VLAN.	4
VRRP—Max VLAN tracks	Maximum numbers 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 appearing in parentheses are for internal reference and can be ignored.

This chapter contains the following sections:

Clarifications and Known Behaviors

Following are the clarifications and known behaviors in ExtremeWare 6.2.2.

System Related – All Systems

CAUTION

In order for configuration changes to be retained through a switch power cycle or reboot, you must issue a 'save' command. For more information on the 'save' command, refer to the ExtremeWare Software User Guide*.*

Port Mirroring

When a 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).

Software Controlled Redundant Port

You can configure only one redundant port for each primary port.

You cannot configure hardware redundant ports (such as ports 49 and 50 on a Summit48i) as software controlled redundant ports.

Software redundant ports are for use only to peer-to-peer connectivity. They cannot be used to dual-home a port.

Contrary to the information in the *ExtremeWare Software User Guide*, Software Controlled Redundant Port *is* supported on 1000-BaseT ports.

Setting Auto-negotiation Off on a Gigabit Port

When connecting to a device that does not support 802.3z auto-negotiation, it is necessary to 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 include specifying the duplex mode. For example the command:

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

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

Flow Control

Flow control is fully supported only on Gigabit Ethernet ports. Gigabit ports both advertise support and respond to pause frames. 10/100 Mbps Ethernet ports also respond to pause frames, but do not advertise support. Neither 10/100 Mbps or Gigabit Ethernet ports initiate pause frames.

Flow Control is enabled or disabled as part of auto-negotiation. If auto-negotiation is set to off, flow control is disabled. When auto-negotiation is turned on, flow control is enabled. (2815).

Configure Sys-Recovery Level Command

The `configure sys-recovery-level` command monitors 2 tasks for the “critical” level software exceptions – tBGTask and tNetTask.

System Logging

By default, log entries of “warning” and “critical” levels are preserved in the log even after a reboot. Issuing a “`clear log`” command will not remove these static entries. Issuing a “`clear log static`” command will remove all entries of all levels and clear the “ERR” LED on the master MSM module of the BlackDiamond switch (2840).

Enabled IdleTimeouts and Console Connections

If the IdleTimeout feature is enabled, and a telnet session that becomes “timed-out”, a subsequent telnet to the box will be successful but will result in a pause or “hang” an existing direct serial console connection. If the subsequent telnet session is terminated, the console port will resume normal function and subsequent telnet sessions will work correctly (5094).

Configuring the IdleTimeouts Interval

In 6.2.0, the ability to configure the time out interval was added through a new `configure idletimeouts <number>` command. The value was specified in seconds, but did not work correctly. In 6.2.1 the command has been changed to specify the interval in minutes. The range for the interval is 1-240 minutes (1-7SY4D).

The Admin Account

For security reasons, you should not attempt to delete the default “admin” user account. If you delete the default account, it will be automatically restored, with no password, after you download a configuration. Therefore, to ensure security you should change the password on the admin account, but not delete it. The changed password will remain intact through configuration uploads and downloads (1-C5S7B).

If you *must* delete the default admin account, you must first create another administrator-level account before you delete the default admin account. You will need to remember to manually delete the default admin account again every time you download a configuration.

Xmodem Downloads

Though not performed under normal circumstances, an ExtremeWare image can be downloaded using Xmodem through the BootROM menu. Listed below are issues associated with Xmodem download.

Table 7: Xmodem downloads

Extreme Switch Platform	Xmodem download through BootROM
All Summit switches	No issues
BlackDiamond switch	Remove 2nd MSM first

To Xmodem download an image to an MSM using the BootROM menu, remove the second MSM from the BlackDiamond switch prior to beginning the operation (4936).

Show Memory Output

On some systems, the `show memory [detail]` command may show the cumulative memory allocation field as negative (9010).

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).

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).

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:

- CPU-TRANSMIT-PRIORITY
- Summit-chipset module support

- Installed ATM, MPLS, ARM, and PoS modules
- Mirroring
- Dynamic 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).

System Related – BlackDiamond Switch

BlackDiamond 6804 Module Support

The BlackDiamond 6804 does not support the MPLS, ATM, ARM or PoS modules.

Reboot Slave MSM64i After Using synchronize Command

When you synchronize the MSM64i modules, reboot the slave MSM64i so that if it becomes the master MSM64i it uses the synchronized software and configuration (PD2-65213801, PD2-65071101, 1-5E9O9).

Using 110v Power on a BlackDiamond Switch

If 110-volt power is supplied, some BlackDiamond I/O modules might not power up. The MSM64i performs power calculations and powers up the maximum number of I/O modules from left (slot 1) to right (slot 8). A module is skipped if that module is not within the power budget. Using 110 volts, only four modules will typically be powered on (4877).

Enabled IdleTimeouts and Multiple BlackDiamond Console Connections

The idletimeouts feature should not be enabled if serial ports from both MSMs in a two MSM configuration are used for console connections. If the idletimeouts feature is enabled in this scenario, console sessions will not be re-established correctly (5093).

Modem Port on MSMs

The lower 9-pin serial port labeled as “modem” on the MSM does not allow any connectivity to the device at this time. The upper 9-pin console ports of both the primary and secondary MSM can be used as console or modem connection (5179).

Hot Removal of an I/O Module with Traffic

If a BlackDiamond I/O module is removed during traffic flow to the module, several error messages may be written to the log immediately following. These messages should cease to occur after about 10 seconds. Under this circumstance, the error messages can be safely ignored. The error messages may contain one or more of the following (5160, 5082):

```
04/13/1999 17:18.46 <DBUG:KERN> killPacket: HW pqmWaitRx failed

04/13/1999 17:18.46 <DBUG:KERN> pqmWaitKill failed. Card 1 is removed.
```

Removal/Insertion of an I/O Module

The action of inserting or removing a BlackDiamond I/O module should be completed in a reasonable time frame. Be sure to remove or insert the module completely and to avoid partial insertion or connection of backplane connectors (7455).

Removal/Insertion of an MSM

The action of inserting or removing a BlackDiamond MSM will report the following message under certain circumstances. This message can be safely ignored (8547).

```
04/27/2000 12:39.37 <WARN:KERN> ngRxFirst failed WTX1 - (1, eeeeeeee, ffff)
```

Extended Diagnostics

Running the CLI `run diags 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 CLI `run diagnostics normal` command.

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

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

MSM Mismatch on Cold Start

If MSM-A is running 6.1.9 or 6.2.1, and MSM-B is running 6.1.8 or earlier, or 6.2.0, on a cold start with both present, MSM-B will come up faster than MSM-A and will be the master. This is due to the fastpost diagnostic that is run with 6.1.9 or 6.2.1, but that is not run in 6.2.0 or 6.1.8 or earlier (1-841CL).

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).

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).

Normal or Extended Diagnostics on BlackDiamond 6816

The BlackDiamond 6816 only supports the packet-memory option of the run diagnostics command. Users attempting to run either normal or extended diagnostics will receive the following message:

```
Warning: run-time diagnostics is not supported currently on the 6816.
```

Sync of Configurations

When you hot add a slave MSM, the slave will automatically do a sync to bring the master’s images and configurations over to the slave. However, if one of the configurations on the master MSM is empty, the sync process will not overwrite the corresponding configuration on the slave. If the configuration on

the slave MSM is an older configuration, this can cause problems if the switch is rebooted using the outdated configuration (1-AJP7P).

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).

System Related – Alpine Switches

Configuring Slots for the GM-4Xi and GM-4SXi

On the Alpine 3808 and 3804 switches, the only configurable option for The Alpine 1000BaseX I/O modules is the “GM-4Xi” option. When using EPICenter to manage the switch, EPICenter will display a slot mismatch for the GM-4SXi modules when configured as a GM-4Xi. The GM-4SXi will be fully operational and recognized as a “GM-4Xi” for the configured type (9884).

System Related – Summit Switches

Summit 48i 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).

The Summit 48i is currently not able to detect a single fiber strand signal loss due to the HW based Auto Negotiation parameters (10995).

Summit Stackables and SNMP results for Power Sources

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

Summit 48Si MIB value for Input Power Voltage

On the Summit 48Si, 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)

Don’t Use the Encrypted Option from the CLI

There is an option available in the CLI for encrypting a password in commands that specify access or authentication. This includes commands to create and configure accounts, to set the shared secret for RADIUS or TACACS+, for setting the SNMP community strings, for access to various services related to

SLB, and others. *Do not use the encrypted option in these commands*. It is for use only by the switch when uploading and downloading an ASCII configuration file, so that passwords are not indicated in clear text within the configuration file (4229, 4719).

CLI Parser Limitation

The CLI parser is limited to 200 characters, including spaces and tabs. If the number of characters exceeds 200, the switch will return a "stack overflow" error. While the excess characters are clipped, the first 200 characters are correctly processed.

"show iproute" Command

The "show iproute" display has a special flag for routes that are active and in use, these routes are preceded by a "*" 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.

Cosmetic PING Errors

When a ping is unsuccessful, the initially reported number of transmit frames is four, but in actuality the switch will continue to try beyond the four frames. Accurate statistics are reported after hitting a carriage return to terminate the ping function (5132).

When a ping is redirected, the statistics for the last packet received is reported as lost but in fact the ping was successful (5170).

If during the execution of a PING command, the switch receives any ICMP messages that are not an echo reply (e.g. IDRP, Time to Live expired, destination unreachable); an error message is displayed on the console. The error message can be safely ignored (2082).

Cosmetic Configuration Download Warnings

During the execution of the ASCII configuration file during the download configuration process, warning messages may appear when attached to the console port. If you scroll back to review these warnings, the indications are harmless and the desired configuration should have taken place (4931).

"Interrupt messages lost" message

For the BlackDiamond switch, an error message might display to the screen if a command or routing protocol processing requires significant processing time. The error message can be safely ignored (3427). The error message will resemble:

```
0xXXXXXXX (tExcTask): XX messages from interrupt level lost
```

Console Appears Locked after Telnet Attempt

If you telnet to an unresponsive device from the CLI, the console may appear to be locked or frozen. Pressing the <ctrl>] (control and right bracket) keys simultaneously will close the frozen telnet session (4557).

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 port mgmt config (8604).

Auto Negotiation and 1000BaseT Ports

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

Switching and VLANs

This section describes issues associated directly with Layer 2 switching and VLANs.

Management Port IP Address

Do not assign an in-band IP address to the management port VLAN. The management port VLAN is an out-of-band VLAN, so if it is assigned an in-band IP address (an address where the source and destination are in the same subnet), the switch will treat it as a normal VLAN and attempt to route traffic through it (14426).

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 6.2.0 switch. In v 6.2.1 the default value has been changed back to 300 seconds. However, when upgrading from 6.2.0 to 6.2.1, the default value will remain and 1800 seconds. For upgrades from earlier versions of 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).

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.

Modifying the Protocol “IP”

If you wish to modify filters associated with the pre-defined “IP” protocol, use the full syntax of the command. For example “`config ip add ...`” will produce an error message but the command “`config protocol ip add ...`” will work correctly (2296).

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).

GVRP/GARP

GVRP is currently not supported in EW 6.1.x and above software.

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.

Maximum Number of VLANs Supported

The maximum number of VLANs supported on the BlackDiamond, Alpine, and Summit "*i*"-series switches is now 3000. To configure more than 1024 VLANs, the CPU-transmit-priority level must be set to "normal". The CPU transmit priority is set to "high" by default to control the priority in which packets are transmitted from the switch in the event that lower priority queues are congested. This mechanism uses internal resources and limits the number of VLANs that can be configured on a switch. The following CLI command must be used to set the CPU-transmit priority:

```
config cpu-transmit-priority [high | normal]
```

To view the configured CPU-transmit priority, use the following command:

```
show switch
```

Note that the switch must be rebooted for this change to take effect. The default setting for the CPU-transmit priority is "high" (7120).

If non-"*i*" series I/O modules are installed in a BlackDiamond Chassis, the maximum number of VLANs supported will be 1024 (8908).

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

Round Robin Load Sharing. If a port in a round robin load share group is removed, the traffic that was being transmitted on that li*n*k 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 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 when the system is rebooted (8589).

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 Port Configuration. All the ports in a load sharing group must have the same exact configuration, including auto negotiation on/off, duplex half/full, ESRP host attach (don't-count), and so on.

Load Sharing and Software Controlled Redundant Port. When both Software Controlled Redundant Port and Load Sharing are configured, the following behaviors will apply.

For fail-over to the redundant group of load shared ports:

- If the primary group of load shared ports is active, and one or more links in that group fail, the entire load shared group of ports will fail-over to their corresponding software controlled redundant ports configured as a redundant group of load shared ports (provided that enough links can be established on the redundant group of load shared ports to be greater than the number of active links remaining in the primary load shared group of ports).

For fail-back to the primary group of load shared ports:

- If one link in the redundant group of load shared ports fails, and its corresponding primary port is up, the entire group of load shared ports will fail-back to the primary load shared group of ports.
- If one link in the redundant group of load shared ports fails, and the corresponding primary port is down, there will be no fail-back to the primary group of load shared ports.
- It is possible for the primary group of load shared ports to be up when the number of active primary ports is less than that of the redundant group of load shared ports. For example, assume ports 4, 5, and 6 are configured as a redundant group of load shared ports for the primary group of load shared ports 1, 2, and 3. If primary port 1 fails, the primary group of load shared ports will fail-over to redundant group of load shared ports with ports 4, 5, and 6 up. Then, if primary port 2 fails, the redundant group of load shared ports (ports 4, 5, and 6) will remain up. However, if redundant port 6 fails, the redundant group of load shared ports will fail-back to the primary group of load shared ports because the corresponding primary port (port 3) for redundant port 6 is still up.

Spanning Tree

STP not Supported with ESRP. Spanning Tree is not supported and should not be attempted in conjunction with ESRP.

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 Default Domain Association. Newly created VLANs are no longer associated with STPD "s0" or any other domain by default.

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).

MAC Security

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

CLI Changes for MAC Security

The command for specifying a MAC address learning limit has been changed in v 6.2.1. In v 6.2.0, the command `config vlan <vlan> add port <port> mac-limit [no-limit | <number>]` was used to set a limit on the number of entries that could be learned. In 6.2.1, the mac-limit option has been removed from the `config vlan` command. Instead, a MAC address learning limit can be set using the following command:

```
configure port <port-list> vlan <vlan> [limit-learning <number> | unlimited-learning]
```

If you do a direct upgrade from 6.2.0 to 6.2.1, the 6.2.0 command will be converted to the 6.2.1 command. However, a 6.2.0 configuration that is downloaded to a switch running the 6.2.1 image will not automatically be converted, and will produce syntax errors related to the MAC-limit configuration (1-85QDI).

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

Access Lists on BlackDiamond I/O modules

Currently, access lists function only on i-series I/O modules and do not function on the G4X, G6X, F32T and F32F I/O modules.

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.

VLAN QoS Between I/O BlackDiamond Modules

When using VLAN QoS on a tagged VLAN between i-series I/O modules and non i-series I/O modules (G4X, G6X, F32T, and F32F), the "`show ports qosmonitor`" will display the active ports between the new and existing I/O modules as using different queues (7116).

MAC QoS

Broadcast MAC QoS does not take effect on non-"*i*" series I/O modules on a BlackDiamond. If an FDB entry is created with a broadcast MAC address assigned to a QoS Profile, the entry will be ignored against that QoS Profile on non-"*i*" series I/O modules (8841).

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

Configuring access lists with large intervals (greater than 10) between precedence values may result in very long delays (several minutes) for each add transaction. During this time, you will not have access to the session from which you issued the configuration command. In addition, rebooting or downloading a configuration with large precedence intervals may be very slow (on the order of 20 minutes or more). To avoid this problem, configure ACL precedence with an interval of less than 5 between each rule (1-B6F48).

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

When monitoring QoS, do not use the `show port qos` and `enable qosmonitor` commands on the same port at the same time (PD2-64202681).

Bi-Directional Rate Shaping

1000BaseT Ports as Loopback Ports

If the loopback port for bi-directional rate shaping configurations is configured on 1000BaseT 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.

EAPS

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.

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

ESRP

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 revision of ExtremeWare. If it becomes necessary to mix ExtremeWare revisions, do not use any of the new ESRP features associated with ExtremeWare 6.1. These include route tracking and the ability to modify the election algorithm.

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).

Ensure that EDP is Enabled

The Extreme Discovery Protocol must be enabled on the ports involved with ESRP in order to function correctly. By default EDP is enabled on all ports. To verify this, use the command "`show port`

<portlist> info". To enable EDP on a port, use the command "enable edp ports <portlist>" (4072).

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).

ESRP and SLB Health Check

In some situations when a ESRP transition occurs and SLB health checking is enabled, the new ESRP master may send out SLB ping health checks, which may cause a very brief broadcast storm on the network (10-20 msecs). Disabling the SLB health checks will prevent this problem (1-9VHPA, 1-9VHOP).

Gigabit Port Restart

When an ESRP master transitions to a slave, it does not send out a port restart on a gigabit port. It sends out the port restart only after it transitions back to master. This can cause a brief set of transitions between the master and save switches (1-9W4SH, 1-9W4R9).

ESRP Ping Tracking

The ESRP Ping Tracking option cannot be configured to ping an IP address within an ESRP VLAN subnet. It should be configured on some other normal VLAN (across the router boundary) (1-C5S6U).

IP Unicast Routing

VLAN Aggregation

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).

Multinetting

Multinetting and IP Multicast Routing. Combining any type of IP multicast routing on VLANs that are also part of an IP multinetted group is not supported (4418).

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 CLI "`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.

IP Multicast Routing and Snooping

Listed below are issues specific to running IP Multicast routing using PIMv2, DVMRP or IGMP Snooping of IP Multicast traffic.

Cisco Interoperation

NOTE

For proper Cisco interoperation, you must run 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).

IGMP & IGMP Snooping with IP Unicast and Multicast Routing

IGMP snooping and IGMP must be enabled when unicast IP routing or multicast routing is configured on the switch. By default, both IGMP and IGMP snooping are enabled. You can check this using the `show ipconfig` command (5112).

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

RADIUS

When RADIUS authentication is configured on a BlackDiamond switch, upon reboot, the user 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."
```

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.

Server Load Balancing

Default Ping Health Checking

For Transparent and Translational modes, the layer 3 PING is enabled for all members of a pool when it is defined. If a server is configured not to respond to ICMP Echo Requests, the server will be marked "down" after the first ping check interval of 30 seconds. The ping health checking can be disabled using the command:

```
disable slb node {all | <ipaddress>} ping-check
```

Server Load Balancing with 3DNS

3DNS is used as a global load balancing and site redundancy tool. Additional information concerning individual server health and performance can be gathered by 3DNS from the SLB services within the Extreme switch for more granular and accurate decision making by the 3DNS device. These additional

functions apply when using Transparent or Translational modes. To enable responses to F5's 3DNS i_query requests from Extreme's SLB services, use the command:

```
enable slb 3dns iquery-client
```

To see what 3DNS devices are currently communicating with the SLB enabled switch, use the command:

```
show slb 3dns members
```

To disable responses to 3DNS queries, use the command:

```
disable slb 3dns iquery-client
```

The SLB enabled switch responds to directed queries from 3DNS. To direct 3DNS queries to the switch, you add a "Big/IP" device to the 3DNS configuration. Encrypted communications with 3DNS is currently not supported. These functions were tested with 3DNS v2.x and should function correctly with v3.x.

Web Cache Redirection/Policy Based Routing

Health Checking

Under very high sustained loads a Web Cache Redirect may fail and a cache server is set to the "down" state and then brought 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 may 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
```

The FDB table will time out before the IPARP table on the ports connected to the cache servers. To work around this configure the switch to have a higher FDB time-out than the IPARP time-out.

An ICMP PING of the next hop address is turned on by default and cannot be disabled.

VLAN boundary

Web Cache Redirection traffic must come in on an "*i*"-series switch running version 6.1 or better software. Traffic that satisfies a flow redirection must otherwise have been forwarded at layer 3 (packets must cross a VLAN boundary). For example, in a Cache Redirection application the client traffic and the ultimate destination they wish to go to needs to cross a VLAN boundary within the switch, however the caches themselves may reside on the client VLAN or any VLAN on the switch. In instances where the clients and servers belong to the same subnet, the functionality can still utilized by using the proxy ARP functionality in the switch with minimal configuration changes to clients or servers.

WCR and SLB on the Same Switch

When configuring switches to use SLB and WCR simultaneously, ensure that no overlapping layer 4 IP ports exist in the configuration. TCP/UDP ports must be completely independent for WCR and SLB parameters. In this configuration, a request to a cache box cannot initiate a request for information from a SLB VIP as this would violate the overlap of L4 ports.

NOTE

Extreme Networks strongly recommends running SLB and WCR on separate switches.

Precedence of Flow Redirection Rules

Multiple flow redirection rules can overlap in making a redirection decision. In these cases, precedence is determined by “best match” where the most specific redirection rule that satisfies the criteria will win. The criteria for best match is determined in the following order:

- Destination IP address / mask
- Destination IP Port or Source IP port
- Source IP address / mask

In general, the following rules apply:

- If a flow with a comparatively better matching mask on an IP address satisfies the content of a packet, that flow will be observed.
- If one flow redirection rule contains 'any' as an L4 protocol and a second flow redirection rule contains explicit L4 port information, the second will be observed if the packet contains matching L4 information.
- If one flow has a comparatively better match on source information and a second flow has comparatively better match on destination information then the rule with the better match on the destination information will be selected.

For example, in the following 2 cases, the rule with the best match (using the above criteria) is the rule that is selected.

Table 8: Flow rule example 1

Rule #	Destination IP Address	Destination IP Port	Source IP Address	Priority Selection
1	192.0.0.0/8	80	ANY	1
2	192.168.0.0/16	ANY	ANY	2

In this case, Rule 1 is the rule with the best match as it contains an explicit Destination IP Port even though the mask for the Destination IP Address is less specific.

Table 9: Flow rule example 2

Rule #	Destination IP Address	Destination IP Port	Source IP Address	Priority Selection
1	192.168.2.0/24	80	ANY	2
2	192.168.0.0/16	ANY	10.10.10.0/24	4
3	192.168.2.0/24	ANY	10.10.0.0/16	3
4	192.168.2.0/24	80	10.10.0.0/16	1

In this case, Rule 4 is the rule with the best match as it again contains an explicit Destination IP Port.

NetFlow

If a flow record filter is configured on one port with type “match-all-flows” you cannot configure the same flow filter on any other port (1-7G1D8).

WEB Management - VISTA

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

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 100/1000BaseTX ports, the ifMauAutoNegAdminStatus can only be disabled if the ifMauDefaultType is set to a speed of 100Mbps. For 10/100BaseTX 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 will be printed to the system log (11718). This access to the MIB has no adverse effects to any protocol stability (i.e., ESRP, OSPF, BGP).

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 Server

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

DLCS

DLCS is only supported on “i” series modules (8389).

Virtual Chassis

The Virtual Chassis is not supported in ExtremeWare 6.0 or higher.

Documentation

IPX Support

The *ExtremeWare 6.2.1 Software User Guide Rev 4* incorrectly states that basic functionality includes support for IPX routing. The error is corrected in the *ExtremeWare 6.2.1 Software User Guide Rev 5*.

Issues Resolved in ExtremeWare 6.2.2b18

The following issues were resolved in ExtremeWare 6.2.2b18. Numbers in parentheses are for internal use and can be ignored.

General

The `show diagnostics packet-memory slot` command now correctly records packet scanning and mapping defects (1-FHG21).

Exported OSPF, RIP, PIM, and route map configurations are now interpreted correctly when upgrading from ExtremeWare 6.1.9 to ExtremeWare 6.2.1 (1-AHM25).

Software redundant 1000BASE-T ports now operate correctly with load sharing (1-AWZMN).

The show log command output no longer truncates the following error message (1-8EFVX):

```
<INFO:SYST> Checksum error (ffff in EEPROM vs 0 calculation) in card 1 memory defect
information, restore defau
```

Configurations that create large ACL or static FDB port lists now format the saved configuration so that a large list of ports no longer generates a “stack overflow when downloading” error (1-CND21).

If you disable smart redundancy, active backup links no longer fail over to the primary link (1-D3XCT).

Multiple simultaneous ARP entries are now handled correctly (1-DHY7W).

When you change the system time, ARP aging is computed properly (1-E3MRP).

Summit

The PSU status LEDs now operate as described in the *Consolidated Hardware Guide* (1-BMDGW, 12960).

Attempting to load BootROM 7.6 on a system with less than 128 MB of memory now generates an error message. In addition, attempting to run ExtremeWare 6.2.2 on system with less then 128 MB of memory now generates an error message and puts the system into limited mode (1-EIMNT).

Alpine

The `show diagnostics` command output now correctly displays information on the Alpine 3804 CPU (1-8TS79).

The `show slot` command output now correctly displays “FM32T” instead of “FM32” (1-AMDQP).

BlackDiamond

Enabling IGMP snooping on a BlackDiamond 6816 no longer generates an error message (1-BP5CD).

The power supply LEDs now consistently operate as described in the *Consolidated Hardware Guide* (1-9RC85).

The LEDs on the F48Ti now operate correctly when diagnostics take the card offline (1-9JX6Q).

Executing commands while the chassis is in “limited commands” mode no longer causes the MSM64i to crash (1-FAXIE).

The BlackDiamond 6816 does not power up using non iPower power supplies (1-FTW7L).

The `clear slot` command no longer causes the switch to reboot when a port on the I/O module belongs to 3,000 VLANs (1-FGFK5).

IPX

When a SAP entry was advertised with a network ID of zero, it was not added to the IPX services table. A value of zero indicates a local network and is now added with the local network number of the receiving interface (1-BTWKQ, 1-BTWK9).

VLANs

Two domains sharing a single VLAN now each have a unique VLAN tag (1-6EJM5).

If you create a VLAN with “mgmt” in the name, you can now delete that VLAN (1-EEUPY).

Mirroring

Enabling jumbo frames on mirroring ports without jumbo frames enabled on the monitor port now generates an error. To properly enable jumbo frames with mirroring, first enable jumbo frames on the monitor port and mirroring ports, then enable mirroring (1-8SJAL).

Load Sharing

If the master port is down, a restart now automatically restarts the next operational master (1-994B4).

If the master port is down, the FDB now correctly learns new MAC addresses (2-HAQAD, 2-H8KID, 2-H182X).

Spanning Tree

Do not configure a topology change time less than 15 seconds (PD2-64997595).

ESRP

On a Gigabit Ethernet port, port restart now correctly flushes the layer 2 ESRP-aware switch’s FDB (1-8S74I, 1-994AP, 11738).

EDP no longer floods packets when 3,000 or more VLANs are configured (1-6EH2H).

ESRP now correctly waits for the specified timeout before making state changes for a VLAN (1-E3TCC, 1-E3TE3).

VRRP

The performance of VRRP with large groups of VRIDs has been improved (1-5GS3X).

Attempting to configure two VLANs with the same port and VRID now generates an error message (1-F7C61, 1-DO0QD).

EMISTP

If you configure dot1d mode on tagged ports and delete the ports from the VLAN, the following error message is generated in the output for the `show stpd ports` command (1-81BID):

```
WARNING: s2 has dot1D port with no default vlan.
BPDU received on this port will not be flooded to other ports in the domain.
```

IP Multicast Routing

PIM is now correctly updated when the RP changes (1-CN8F8, 1-CN8FD).

Unconfiguring PIM now correctly sets the Register-Rate-Limit-Interval to the default (1-CMY1R).

BGP

The `show bgp neighbor <ip address> transmitted-routes all` command no longer crashes the sytem (1-AIDPH).

The `config debug trace bgp-misc` output now correctly displays a “Deleting Network” message (1-BZC4Z).

The transmitted route statistics for Aggregated Routes now displays the correct route statistics (1-CCDZT).

The output policy can no longer set the next hop while advertising the route to the IGP session (1-EP9V1).

The show bgp neighbor command output was modified to change “EBGP” and “IBGP” to “EGP” and IGP” (2-H09TL).

OSPF

OSPF SPF no longer make a directly attached network unreachable if that network moves multi-hops away (1-9D242, 1-9D23T).

The error message when VLAN deletion fails is more understandable (1-8SCKX).

OSPF now supports a minimum cost of 1, instead of 0 (1-8JFIL).

The `show ospf lsdb stat` command now operates correctly (1-81GWT).

Saving the configuration now correctly saves the OSPF filter (1-FQUPP).

EAPS

You now receive the following error message when attempting to add more than 64 EAPS domains (1-87LQ3):

```
<CRIT:EAPS> eaps.c 2639: Error! Reached maximum limit of EAPS instances
```

The EAPS Master Node now fails immediately when you disable a Primary or Secondary port (1-87LQI).

When the Primary port for an EAPS Master goes down, the seconday port comes up before the Primary port is deleted from the list of active ports, so that OSPF no longer re-converges (1-90L44).

You can no longer run EAPS without enabling EDP. Disabling EDP on an active EAPS Domain generates an error message and causes EAPS to be idle. You cannot configure a port that has EDP disabled as an EAPS ring port (1-9YM15).

NetFlow

You can now configure multiple ports with the `match-all-flows` type (1-7G1D8).

NAT

The NAT and UDP timeout values now work correctly (1-GAWX4, 1-5AJ9H).

SNMP

You can no longer use SNMP to rename or change the tag setting for the “default,” “Mgmt,” or “MacVlanDiscover” VLANs (1-ADUKD).

You can now create, configure, and delete a VLAN with a tag of 1 (16394).

You can now use SNMP to reset all QoS settings to the factory default (1-7YDFC).

The System ID now displays correctly (1-5AOFH, 14113).

The QoS profile priority is now displayed consistently in both the CLI and the SNMP MIB (1-7XD2X).

The extremeEdpPortIfIndex now records the correct value for port 4:1 (1-8SMO8).

You can now correctly configure WDMi modules in slots 4 and 8 using SNMP (1-CHMKY).

Flow Redirection

You can now enter any size netmask when creating a flow rule. If you configure subnet-based forwarding, all of the flow rules must use the same size netmask (1-64K4J, 11464, 16123).

Issues Resolved in ExtremeWare 6.2.1b27

The following issues were resolved in ExtremeWare 6.2.1b27. Numbers in parentheses are for internal use and can be ignored.

General

The memory mapping feature now properly remaps SRAM errors (1-EER8D).

If you download a configuration with the system watchdog timer disabled, save the configuration, and reboot the switch, the system watchdog timer now remains disabled (1-94E8I, 16057).

System block memory is now checked for and protected against double free conditions (1-ASJW0, 1-B98WH, 1-D3TKO). ExtremeWare now checks memory and protects it against conditions that cause a tnetTask EPC in the ARP functions.

If you configure the management port with an IP address, a link transition no longer causes the IP FDB to be cleared (13787).

If you delete a duplicate precedence IP permit access rule, you can now successfully create another permit rule (13525).

BlackDiamond

The switch no longer detects or reports false Interrupt Service Requests on BlackDiamond modules. (1-EV2T6).

Hot-swapping an "*i*" series I/O module no longer generates an error message (8840).

Port- and VLAN-based mirroring now functions correctly (9892, 10548, 10642, 11102, 11122, 11123, 11350, 11497).

Alpine

Configuring load sharing across modules with four or more modules no longer creates a slot mismatch (1-9W4T6).

The `show diagnostics` and `show tech support` commands no longer cause the Alpine 3804 to lose Telnet access. (1-BTXN5).

Network Login

If you log in to a switch using Network Login and change your port, Network Login now correctly authenticates the new port instead of the previous port (1-AWYJ1).

Network Login now successfully times out 20 mintues after the connection is closed (1-EBWGD).

SNMP

HP OpenView now correctly recognizes link up, link down, OSPF, BGP, VRRP, ping, and traceroute traps (1-EY6RQ).

Spanning Tree

In large, concentrated STP environments, STP BPDU's are now sent from the switch in the correct timeframe (1-AKFAL, 1-9FKJD).

IP

If you configure a new IP address for the MGMT port or create a new VLAN with a new IP address, a link state transition no longer clears the IPFDB table (13787). Instead, only the IPFDB entries that egress through the new subnet are cleared.

A high rate of UDP pings (greater than 300/sec) no longer corrupts memory. The corrupted memory could lead to suspended tasks or a system crash (1-BZB3W, 1-CP2ZI).

IP Multicast Routing

When Spanning Tree converges and a port changes from blocking to non-blocking, the IPMC FDB now correctly flushes the old entry and adds the new entry (1-A5LKP).

You can now add a new router port when there is continuous multicast traffic traversing the switch (1-BTWLK).

A software exception no longer occurs when removing an IGMP snooping entry when internal system port tags are not available (1-B0979).

The sender entries installed by PIM now correctly age out (1-CHLZ8).

OSPF

Unplugging the port in a single-port OSPF VLAN no longer causes CPU utilization to spike to 97 percent (1-9FFF1).

Server Load Balancing

The layer 7 service check no longer checks explicitly for the phrase “server ready” from POP3 servers, allowing it to work correctly with all POP3 servers (1-8II6L).

Flow Redirection

You can now enter a netmask greater than /20 when creating a flow rule. (1-DSWIP). We recommend that you only configure netmasks of /16 or greater in a lab environment, as a netmask of /16 uses all 64,000 available IP source addresses.

The source flow no longer changes to the destination flow after rebooting the switch (1-9YEP9).

Disabling the access-list log no longer causes flow redirection to stop working (1-9FFEL).

Disconnecting or reconnecting the flow redirection next hop no longer causes the CPU utilization to spike to 100 percent. This made the switch unresponsive and occasionally caused the switch to reboot (16056).

On a system with a combination of flow rules, populating the IPFDB table with 200,000 entries while constantly modifying health check rules no longer causes a BlackDiamond to become unresponsive (1-CJPVT).

Issues Resolved in ExtremeWare 6.2.1b20

The following issues were resolved in ExtremeWare 6.2.1b20. Numbers in parentheses are for internal use and can be ignored.

General

Cisco Interoperation

The switch did not detect a topology change because the Topology Change Notification (TCN) from a Cisco switch was dropped. Since it failed to detect a topology change, the switch did not fast age or clear the FDB. Cisco's TCN does not include the TLV (type Length Value), unlike the Config BPDU. This caused unpredictable traffic flow. The switch now correctly recognizes a Cisco TCN (1-7F5A9).

Daylight Savings Time

Worldwide daylight-saving-time is now supported (1-6Z6SS, 14239).

If you configure the switch with a timezone and daylight savings time (dst) name, then upload and download the configuration, the dst name is now saved (1-837F2).

DNS Lookup

The nslookup function in ExtremeWare did not correctly handle a multi-answer DNS response. In turn, this meant that ExtremeWare `ping` and `traceroute` commands used with name arguments were unreliable since there was no way to tell whether there were multiple answers to the DNS request. This problem has been fixed (14903, 15097).

When a DNS host lookup fails, this is now logged as a warning rather than as a critical error (1-7NSE9, 13872).

Ping

Executing a ping request from a console session and a Telnet session will no longer fail on the second attempt to start the ping from both sessions (1-60H33).

Closing a Telnet session now terminates a continuous ping request (1-79YGL).

Console Session

The `show vlan dhcp-address-allocation` command no longer crashes a Console session (1-8MT4K).

Auto Negotiation

Disabling auto negotiation on a link between two 1000baseT links no longer brings down the link (12821).

A loopback port cannot be configured with auto-negotiation set as enabled (1-5XHF7, 1-57DUV).

Show Command Output

The show qosprofile command, in specific situations, would continuously scroll through the ports assigned to QoS profile 2. The show qosprofile command now operates correctly (14960, 1-5XHG6, 12550, 14959).

The show switch output no longer shows an extra timezone parenthesis for the timezone section (1-7TC77).

The show vlan stats command no longer stops the port from forwarding traffic (1-8AF1T).

The show ipstats command no longer increments the "Bad Length" packet counter for short packets instead of the "Short Packet" counter (1-5OZ69).

BlackDiamond

If you disable sharing on a switch with no slave MSM64i, when the switch is rebooted the tConsole task no longer crashes (1-5I5K7).

When viewing statistics or configuration of PIM-DM, the BlackDiamond no longer reports a high CPU utilization for the HTTP task and the CLI no longer becomes unresponsive (1-7S1K5, 16430).

The master mgmt port now uploads a configuration with 96 10/100 Mbps ports in the chassis (1-51J9H).

In minimum access mode on a BlackDiamond 6816, the tConsole no longer crashes when the user logs out (1-998LT).

ESRP-aware switches now display the correct port number (1-5VTH5, 1-5VTHF, 1-5VTHL).

If you create more than 73 ACLs, you no longer experience slow ARP response (15899).

Alpine

When configuring and enabling load sharing across I/O modules with the slave port module removed, the CLI no longer sends an incorrect message indicating that the port types in the group are different. It now correctly indicates that the module was not present (1-5NNBY, 1-5NNBT, 1-5NNCX).

The Alpine status LED now changes from blinking green when a critical failure is detected (1-5MP5J).

When upgrading a switch from ExtremeWare 6.1.8b12, load share slave ports on a FM-32Ti module of the newly upgraded switch become active (1-5I5IZ, 1-5LU1U).

An Alpine with load sharing configured no longer crashes if rebooted without the slave module (1-5NNCB, 1-5NFZX).

Transmission collision errors no longer occur on an FM32Ti module installed in slots 5 - 8 if it is the last active slot (15772).

Port mirroring now correctly mirrors additional learned ports (1-53ISR).

Summit

On a Summit48i or Summit48si, heavy traffic no longer locks up the switch (1-5P2N1).

A reboot now properly takes a Gigabit link down, so an attached device recognizes a link transition and proceeds with a clearing of the forwarding database (1-5CIPQ, 1-5CV53, 1-5CV4T).

On a Summit48i, if port 49 or 50 is connected to any other device and then disabled and re-enabled, the switch no longer displays an active LED on port 49r/50r rather than the primary port (15775, 15774).

On a Summit48i, ports 49R and 50R now function correctly if you disable auro-negotiation and connect the ports to switches that do not have redundant hardware ports (1-5I5IK).

FDB

After changing the FDB aging timer value, the age values of entries in the FDB table could become incorrect. This problem has been fixed (16043, 12632).

The `clear fdb` command now works correctly (15616).

The show ipfdb command would display an incorrect MAC address for the destination IP address for that IP FDB entry. This occurs with a blackhole association after multiple link transitions between a host and that destination. The show iparp command displays the correct MAC address for the entry. The show ipfdb command now displays the correct MAC address for the entry (1-81YB2, 1-60H3B).

IP Forwarding

When sending an IP frame with an invalid IP header (but where the IPF entry already exists) the packet was redirected to CPU but was then incorrectly forwarded by CPU. These packets are no longer forwarded (1-5OZ61).

Clearing the IPFDB on a system with 50,000 or more IPFDB entries no longer causes ESRP to flap (11070).

When a port loses connectivity, the IP address is now completely removed from the IP FDB (11700).

SNMP

SNMP vulnerabilities as detailed in CERT Advisory 2002-03 have been fixed (1-9I8DD).

Executing an `snmpwalk` command in certain situations would cause the switch to become unresponsive to any other SNMP requests until the switch was rebooted (12890, 12904).

After a "`clear iparp`" command was issued, which removed the next hop MAC from the IP ARP table, the default route would be used, and an SNMP trap would be sent only to the last trap receiver (14939, 14929, 14932).

The appearance of ports in the dot1dBasePort table is now independent of STP; thus, all Ethernet ports can appear in this table. In addition, only ports with an STP port mode of "dot1d" appear in the dot1dBasePort table. Previously, only ports belonging to the STP domain "s0" would appear (1-6DOIP).

SNMP `get` and `getnext` commands now return the same values for a given MIB variable (1-6432L, 1-60SXJ).

Fixed an SNMP crash that occurred when HP OpenView polled the switch with an invalid OID (1-71SID, 1-6YTJX).

SNMP now correctly reports the configured MTU size values for ports and VLANs (1-7UOUU).

The control of what UDP port number to be used when sending SNMP traps can now be controlled through the CLI and is stored as part of a configuration (4914).

The dot1dBasePortTable now displays all 8 slots in a BlackDiamond switch (6918).

Creating an entry in the extremePortLoadShare2 table with an active port and subsequently destroying the row for the entry no longer results in a software exception (1-57P9C).

The switch now responds correctly to snmpwalk on extremeCpuTaskMaxUtilization.1 (12906).

The linkUp/Down trap OID now fully complies with RFC 1573 (1-5XHGQ).

The ipNetToMedia table now returns the correct value (1-7YHZ0).

Vista

If RADIUS is enabled on a switch, Vista is now fully supported (8887).

The Vista VLAN display now includes loadsharing port 1:1 (1-6433N).

Flow Redirection

When switching from pinging to layer 4 port health checking only one next hop would stay up. When switching back to ping health checking, one next hop would continue to stay down. This problem has been fixed (14879, 14881, 14882).

Disconnecting and reconnecting one next hop no longer causes the next hop to stay down (16363, 16364).

Disabling flow redirection, setting a service check, and enabling flow redirection, no longer causes the switch to crash (14020).

If you enable ping checking, enable a layer 4 health check, and enable ping checking again, the next hop is no longer marked down (14882).

The layer 4 health check maximum timeout is now 9 seconds (1-5BG7K). Thus, the next-hop is taken down immediately when its FDB entry is deleted, or the router interface to get to the next hop goes down. Next hops will be marked up as soon as they pass the configured health check.

Server Load Balancing

Clearing a UDP connection can now be done using the IP or VIP options with a `clear slb connections` command (14954, 14964).

Web Cache Redirection

A VLAN with WCR next-hops can no longer be deleted (15455, 10167).

Removing or disabling a port no longer causes task crashes (16141).

WCR now fully supports jumbo frames (1-5XHGL).

IGMP Snooping

In ExtremeWare 6.1.8, IGMP was always enabled on the management VLAN (mgmt) and there was no way to disable it. It is now disabled for the management VLAN (15643, 13745, 13972).

On Alpine 3804 and 3808 systems, after disabling IGMP snooping and performing a TFTP upload, the configuration file is uploaded correctly so that IGMP is still disabled after a subsequent download (1-65571, 1-6556W).

A port is no longer removed because of an IGMP leave if there is a router detected on the port (1-632YR).

IGMP and EAPS

On an EAPS ring supporting L3 VLANs, with PIM-SM running on each node, the PIM-SM (S,G) entry, as well as IGMP sender entry and IPMC FDB, no longer get deleted on a last-hop node if an IGMP leave comes onto another last-hop node (1-5Q04L).

If there is no routing configured in a pure EAPS Layer 3 ring, the IGMP snooping entry for routers would be missing. This has been fixed (1-5DV8T).

QoS

Rate Shaping

When rate shaping routed traffic on 10/100 ports, rate shaping ports are no longer restricted from belonging to the same block of 8 ports as loopback or normal ports.

Multicast Control Packets can now pass through a rate shaped port on a Layer 2 VLAN (1-7S9PL, 1-7S9OT).

Access Lists

When ACLs in a large ACL configuration were listed in ascending precedence order, the 154th ACL would cause a software exception when configured through a Telnet or Console session. This problem has been fixed (1-7G4MH, 12174, 1-5XHFW).

SSH

Having two simultaneous SSH sessions open no longer causes the SSH task to crash (1-7CFJI, 1-7CFIT).

Once SSH was enabled, the configuration always showed SSH as enabled, even after a disable SSH command. The configuration now correctly shows SSH as enabled or disabled as appropriate (1-7AHWM).

IP Unicast Routing

The default route's gateway is no longer allowed to configure to a network address (16130).

IP Multicast Routing

IP multicast traffic is now supported on a sending IRS port. (15076, 12539).

IP multicast control packets sent to a Layer 2 interface on a switch with IP forwarding enabled could result in a Layer 3 interface on the same switch losing router adjacency. The console could also appear to be responding sporadically to commands when these packets were being transmitted. This problem has been fixed (1-7L3YZ, 1-6XGCK, 13457, 12634).

The “prune sent” messages are no longer shown in the upstream IPMC after a PIM-join is received (15758).

With PIM-SM, after a rendezvous point (RP) failover, the backup RP now correctly sends register stop packets (1-59M8H).

IPX Routing

IPX LHF FDB has been updated to support dynamic learning (1-5Q07E).

IPX and spanning tree now complete fast aging. FDB entries with the "*i*" flag on IPX VLANs could not do fast aging when spanning tree detected a topology change. Any "*i*" flagged entry in the FDB now stays regardless of a topology change (1-6XUAH).

Load Sharing

Fixed a problem with port-based load sharing across modules that would cause traffic to fail to pass after saving the configuration and rebooting (1-5VKC4).

Fixed problem with load sharing configuration across modules that caused switch to roll with the message “Need to translate between vpst’s” (1-60SVS).

ExtremeWare Vista now correctly displays port 1:1 (on a BlackDiamond or Alpine) or port 1 (on a Summit) when the port is configured for loadsharing (1-6433N), 1-6433U).

A static FDB entry configured to a load sharing master port will now be linked correctly to other ports in a load sharing group when the master port fails (1-60SVG).

When loadsharing is enabled on a BlackDiamond and the module is removed or is set “down” by the system health checker, iproute no longer shows the interface in the “up” state (1-7YI0Z, 1-7YI0J).

OSPF Routing

OSPF no longer becomes unstable when the message task (tospfMsgTask) utilization becomes greater than 80% (1-6C396).

Static routes now export correctly to OSPF after an ESRP transition (15958).

Downloading via TFTP an uploaded configuration with a user-configured OSPF cost no longer reset the cost to the default value. The user-configured cost is preserved (1-6BFWA, 1-6BFW1).

Route sharing works correctly for exported static routes (10746).

The ABR now changes the type 7 LSA to type 5 in all cases (13598, 16259).

BGP

When a BGP peer opened a connection to a multihop peer without an IGP path, TCP Connect returns an error. The neighbor becomes active without closing the connection. The other side will accept the connection and move to the open sent state. This causes the other side to reject any retries as it already has a connection in the open sent state, causing a delay in establishing the connection. The connection is now cleared when the neighbor is moving to the active state (1-6B8VB).

ESRP

ESRP will no longer fail over when displaying a long show command output (16318).

ESRP no longer flaps when using bootprelay. This problem occurred when ESRP and bootprelay were using a default route to get to a DHCP server. If the link goes down, though the DHCP server is reachable by ping, there is no DHCP reply (1-5VKBE).

A fully-populated BlackDiamond configured as an ESRP slave no longer transitions to become a second ESRP master upon bootup (1-606HH).

When configuring an ESRP VLAN with HA and RIPv1, the slave ESRP switch no longer generates a spoofing attack message when it receives a RIPv1 broadcast packet across the HA port (1-60H2B).

The switch no longer loses connectivity with certain hosts after an ESRP slave-to-master transition, and ARP requests to the new master are processed properly (1-6DOON).

The `show configuration` command no longer appends the port display string to ESRP port mode configuration commands, and the ESRP configuration commands now work correctly after saving and downloading the configuration (1-74MLY).

Removing and re-inserting the slave MSM64i in a BlackDiamond chassis no longer causes ESRP to transition (11066).

Bidirectional rate shaping no longer causes an ESRP transition (16374).

The switch no longer drops traffic for 20-30 seconds during an ESRP transition (1-8SMML, 1-8SMM1).

When using flow redirection, if a client is on a server ESRP VLAN and the cache server does not have ESRP enabled, the traffic from the client is forwarded to the cache server. The cache server updates its host entry with the ESRP MAC address. When a reply comes from the cache server containing the ESRP MAC address, the packet would be dropped because the cache VLAN did not have ESRP enabled. This continued until the ARP table was cleared, which then sent an ARP request. This ARP request updated the cache server with the proper system MAC address. This ultimately caused a slight delay while accessing web pages. The cache server no longer drops the packet (1-5KZ15).

Domain Members and sub-VLAN Support on Host Attached (HA) Ports

Earlier restrictions that required HA ports and sub-VLANs or Domain-member VLANs to be mutually exclusive have been lifted in 6.2.1.

System Health Checking

With system health checking enabled and set for autorecovery, the link will no longer show active after the module (G8Xi) is removed (1-5ZZ2T).

The config sys-health-check card-down command will no longer bring the neighbor's link down (1-842DL).

Mirroring

Mirroring filters are now added correctly when a configuration is downloaded (1-5ELMT).

In cross-module mirroring, warning is no longer given about mirroring being allowed only within one slot (1-5I8YL).

The monitor port now captures packets after the VLAN tag ID has been changed (1-5K1P5).

Mirroring a blocked Spanning Tree port no longer causes a flood (12786).

NAT

Network Address Translation now works properly with ESRP enabled (1-58YPU).

NAT rules will be preserved across a configuration save and reboot even with more than 50 rules (1-5JE5X).

Output from a `show nat rule` command now works properly when the `clipaging` setting is enabled (1-5JE62).

EAPS

EAPS control messages are no longer dropped if a broadcast storm occurs (12302).

STP

The STP Hold Timer now delays the forwarding of Configuration BDPUs for the correct amount of time (default of 1 second) to prevent topology changes and connectivity loss (1-61LCM).

When a switch was proceeding through the Listening/Learning steps on participating ports, major flooding could occur on a non-root designated port that was transitioning to a blocked state, creating a broadcast storm for a period of a few seconds to several minutes. This problem has been fixed (1-6ZO4P).

VLANs

Adding a disabled port to a VLAN as a loopback port now results in an error indicating that the port is in disabled state (1-5E7LJ, 1-5E7L9).

When you configure 3,000 VLANs, you can now add all ports as active interfaces (9799, 9807, 9629, 9474, 9301, 10357, 10961).

You can no longer rename system-created VLANs (13820, 13817, 13818, 13819).

You can now configure a VLAN with a tag identical to the intenal tag of the VLAN (1-84DOA).

If a port was added to a VLAN as tagged and a restart port, and the configuration was subsequently uploaded using TFTP, upon downloading the configuration the port would be considered in conflict with the untagged port on the Default VLAN. The configuration is now processed correctly and the

port is correctly configured in both the tagged and untagged VLANs with the restart option (1-74MN3, 1-74MMN).

The switch no longer moves MAC addresses from non multinetted VLANs (1-7UOVY, 1-7X90P).

Adding or deleting a port to/from a VLAN no longer causes the IPFDB to be flushed (1-81DE0).

VLAN Aggregation

When using bootprelay and VLAN aggregation, the ARP entry now learns the port number when a client gets a DHCP address (1-6Z6TG, 16379).

When using bootprelay and VLAN aggregation, upon receipt of a DHCP IP address at a client, the associated ARP entry did not bind the physical port number, causing communication problems between subVLANs and unknown destinations (1-81YAF).

Debug Tracing

The config debug-trace ip-forwarding command now captures the IP header (1-579HT).

TACACS

When downloading a configuration file that configures and enables TACACS accounting, there is no longer a delay in processing each line of the configuration after the TACACS accounting configuration command is initialized (1-7TD7B, 1-7TD7U).

Enabling TACACS accounting no longer causes a delay in downloading configurations (1-7TD7B).

Under TACACS, a priv-lvl 15 user granted only “show” commands is no longer prompted to save the configuration when logging out (1-7UOVE, 1-7X907).

NetFlow

After configuring a NetFlow export filter with multiple UDP ports on an IP address, only the first UDP port was being shown as “up” in a `show flowstats detail` display. Now, all ports are correctly shown as up (1-59MXH).

NetFlow filters no longer appear in the `show access-list` display (1-8AOJP).