




ExtremeWare Release Notes

Rev 2

Software Version 6.1.8



Contents

Overview	1
Supported Hardware	1
BlackDiamond Module Support	2
Alpine Module Support	3
Summit Module Support	3
GBIC Support	3
Upgrading from v6.0 to v6.1	4
Upgrading BootROM	4
Upgrading ExtremeWare	4
Upgrading Switches.....	4
Downgrading Switches.....	5
Features Unique to the “ i” Chipset	6
Staying Current	6
New Features in ExtremeWare 6.1	6
General	6
*show tech-support Command	6
Image Filename	7
*show vlan Command	7
Mirroring	7
*Enhancements to Subnet Directed Broadcast Forwarding.....	7
*Packet Forwarding Options for IP interfaces.....	8
BlackDiamond	8
System Health Checking	8
Runtime Diagnostics	9
Disabling G1 Support.....	9

*Multicast Performance Enhancements.....	10
Layer 2 Switching and VLANs	10
Maximum Number of VLANs Increased to 3,000 on all “i” Series	10
*STPD BPDU Tunneling.....	10
Renaming a VLAN	10
VLAN Statistics.....	10
Jumbo Frames	11
vMANs - VPN Services for Metropolitan Area Providers	12
Spanning Tree Rapid Root Failover	13
General IP Functionality	13
IPARP Address Checking.....	13
QoS.....	14
Bi-directional Rate Shaping for Routed VLANs	14
Configuring Bi-Directional Rate Shaping.....	14
Bi-Directional Rate Shaping Limitations.....	15
Bi-Directional Rate Shaping Commands.....	19
Maximum QoS Buffer	19
ESRP	20
*ESRP Multiple Ping Tracking	20
*ESRP Port Restart.....	20
*dont - count Parameter.....	20
ESRP Environment and Diagnostic Tracking.....	20
Increased Number of ESRP Domain Member VLANs	21
IP Unicast Routing.....	21
Route Map Support.....	21
VLAN Aggregation SubVLAN Address Range Checking.....	22
IP FDB Performance.....	22
OSPF	23
*OSPF Point-to-Point Support.....	23
*Configurable OSPF Wait Interval.....	23
*OSPF CLI Display Enhancement	24
OSPF Database Overflow	24
OSPF Password Encryption.....	25
OSPF Passive Interface.....	25
Route Map Support for OSPF Export.....	25

BGP	26
BGP Peer Groups	26
BGP Route Selection.....	27
BGP MD5 Authentication.....	28
BGP Password Encryption.....	28
BGP Route Flap Dampening.....	28
IP Multicast Routing and Snooping	28
Static Rendezvous Points RPs.....	28
PIM Mode Translation	28
IP Multicast Cache Display.....	28
IGMP Snooping.....	29
IPX	29
IPX Routing Design Restrictions Lifted.....	29
Security and Access Policies	30
ICMP ACL Precedence.....	30
Access List Display	30
IPX Routing Access Policies.....	30
BGP and OSPF Route Map Support for Tagging and DSB Accounting.....	31
Server Load Balancing	31
Health check definitions.....	31
GoGo Mode Health Checking.....	32
SLB Global Connection Timeout.....	33
Combined SLB and ESRP Failover.....	33
SLB Pool and VIP Statistics.....	34
SLB Pool Member Configuration.....	34
SLB Proxy Client Persistence.....	34
Web Cache Redirection/Policy Based Routing	34
Health Checks.....	34
Support for 'any' Layer 4 Flows.....	35
Policy-Based Routing with Route Load-Sharing.....	35
SNMP	36
MIB Support.....	36
*SNMP ifDescription Enhancements.....	36
*SNMP ifTable Enhancement.....	36
SNMP ifMib Enhancements.....	36
SNMP ifType Enhancements	36

*SNMP Trap Receiver Changes	37
*RADIUS and TACACS Password Length.....	37
OSPF Traps.....	37
*SNMP dot1dTpFdbTable Enhancements.....	37
Supported Limits	37
Clarifications, Known Behaviors, and Problems	40
System Related – All Systems	40
Setting Autonegotiation off on a Gigabit Port	40
Flow Control.....	40
Config Sys-Recovery Level Command	41
System Logging.....	41
Enabled IdleTimeouts and Console Connections	41
Xmodem Downloads	41
Show Memory Output	41
EDP Packet Length.....	41
TFTP Download of Configuration Files.....	42
System Related – BlackDiamond Switch	42
Using 110v Power on a BlackDiamond Switch.....	42
Enabled IdleTimeouts and Multiple BlackDiamond Console Connections.....	42
Modem Port on MSMs	42
Hot Removal of an I/O Module with Traffic	42
Removal/Insertion of an I/O Module.....	42
Removal/Insertion of an MSM	42
Extended Diagnostics	43
System Related – Alpine Switches	43
Configuring Slots for the GM-4Xi and GM-4SXi.....	43
System Related – Summit Switches.....	43
Summit 48i Redundant PHY	43
Summit Stackables and SNMP results for Power Sources	43
Command Line Interface (CLI).....	43
Don't Use the Encrypted Option When Creating an Account	43
"Show Iproute" Command.....	44
Cosmetic PING Errors.....	44
Cosmetic Configuration Download Warnings.....	44
"Interrupt messages lost" message.....	44

Console Appears Locked after Telnetting	44
Serial and Telnet Configuration	44
Displaying Management Port on a Dual MSM System.....	45
Displaying Management Port with “Show Port Config”.....	45
AutoNegotiation and 1000BaseT Ports	45
Switching and VLANs.....	45
Default Routes or Static Routes.....	45
Modifying the Protocol “IP”	45
Configuring a Protocol Filter with ‘ffff’.....	45
GVRP/GARP.....	45
GVRP/GARP and Load-Sharing.....	45
Deleting Protocols from a VLAN.....	45
Maximum Number of VLANs Supported.....	46
VLAN to VLAN Access Profiles	46
Load Sharing	46
Spanning Tree	47
Mirroring	47
QoS.....	48
Bandwidth Settings and their impact.....	48
QoS Profile minimum Bandwidth should not exceed 90% totals.....	49
Access Lists on BlackDiamond I/O modules	50
Access Lists Using the IP Deny Any Rule	50
VLAN QoS Between I/O BlackDiamond Modules.....	50
MAC QoS.....	50
Access Lists and IP Fragmentation	50
Bi-Directional Rate Shaping	50
1000BaseT Ports as Loopback Ports	50
ESRP	50
ESRP Port Count.....	50
Multiple ESRP VLANs.....	50
ESRP Interoperability.....	50
Mixing Clients and Routers on an ESRP-Enabled VLAN	51
Ensure that EDP is Enabled.....	51
ESRP and Host Attached Ports.....	51
ESRP host attached ports are not supported in subVLANs or domain-member VLANs (10728). In addition, it is recommended that the use of ESRP host attached ports be limited on normal ESRP	

VLANs, superVLANs, and domain-master VLANs to limit the amount of processing placed on the system CPU to manage these ports.ESRP and 3000 VLAN Configurations.....	51
ESRP and Bi-Directional Rate Shaping.....	51
ESRP and the “Save” Command	51
IP Unicast Routing.....	51
VLAN Aggregation	51
Multinetting	52
RIP Routing.....	52
RIP V2 Authentication.....	52
Disabling RIP	52
Routing with OSPF.....	52
Set the RouterID	52
OSPF Default Cost	52
IP Multicast Routing and Snooping	53
Cisco Interoperation.....	53
IGMP Settings	53
IGMP & IGMP Snooping with IP Unicast and Multicast Routing.....	53
IPX Routing.....	53
Tuning.....	53
IPX and Round-Robin Loadsharing.....	53
IPX Performance Testing Using Traffic Generators.....	53
IPX and Bi-Directional Rate Shaping.....	53
Security and Access Policies.....	54
RADIUS	54
Server Load Balancing	54
Default Ping Health Checking	54
Server Load Balancing with 3DNS.....	54
Web Cache Redirection / Policy Based Routing.....	54
Health Checking	54
VLAN boundary	55
WCR and SLB Not Supported on the Same Switch	55
Precedence of flow redirection rules.....	55
WEB Management - VISTA	56
WEB Server Busy.....	56
Closing Internet Explorer 4.0	56


Configuring VLANs Through VISTA on GVRP Ports.....	56
Default QoS Profile does not Appear	56
Log Entry Order	56
Vista and RADIUS.....	56
Vista and Management Port	57
SNMP.....	57
Trap Receivers as Broadcast Entry.....	57
Control of UDP Port used in Sending Traps	57
Bridge MIB Attributes.....	57
SNMP Timeout Setting	57
SNMP Access Profile.....	57
SNMP and Auto-negotiation Settings	57
SNMP and the BGP MIB	58
DLCS	58
Virtual Chassis.....	58
Issues Resolved.....	58
Issues Resolved from v6.1.8b7	58
General	58
BlackDiamond.....	59
Alpine	60
Port Mirroring.....	60
Spanning Tree	60
Load Sharing	60
General IP.....	60
ESRP	60
IP Multicast	60
BGP	61
OSPF	61
RIP.....	61
DVMPRP.....	61
PIM.....	61
Bi-Directional Rate Shaping.....	61
QoS.....	62
Access Control Lists.....	62
SLB.....	62
WCR.....	62

SNMP	62
Issues Resolved from v6.1.7b7	63
General.....	63
BlackDiamond	64
Alpine	64
STP	64
Port Mirroring.....	64
BGP	64
ESRP.....	65
OSPF.....	65
RIP.....	65
PIM.....	65
General IP	65
IP TOS	66
QoS.....	66
Access Control Lists.....	66
SLB.....	66
SSH.....	66
TACACS	66
RADIUS	67
SNMP	67
Issues Resolved from v6.1.7b5	67
General.....	67
IP Multicast.....	67
Issues Resolved from v6.1.6b19	68
General.....	68
Alpine	68
Bi-Directional Rate Shaping.....	68
Access Control Lists.....	68
ESRP.....	68
IPX.....	68
Spanning Tree.....	69
OSPF.....	69
BGP	69
SNMP	69
Vista.....	69

Issues Resolved from v6.1.5b20	69
General	69
BlackDiamond.....	70
ALPINE.....	71
VLANs	71
Port Mirroring.....	71
General IP.....	71
General Routing	72
IP Multicast	72
OSPF	72
BGP	72
RIP.....	73
DVMPRP.....	73
PIM.....	73
ESRP	73
Access Control Lists.....	73
Server Load Balancing	74
Web Cache Redirection.....	74
SNMP	74
DLCS.....	74
TACACS+	74
VISTA	74
Issues Resolved from v6.1.4b20	74
General	75
BlackDiamond.....	75
Alpine.....	76
VLANs	76
ESRP	76
OSPF	76
RIP.....	76
PIM.....	77
General Routing	77
General IP.....	77
IP Multicast	77
IPX	78
QOS.....	78
Bi-directional Rate Shaping.....	78

Server Load Balancing.....	78
Policy-Based Routing and Web Cache Redirection.....	78
SNMP	79
Access Control Lists.....	79
Port Mirroring.....	79
DLCS	79
Radius.....	79
Vista.....	80
Issues Resolved from v6.1.4b12	80
General.....	80
BlackDiamond	80
Alpine	80
VLAN Aggregation	80
SLB.....	81
Port Mirroring.....	81
Issues Resolved from v6.1.3b11	81
General.....	81
BlackDiamond	81
VLANs.....	82
Bi-directional Rate Shaping.....	82
ESRP.....	82
OSPF.....	82
RIP	83
BGP	83
IP	83
IP Multicast.....	83
IGMP	83
QOS	83
SNMP	84
Issues Resolved from v6.1.2b7	84
General.....	84
BlackDiamond	84
IP	84
Access Control Lists.....	84
VLAN Aggregation	85
RIP.....	85

PIM-SM.....	85
SLB.....	85
SNMP.....	85



Release Notes for ExtremeWare v6.1.8

These release notes contain information on features and issues specific to this release of ExtremeWare v6.1 not covered in the *ExtremeWare Software User Guide v6.1*. Numbers that appear in parenthesis are used for internal reference and can be ignored.

Overview

This document contains the following sections:

- Supported hardware
- Instructions for upgrading from v6.0
- “i” Chipset unique features
- New features in ExtremeWare v6.1
- Supported limits
- Clarifications, known behaviors, and problems
- Issues resolved from previous releases

For information on resolved issues going back farther than those documented here, 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>.

Supported Hardware

This release of ExtremeWare v6.1 is designed to support products using the "i" chipset *only*.

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):

Extreme Switch Platform	ExtremeWare Filename/Version	BootRom Filename/Version
BlackDiamond switch using MSM64i MSMs*	v618b12.xtr or v618b12ssh.xtr/v6.1.8b12	Ngboot65.bin*/ v6.5
Alpine 3808	v618b12.xtr or v618b12ssh.xtr/v6.1.8b12	Ngboot65.bin*/ v6.5
Alpine 3804	v618b12.xtr or v618b12ssh.xtr/v6.1.8b12	Ngboot65.bin*/ v6.5

Supported Hardware

Summit 7i/7iT	v618b12.xtr or v618b12ssh.xtr/v6.1.8b12	Ngboot65.bin*/ v6.5
Summit 1i/1iT	v618b12.xtr or v618b12ssh.xtr/v6.1.8b12	Ngboot65.bin*/ v6.5
Summit 5i/5iT/5iLX	v618b12.xtr or v618b12ssh.xtr/v6.1.8b12	Ngboot65.bin*/ v6.5
Summit 48i	v618b12.xtr or v618b12ssh.xtr/v6.1.8b12	Ngboot65.bin*/ v6.5

* ExtremeWare v6.1.5 or above requires version 6.5 BootROM. Note that BootROM v6.5 is backward compatible with all versions of Extremeware v6.x. ExtremeWare v6.x requires that the BlackDiamond switch use only the MSM64i in MSM slots marked "A" and "B". It is not possible to use MSM32 modules with ExtremeWare v6.x or higher.

** Please see "Issues Resolved from v6.1.7b5" section for special upgrade instructions when upgrading from v6.1.7b5.

*** Please see "Issues Resolved from v6.1.8b7" section for special upgrade instructions when upgrading from v6.1.8b7.

BlackDiamond Module Support

BlackDiamond modules supported with ExtremeWare v6.1.5 and the MSM64i include:

BlackDiamond Module	ExtremeWare v6.1.5 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 6.1.8 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

*As documented in Chapter 1 of the *ExtremeWare Software User Guide v6.1* and within these release notes, some new features require that the "i" chipset also be present on the I/O module in order for the feature to function.

**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 v6.1.5 and above include:

Alpine Module	ExtremeWare v6.1.5 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 Module Support

Summit modules supported with ExtremeWare v6.1.5 or above include:

Summit Module	ExtremeWare v6.1.5 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 CLI command `show port config`. 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.

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 (v6.1.6 and above)

Upgrading from v6.0 to v6.1

If you are currently running a release of ExtremeWare v6.0 on a Summit or BlackDiamond, simply TFTP download the new image to the primary or secondary image space, then make sure you are configured to use that image space and reboot the switch. We recommend downloading into an image space that is not currently in use. In this way, the currently used image is preserved should you need to go back. For example, if the primary image space is used currently, to upgrade to v6.1 use the commands:

```
download image <ipaddress> <v6.1_filename> secondary
use image secondary
reboot
```



Note: You must upgrade to BootROM v6.5 to run ExtremeWare v6.1.5 or above. See below for instructions on bootrom upgrades.

Note: When upgrading from ExtremeWare 6.1.7b5 to ExtremeWare v6.1.7b7, it is required that you upgrade directly from v6.1.6b19 to v6.1.7b7 or re-download the desired configuration to the switch after upgrading to v6.1.7b7 from v6.1.7b5. Upgrades directly from v6.1.7b5 to v6.1.7b7 may result in the Radius, TACACS, and SNMP trap receiver configuration parameters being modified after reboot.

Note: When upgrading from v6.1.8b7 to v6.1.8b12, it is required that you upgrade directly from v6.1.6b19 to v6.1.7b7 or re-download the desired configuration to the switch after upgrading to v6.1.8b12 from v6.1.8b7.

Upgrading BootROM

This release is also supplied with a new BootROM image for the Summit and BlackDiamond switches. The new BootROM release is critical to the upgrade aspects of ExtremeWare v6.1 and is backward compatible with ExtremeWare v6.0 and previous ExtremeWare v6.1 releases. Be sure to perform the BootROM upgrade *before* upgrading to ExtremeWare v6.1 using the command:

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

Upgrading ExtremeWare

Below are instructions specific to upgrading to, and downgrading from, ExtremeWare v6.1 for Summit and BlackDiamond switches.

Upgrading Switches

ExtremeWare v6.1 can read a stored configuration saved by ExtremeWare v6.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. Commands are "save primary" and "save secondary".

2. Configure the switch to use the primary image and the primary configuration. Commands are “use image primary” and “use config primary”.
3. Verify that all of the above procedures were completed successfully with the command “show switch”.
4. Upload the configuration of the switch to a TFTP server for safekeeping. Command is “upload config <ipaddress> <filename>”.
5. If not already running bootrom v6.5, TFTP download bootrom v6.5 to the switch. An example command is “download bootrom <ipaddress> ngboot6x.bin”. Reboot the switch to come up with bootrom v6.5.
6. TFTP download desired version of ExtremeWare v6.1 to the primary image space. An example command is “download image <ipaddress> v61xby.xtr primary”.
7. Reboot the switch. The previous configuration of the switch will be preserved going from the previous version of ExtremeWare to the desired version of ExtremeWare v6.1. Verify that the switch is operating as expected. After verification, you may configure features specific to the current version of ExtremeWare. Save the configuration to the primary space and do NOT save to the secondary configuration space unless 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 v6.0 or previous v6.1 image, configure the switch to use the secondary configuration. Command is: “use config secondary”. 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 by using the command: “use image secondary”.
3. Verify that the above procedures were completed successfully with the command “show switch”.
4. 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 EW, you must ensure that the switch configuration matches that version of EW or below. Pointing the configuration to a new version of EW and using a previous version of EW is not supported.

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.
- ESRP options - Support for ESRP Groups, ESRP Domains and ESRP Host Attach.
- Traffic statistics on a per VLAN basis.
- Subnet directed broadcast forwarding improvements.
- System health-checker on the BlackDiamond.

Staying Current

For support purposes, we always recommend operating the most current release of ExtremeWare. 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/>

New Features in ExtremeWare 6.1

Following are descriptions of features added or enhanced since the publication of the *ExtremeWare User Guide v6.1*. An “*” denotes features added since 6.1.5.

General

*show tech-support Command

To help Extreme Networks Technical Support diagnose technical problems, please use the following command:

```
show tech-support
```

This command provides the output for the following commands:

- show version
- show switch
- show config

- show diag
- show slot
- show fdb
- show iparp
- show ipfdb
- show ipstats
- show iproute
- show ipmc cache detail
- show igmp snooping detail
- show mem detail
- show log
- additional internal debug commands

This command disables the CLI paging feature. When using this command with large configurations, you must wait for the system to display the full configuration before regaining access to the CLI.

Image Filename

The switch software filename extension has been changed from “.Z” to “.xtr” to avoid representation of Extreme switch software as WinZip files.

*show vlan Command

The show vlan command now displays the total number of VLANs configured on the switch.

Mirroring

The mirroring port now has the option to transmit tagged or untagged frames. This allows the mirroring of multiple ports and/or VLANs to a mirror port while preserving the ability of a single protocol analyzer to track and differentiate traffic within a broadcast domain (VLAN) and across broadcast domains (e.g. across VLANs when routing). The new syntax for the command is:

```
enable mirroring to <port> [tagged | untagged]
```

*Enhancements to Subnet Directed Broadcast Forwarding

You can enable or disable the hardware forwarding of subnet directed broadcast IP packets. This allows wire-speed forwarding rates for subnet directed broadcast packets.

To enable or disable hardware forwarding, use the following command:

```
<enable|disable> ipforwarding fast-direct-broadcast [vlan <vlan_name>]
```

The entries will be installed in the IP forwarding table as standard entries, and you can view them using the show ipfdb command.

In addition, you can configure the VLAN router interface to either forward and process all subnet-directed broadcast packets, or to simply forward these packets after they have been installed in the IP forwarding database. This option allows you to improve CPU forwarding performance by having

New Features in ExtremeWare 6.1

upper layers, such as UDP and TCP, ignore broadcast packet processing (for example, if the packets have IP-options configured).

To enable or disable this functionality, please use the following command:

```
<enable | disable> ipforwarding ignore-broadcast vlan <vlan_name>
```

To view the configuration settings, use the following command:

```
show ipconfig [detail | vlan <vlan_name>]
```

The default setting for both commands is “disabled.” Using these commands, you can achieve a 30-50% reduction in system processing cycles in forwarding subnet directed broadcast traffic on a BlackDiamond switch, and a 100% reduction on the Alpine and Summit switches.

Note that although forwarding performance is improved in the BlackDiamond switch, the CPU will continue to observe the subnet directed broadcast packets and will not be able to ignore such packets when traversing modules in a BlackDiamond. Only “i” series modules are supported for this command on the BlackDiamond switch.

*Packet Forwarding Options for IP interfaces

You can now configure the forwarding options for packets destined to IP interfaces that are down on a switch. This is a global setting for all interfaces on the system. There are 3 options for configuration:

- Consume – packets will be sent to the appropriate upper layer protocols and forwarded accordingly. An example of this would be for ICMP requests and responses.
- Drop – packet will be dropped by the interface and not forwarded to the destination(s).
- Forward – default option. Packets such as DHCP responses would be forwarded to the appropriate destination.

The default configuration is to forward packets destined to down IP interfaces.

To enable or disable this functionality, please use the following command:

```
config ip-down-vlan-action <consume | drop | forward>
```

To view the configuration settings, use the following command:

```
show ipconfig
```

BlackDiamond

System Health Checking

The system health checker tests both the backplane and the CPU by periodically forwarding packets and checking for the validity of these packets. All error messages are logged in the syslog and the diagnostics CLI show output. A “CRIT” message will be posted to the log if any of the packet tests fail. If you observe a failure, please contact Extreme Technical Support. Please see note in the “Known Issues” section for the system health checker.

To enable the system health checker, use this command:

```
enable sys-health-check
```

To disable the system health checker, use this command:

```
disable sys-health check
```

To configure the system health checker, use this command:

```
configure sys-health-check alarm-level [card-down | default | log |
system-down | traps]
```

This command allows you to configure the switch's reaction to a failed health check, and provides the following options:

- `log`—post a CRIT message to the log
- `traps`—post a CRIT message to the log and send a trap
- `card-down`—post a CRIT message to the log, send a trap, and bring the module down (this option is not yet available)
- `system-down`—post a CRIT message to the log, send a trap, and bring the system down

The default option is `log`.

To view the status of the system health checker, use this command:

```
show diag
```

Note that you cannot use both mirroring and the system health checker at the same time. If you configure mirroring with the system health checker enabled, the health checker will indicate that it has been disabled by sending a message to the syslog.

Runtime Diagnostics

The runtime diagnostics perform a single test on a single I/O blade. Runtime diagnostics are not supported on management modules. All error messages are logged. To perform diagnostics on an I/O blade, use this command:

```
run diagnostics [normal | extended] slot <slot number>
```

Use the `normal` option when you want a fast (30 – 60 seconds) hardware status check. Use the `extended` option when you want a more thorough test. The `extended` option requires significantly more time to complete, depending on the number of ports on the blade.



Note: Only run extended diagnostics when the switch can be brought off line. The tests conducted during extended diagnostics are extensive and can affect traffic that must be processed by the system CPU (9838). This includes the switch's ability to maintain ESRP Master/Slave state and the ability to process other control protocols (10405).

To view results of the diagnostics test, use this command:

```
show diag
```

Disabling G1 Support

On BlackDiamond switches, you can now disable support for G1 (first generation, or non “i” series) modules to optimize control data performance for the system CPU. When you disable support for non-“i” series modules, they will neither be powered up nor pass traffic in a BlackDiamond system. You must save and reboot for these changes to take effect. The default setting is enabled.

To enable non-“i” series I/O module support, use this command:

New Features in ExtremeWare 6.1

```
enable g1-module support
```

To disable non-“i” series I/O module support, use this command:

```
disable g1-module support
```

*Multicast Performance Enhancements

The BlackDiamond switch can optimize “i”-series multicast data forwarding performance. To increase the performance of multicast applications, you can disable non-“i” series I/O modules in the system. In addition, you can modify the backplane loadsharing policy for more robust support of multicast streams. Note that the round-robin algorithm is not supported on non-“i” series I/O modules. The default backplane loadsharing policy is “port-based”.

To configure the switch backplane load-sharing policy, use this command:

```
configure backplane-ls-policy <address-based | port-based | round-robin>
```

To display non-“i” series module configuration, use this command:

```
show switch
```

Layer 2 Switching and VLANs

Maximum Number of VLANs Increased to 3,000 on all “i” Series

All “i” series switches now support 3,000 VLANs as the maximum number of configured VLANs.

*STPD BPDU Tunneling

ExtremeWare allows a BPDU to traverse a VLAN without being processed by the Spanning Tree process even if STP is enabled on that port. To use BPDU tunneling on a VLAN, use the following command:

```
enable ignore-bpdu vlan <vlan_name>
```

If you have a known topology and have switches outside of your network within your STPD, use this feature to keep the root bridge within your network.

Renaming a VLAN

You can rename an existing VLAN using a new CLI command. Note that the “MacVlanDiscover” or the “Default” VLAN names should not be changed.

To modify a VLAN name, use this command:

```
configure vlan <vlan_name> name <new vlan name>
```

VLAN Statistics

You can collect statistics on a per VLAN basis using a new CLI command. Statistics available include Receive and Transmit Unicast, Receive and Transmit Multicast, Receive and Transmit Broadcast, and Receive and Transmit Byte Count. This is available on “i” series products only

To display VLAN statistics, use this command:

```
show vlan stats vlan <vlan_name> <vlan_name>
```


Note that multiple VLAN names can be used in this syntax for multiple VLAN displays.

Jumbo Frames

Path MTU Discovery

ExtremeWare now supports path MTU discovery. In path MTU discovery, a source host will assume that the path MTU is the MTU of the first hop, which is known. The host will send all datagrams on that path with the DF bit set, restricting fragmentation. If any of the datagrams must be fragmented by an Extreme switch along the path, that switch will discard the datagrams and return ICMP Destination Unreachable messages with a code meaning "fragmentation needed and DF set". Upon receipt of such a message (sometimes called a "Datagram Too Big" message), the source host reduces its assumed path MTU for and can retransmit.

The path MTU discovery process ends when one of the following is true:

- The host sets the path MTU low enough that its datagrams can be delivered without fragmentation
- The host does not set the DF bit in the datagram headers

A host can choose not to set the DF bit because it is willing to have datagrams fragmented. Normally, the host continues to set DF in all datagrams, so that if the route changes and the new PMTU is lower, the host can perform PMTU discovery again.

IP Fragmentation with Jumbo frames

ExtremeWare now supports IP Fragmentation. If an IP packet originates in a local network that allows large packets and that packet traverses a network that limits packets to a smaller size, the packet will be fragmented instead of discarded. This is designed for use in conjunction with Jumbo frame support. Frames that are fragmented are not processed at wire-speed within the switch fabric. Also note that Jumbo frame to Jumbo frame fragmentation is not supported – only Jumbo frame to normal frame fragmentation is currently supported (9148).

To configure VLANs for IP fragmentation, you must do the following:

1. Enable jumbo frames on the incoming port
2. Add the port to a VLAN
3. Assign an IP address to the VLAN
4. Enable ipforwarding on the VLAN
5. Set the MTU size using the following new command:

```
config ip-mtu <size> vlan <vlan name>
```

The ip-mtu value can be 1500 or 9216, with 1500 the default. If you enter a value other than 1500, the switch will recognize that value as 9216.



Note: To set the MTU size greater than 1500, all ports in the VLAN must be jumbo-frame enabled.

IP Fragmentation within a VLAN

ExtremeWare also supports IP Fragmentation within a VLAN. This feature does not require you to configure the MTU size. To use IP fragmentation within a VLAN, you must do the following:

New Features in ExtremeWare 6.1

1. Enable jumbo frames on the incoming port
2. Add the port to a VLAN
3. Assign an IP address to the VLAN
4. Enable ipforwarding on the VLAN



Note: If you leave the MTU size as the default value, when you enable jumbo-frame support on a port in the VLAN you will receive a warning that the ip-mtu size for the VLAN is not set at max jumbo frame size. You can ignore this warning if you want IP fragmentation only within a VLAN. For inter-VLAN IP fragmentation, all ports in the VLAN must be configured for Jumbo frame support. For intra-VLAN IP fragmentation, all ports in the VLAN should not be configured for jumbo frame support.


vMANs - VPN Services for Metropolitan Area Providers

vMAN services allow the "tunneling" of any number of 802.1Q and/or Cisco ISL™ VLANs into a single VLAN which can be switched through an Extreme ethernet infrastructure. A given vMAN tunnel is completely isolated from other tunnels or VLANs. This feature is useful in building transparent private networks that need point-to-point or point-to-multipoint connectivity across an ethernet infrastructure. The VLAN tagging methods used within the vMAN tunnel are transparent to the tunnel. For the MAN provider, the tagging numbers and methods used by the customer are transparent to the provider.

The steps to configure a vMAN tunnel are:

- 1) modify the 802.1Q Ethertype the switch uses to recognize tagged frames
- 2) configure the switch to accept larger MTU size frames ("Jumbo" frames).
- 3) create tunnels by creating VLANs and configuring member ports as tagged on switch-to-switch ports and untagged on the tunnel's ingress/egress ports.

Depicted below is an example configuration with vMANs. Two tunnels are depicted that have ingress/egress ports on each Summit7i.



For all the Summit7i's that are shown, the configuration is:

```

config dot1q ethertype 9100
enable jumbo-frame ports 31,32
config jumbo-frame size 1530
create vlan Tunnel1
config vlan Tunnel1 tag 50
config vlan Tunnel1 add port 1-4 untag
config vlan Tunnel1 add port 31,32 tagged
create vlan Tunnel2
config vlan Tunnel2 tag 60
config vlan Tunnel2 add port 5-8 untag
create vlan Tunnel2 add port 31,32 tagged

```

On the BlackDiamond shown, the configuration is:

```

config dot1q ethertype 9100
enable jumbo-frame ports all
config jumbo-frame size 1530create vlan tunnel1
config vlan tunnel1 tag 50
config vlan tunnel1 add port 1:1-1:2 tagged
create vlan tunnel2
config vlan tunnel2 tag 60
config vlan tunnel2 add port 1:1-1:2 tagged

```

Specific to this configuration, a Layer 1 or Layer 2 redundancy method would also employed, such as Spanning Tree or other methods ExtremeWare offers.

Spanning Tree Rapid Root Failover

ExtremeWare now supports rapid root failover for faster Spanning Tree failover recovery times. The default setting is disabled.

To enable or disable rapid root failover, use the following commands:

```
<enable | disable> stpd <spanning tree name> rapid-root-failover
```

To display configuration, use the following command:

```
show stpd <spanning tree name>
```

General IP Functionality

IPARP Address Checking

The IP ARP Address checking feature allows you to configure the option to allow IP ARP entries to be installed in the IP ARP table for addresses that are not within the correct range of IP Address/Mask of the configured VLAN IP interface. Enabling the feature rejects entries that do not

New Features in ExtremeWare 6.1

fall within the IP address range of the VLAN IP interface address. Disabling the feature allows any entry to be installed.

To enable or disable iparp checking, use the following commands:

```
<enable | disable> iparp checking
```

QoS

Bi-directional Rate Shaping for Routed VLANs

ExtremeWare now supports bi-directional rate shaping for VLANs with routed or switched ingress traffic. The following text replaces the text in the *ExtremeWare Software User Guide v6.1*:

Bi-directional rate shaping allows you to manage bandwidth on switch or routed traffic flowing both to and from the switch. You can utilize up to 8 ingress rate shaping queues per VLAN and 8 egress rate shaping queues per physical port. By defining a QoS Profile's minimum and maximum bandwidth corresponding to the physical queue and port, you define committed information rates for each queue and port. Different bandwidth rates can be applied to ingress vs. egress traffic.

You can then provide any supported traffic groupings (e.g. physical port, VLAN, .IP, DiffServ, IP address, Layer 4 flow etc.) for the 8 pre-defined QoS Profiles thereby directing specific types of traffic to the desired queue. The traffic groupings used are not dependent on whether the traffic is switched or routed.

Configuring Bi-Directional Rate Shaping

Each VLAN requires a loopback port; all traffic from rate-shaped ports is directed through the loopback port for that VLAN. To rate-shape ingress traffic, configure QoS normally on the loopback port for the VLAN. The maximum bandwidth and traffic grouping defined in the QoS profile for the loopback port defines the rate limit for ingress traffic on rate-shaped ports in that VLAN.

Use the following guidelines for bi-directional ingress rate shaping:

- You must configure a loopback port before adding rate-shaped ports to the VLAN.
- A loopback port cannot be used by an external device.
- The loopback port must be configured with a unique loopback VLAN tag ID.
- Ingress traffic on a port that is configured to use the loopback port will be rate-shaped.
- Ingress traffic on a port that is not configured to use the loopback port will not be rate-shaped.
- Unicast traffic from a non-rate-shaped port to a rate-shaped port within the VLAN will not be rate-shaped.
- The aggregate forwarding bandwidth of all rate-shaped ports in a VLAN is determined by the traffic groupings and bandwidth settings for the QoS Profiles of the loopback port.
- For 10/100 ports, you can configure the loopback port as a 10 Mbps port to achieve lower bandwidth values.

Use the following guidelines for bi-directional egress rate shaping:


- For egress rate shaping, simply set the maximum bandwidth of the QoS profile on the egress port.

Bi-Directional Rate Shaping Limitations


Consider the following limitations when configuring bi-directional rate shaping:

- When configuring VLAN memberships, delete all rate-shaped ports before deleting the loopback port.
- If rate-shaped ports within a VLAN use different bandwidth parameters, set the priority of the QoS profiles on the loopback port and rate-shaped ports to low.
- Layer 2 switched rate-shaping only affects a single VLAN.
- IP forwarding must be enabled on the VLAN prior to adding the loopback port to a VLAN for L3 rate shaping. If you do not enable IP forwarding first, you must reboot the switch for the rate shaping configuration to take effect.
- On a BlackDiamond switch, the loopback port must be on the same I/O module as the rate-shaped ports (this does not apply to an Alpine switch).
- You cannot use tagged ports for rate shaping.
- You cannot use rate shaping on load-shared ports..

In addition, when rate shaping routed traffic on 10/100 ports, rate shaping ports cannot belong to the same block of 8 ports as loopback or normal ports. The following figures show the rate shaping blocks on the Summit48i switch, Alpine FM-32Ti module, and BlackDiamond F48Ti and F96Ti modules.




Summit48i port groupings

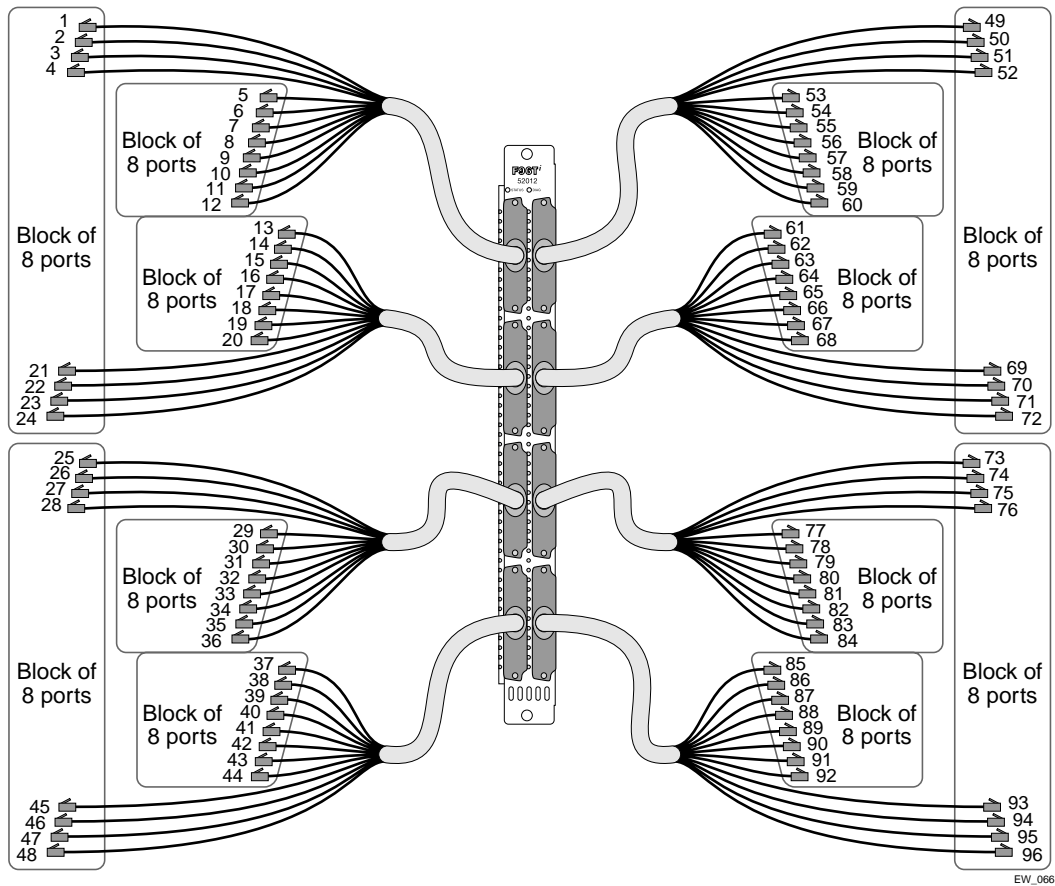


Alpine FM-32Ti module port groupings

New Features in ExtremeWare 6.1



BlackDiamond F48Ti module port groupings



BlackDiamond F96Ti module port groupings

If you have IP routing enabled and you add a rate-shaped port to a VLAN, and the rate-shaped port is in the same port block as loopback or normal ports, ExtremeWare will return one of the following error messages:

```
ERROR:Rate shaped port can't be in the same block as loopback port
```


```
ERROR: Normal port 8:20 cannot share the block with rate shaped port
```



Note: If you configure ports for Layer 2 switching only, and later enable IP routing, ExtremeWare will not return port block conflict errors.

If traffic between rate-shaped ports is routed, as in the following example, you must assign the same QoS profile to both ports (or both VLANs).

New Features in ExtremeWare 6.1



Routing bi-directional rate shaped ports

If you assign the rate-shaped ports to different QoS profiles, the switch will only rate-shape traffic through the loopback port.

Bi-Directional Rate Shaping Commands

To add the loopback port to the VLAN, use the following command:

```
config vlan <vlan name> add port <port> loopback-vid <vlan_tag>
```

To enable the loopback port, use the following command:

```
Restart port <loopback_port>
```

To add rate-shaped ports to the VLAN, use the following command:

```
config vlan <vlan name> add port <portlist> {tagged | untagged}
{nobroadcast} soft-rate-limit
```

To delete rate-shaped ports from the VLAN, use the following command:

```
config vlan <vlan name> delete port <portlist>
```

To configure the rate-shaping parameters of the loopback port, use the normal QoS profile configuration command, as follows:

```
config qosprofile <qosprofile> {minbw <pcnt>} {maxbw <pcnt>} priority
<level> {buffer <pcnt>} {<portlist>} <loopback port number>
```

To remove the rate-shaping parameters of the loopback port, use the normal QoS profile configuration command without the buffer or portlist parameters:

```
config qosprofile <qosprofile> {minbw <pcnt>} {maxbw <pcnt>} priority
<level> <loopback port number>
```

To display the bi-directional rate shaping configuration, use the following command:

```
show vlan {<vlan name> | detail}
```

This command designates rate-shaped ports with an "R" and loopback ports with an "L" next to the port number.

To set the port speed of a loopback port, use the normal port configuration command, as follows:

```
config ports <portlist> auto off {speed [10 | 100 | 1000]} duplex [half
| full]
```

Maximum QoS Buffer

QoS profiles now have an additional buffer parameter: `maxbuf`. The `maxbuf` parameter allows you to set a maximum buffer for each queue, so that a single queue will not consume all of the unallocated buffer space. The `maxbuf` values can be set in kilobit or megabit increments. The minimum value is 0K and the maximum is 16,384K. The default value is 256K. Unless you have explicit reason to modify these parameters, do not modify them. Only unique situations should require any non-default configurations of QoS. You must reboot the switch for a change to the `maxbuf` parameter to take effect (10846).

To set the `maxbuf` value on a queue, use the following command:

```
configure qosprofile <qos profile> minbw <percent> maxbw <percent>
priority <priority> maxbuf <number>
```

To view the `maxbuf` configuration, use the following command:

```
show qosprofile
```

ESRP

*ESRP Multiple Ping Tracking

You can configure ESRP to track connectivity to up to four outside responders using a simple ping. To configure ping tracking, use the following command:

```
config vlan <vlan name> add track-ping <ip address> frequency <seconds>
miss <number>
```

To view the status of the tracked devices, use the following command:

```
show esrp
```

*ESRP Port Restart

You can configure ESRP to restart ports if those ports are members of a VLAN that becomes a slave. To configure port restart for a port, use the following command:

```
config vlan <vlan name> add ports [<portlist> | all] restart
```

To disable port restart for a port, use the following command:

```
config vlan <vlan name> add ports [<portlist> | all] no-restart
```

If a VLAN becomes a slave, ESRP disconnects member ports that have port restart enabled. This causes downstream devices to remove the port from their FDB table, allowing you to use ESRP in networks that include equipment from other vendors. After 3 seconds the ports re-establish connection with the ESRP switch. Note that the “norestart” option is not available in v6.1.6b19. To remove a port from the “restart” configuration, delete the port from the VLAN and re-add it.

*dont-count Parameter

You can use the `dont-count` parameter in the following command to remove host ports from consideration in the active port count:

```
config esrp port-mode host ports <portlist> [dont-count]
```

This parameter is useful if you have host ports on a switch that alternate between active and inactive, forcing frequent ESRP failover.

ESRP Environment and Diagnostic Tracking

ESRP is now capable of tracking hardware status. If a power supply or fan fails, or if the chassis is overheating, the priority for the ESRP VLAN will change to the failover settings. ESRP will also track the diagnostics that run on each blade. If the diagnostics fail, the ESRP VLAN will change the priority to the failover settings.

To configure the failover priority for ESRP VLANs, you must first assign a priority to each ESRP VLAN, using the following command:

```
config vlan <vlan name> esrp priority
```

The range of the priority value is 0 to 254; a higher number has higher priority. The default priority setting is 0.



Note: If you set the priority to 255, the ESRP VLAN will remain in standby mode even if the master ESRP VLAN fails.

You will typically configure both ESRP VLANs with the same priority.

Next, you must give the priority flag precedence over the active ports count, which has precedence by default, by using the following command:

```
config vlan <vlan name> esrp esrp-election priority-ports-track-mac
```

Because the priority of both VLANs are set to the same value, ESRP will use the active ports count to determine the master ESRP VLAN.

Finally, you must set the failover priority. To configure environmental tracking, use the following command:

```
config vlan <vlan name> add track-environment failover <priority>
```

To disable environmental tracking, use the following command:

```
config vlan <vlan name> delete track-environment failover <priority>
```

To configure diagnostic tracking, use the following command:

```
config vlan <vlan name> add track-diagnostic failover <priority>
```

To disable diagnostic tracking, use the following command:

```
config vlan <vlan name> delete track-diagnostic failover <priority>
```

Typically, you will set the failover priority lower than the configured priority. Thus, if one of the VLANs experiences a hardware or diagnostics failure, that VLAN becomes the standby VLAN.



Note: If you set the failover priority to 255, the ESRP VLAN experiencing hardware or diagnostics failure will become the standby VLAN and will remain in standby mode even if the master ESRP VLAN fails.

If a switch has both environmental and diagnostic failures, the higher of the configured priorities will be selected. If no hardware tracking failures are encountered, ESRP will use the user configured priority.

Increased Number of ESRP Domain Member VLANs

The maximum number of ESRP domain member VLANs is now 2000. To configure 2000 ESRP VLANs, you must first disable support for non "i" chipset products using the following command:

```
disable g1-module-support
```

You must then set the CPU transmit priority to normal, using the following command:

```
configure cpu-transmit-priority normal
```

After these two changes you can configure up to 2000 ESRP VLANs.

IP Unicast Routing

Route Map Support

This release includes the ability to apply route maps to routes that are being added to the kernel route table. You can configure the route maps based on the following origins of the route:

- Direct

New Features in ExtremeWare 6.1

- Static
- RIP
- OSPF
- BGP

These route maps match the various characteristics of the route based on the originating protocol, and set the characteristics of the route. The characteristics that can be matched and set depends on the protocol originating the route. Use the following command to configure route maps:

```
config iproute route-map [bgp | direct | e-bgp | i-bgp | ospf | ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static] none <route-map>
```

Use the following command to view the log:

```
show iproute {priority | vlan <vlan> | permanent | <ipaddress> <netmask> | route-map | origin [direct | static | blackhole | rip | bootp | icmp | ospf-intra | ospf-inter | ospf-as-external | ospf-extern1 | ospf-extern2]} {sorted}
```

You can make dynamic changes to the route map. Direct and Static route changes are reflected immediately, while RIP, OSPF, and BGP changes are reflected within 30 seconds.

VLAN Aggregation SubVLAN Address Range Checking

SubVLAN address ranges can be configured on each subVLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

Configuring a subVLAN range:

```
configure vlan <vlan_name> subvlan-address-range <ip_address> - <ip_address>
```

Removing a subVLAN address range:

```
configure vlan <vlan_name> subvlan-address-range 0.0.0.0 - 0.0.0.0
```

Viewing subVLAN range:

```
show vlan [vlan_name]
```

Note that there is no error checking to prevent the configuration of overlapping subVLAN address ranges between multiple subVLANs. Doing so can result in unexpected behavior of ARP within the superVLAN and associated subVLANs.

IP FDB Performance

The IP FDB handling has been enhanced so that only relevant IP FDB entries are flushed when entries are modified in the system routing table. As a result of this enhancement, you will see a performance improvement in situations where there are frequent route changes. Performance is improved because route changes will not affect the traffic that is not relevant to the route change.

OSPF

*OSPF Point-to-Point Support

ExtremeWare now allows you to manually configure the OSPF link type for a VLAN. Previously, ExtremeWare automatically determined the link type. The following table describes the link types.

Link Type	Number of Routers	Description
Auto	Varies	ExtremeWare automatically determines the OSPF link type based on the interface type. This is the default.
Broadcast	Any	Routers must elect a designated router (DR) and a backup designated router (BDR) during synchronization. Ethernet is an example of a broadcast link.
Point-to-point	Up to 2	Will not operate with more than 2 routers on the same VLAN. Synchronizes faster than a broadcast link because routers do not elect a DR or BDR. PPP is an example of a point-to-point link.

An Ethernet link is typically an OSPF broadcast link, which supports from 0 to n OSPF routers. An OSPF broadcast link must first synchronize the routers, then elect a Designated Router and a Backup Designated Router. An OSPF point-to-point link supports only 0 to 2 OSPF routers, and does not elect a Designated Router or a Backup Designated Router. An OSPF point-to-point link will therefore synchronize faster than a broadcast link. However, if you have three or more routers on the VLAN, OSPF will fail to synchronize if the neighbor is not configured.



Note: The number of routers in an OSPF point-to-point link is per-VLAN, not per-link.



Note: All routers on the VLAN must have the same OSPF link type.

To specify the new link-type parameter, use the following command:

```
config ospf add vlan <vlan name> area <aid> [link-type
auto|broadcast|point-to-point] [passive]
```

To prevent the inadvertent disruption of OSPF due to the addition of other routers on an OSPF point-to-point link, you can explicitly configure the point-to-point neighbor using the following command:

```
config ospf vlan <vlan name> neighbor [add|delete] <ipaddress>
```

*Configurable OSPF Wait Interval

ExtremeWare allows you to configure the OSPF wait interval, as opposed to the fixed Router Dead Interval period. You can now configure the following parameters:

- Retransmit interval (RxmtInterval)
- Transit delay (TransitDelay)
- Hello interval (HelloInterval)
- Dead router wait interval (RouterDeadInterval)
- Router wait interval (WaitInterval)



Caution: Do not configure OSPF timers unless you are comfortable exceeding OSPF specifications.

Retransmit interval—This is the length of time that the router will wait before retransmitting an LSA that is not acknowledged. If you set an interval that is too short, you will create unnecessary retransmissions. The default value is 5 seconds.

Transit delay—This is the length of time it takes to transmit an LSA packet over the interface. The transit delay must be greater than 0.

Hello interval—This is interval at which routers will send hello packets. Smaller times allow routers to discover each other more quickly, but also increase network traffic. The default value is 10 seconds.

Dead router wait interval—This is the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. This interval should be a multiple of the hello interval. The default is 40 seconds.

Router wait interval—This is the interval between the interface coming up and the election of the designated router and backup designated router. This interval should be greater than the hello interval. If it is close to the hello interval, the network will synchronize very quickly, but might not elect the correct designated router or backup designated router. The default is the dead router wait interval.



Note: The OSPF standard specifies that wait times are equal to the dead router wait interval.

To specify the timer intervals, use the following command:

```
config ospf vlan <vlan name> timer <retransmitint> <transitdly>
<helloint> <deadint> [<waitint>]
```

*OSPF CLI Display Enhancement

ExtremeWare provides several new filtering criteria for the `show ospf lsdb` command. You can specify multiple search criteria, and only results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

To display the current link-state database, use the following command:

```
show ospf lsdb [detail | summary | stats] [area all | <aid>[/<len>]]
[lstype all | as-external | external-type7 | network | router |
summary-asb | summary-net] [lsid <id>[/<len>]] [routerid <id>[/<len>]]
```

The `detail` option displays all contents of the link-state database. The `summary` option displays one line per LSA. The `stats` option displays the number of matching LSAs.

In addition, you can use a shortened form of the command, as follows:

```
show ospf lsdb
```

This shortened form displays all areas and all types in a summary format.

OSPF Database Overflow

The OSPF Database Overflow feature allows you to both limit the size of the link-state database and to maintain a consistent link-state database across all the routers in the system. Maintaining a consistent link-state database across all the routers in the domain ensures that all routers have a consistent view of the network.

Consistency is achieved by the following:

- Limiting the number of External LSAs in the database of each router
- Ensuring that all routers have identical LSAs

Use the following command to configure OSPF Database Overflow:

```
configure ospf ase-limit <number> {timeout <seconds>}
```

This command takes two parameters.

1. A limit specifying the number of External LSAs (excluding the default LSAs) that the system will support before it goes into overflow state. A limit value of zero disables the functionality.
2. The timeout in seconds after which the system will come out of overflow state. A timeout value of zero leaves the system in overflow state until OSPF is disabled and enabled.

When the link-state database size limit is reached, OSPF Database Overflow flushes external LSAs originated locally from the link-state database. OSPF Database Overflow flushes the same LSAs from all the routers, thereby maintaining consistency.

OSPF Password Encryption

The neighbor password for OSPF is now encrypted in upload/download configuration.

OSPF Passive Interface

A new CLI command allows you to configure an OSPF interface as passive. Hello packets are not sent over passive interfaces and adjacencies are not established over them.

To configure an OSPF interface as a passive interface:

```
configure ospf add vlan <vlan name> area <area identifier> passive
```

To reconfigure an OSPF interface as a normal interface:

```
configure ospf add vlan <vlan name> area <area identifier>
```

To display passive interface configuration:

```
show ospf interface [detail]
```

Route Map Support for OSPF Export

The `enable ospf` command has been enhanced to support route maps. The route map will be applied on each and every route that is exported to OSPF. It can be used for filtering or for setting the cost, cost type, and tag of the exported route. You can use this feature to make dynamic changes to the route map.

Use the following commands to enable OSPF route map export:

```
enable ospf export direct [cost <metric> [ase-type-1 | ase-type-2] {tag <number>} | <route map>]
```

```
enable ospf export static [cost <metric> [ase-type-1 | ase-type-2] {tag <number>} | <route map>]
```

```
enable ospf export rip [cost <metric> [ase-type-1 | ase-type-2] {tag <number>} | <route map>]
```

New Features in ExtremeWare 6.1

```
enable ospf export [bgp | i-bgp | e-bgp] [cost <metric> [ase-type-1 |
ase-type-2] {tag <number>} | <route map>]
enable ospf export vip [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route map>]
```

The current format of the command is also supported for backward compatibility.

BGP

BGP Peer Groups

You can use BGP Peer Groups to group together up to 128 BGP neighbors. This simplifies configuring and updating neighbors because all neighbors automatically inherit the parameters of the BGP Peer Group. All neighbors in the Peer Group share the following mandatory parameters:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

You assign a unique name to the Peer Group when you create it. Use the following command to create and delete a peer group.

```
[create | delete] bgp peer-group <peer-group>
```

Use the following commands to configure the parameters of the peer group.

```
config bgp peer-group <peer-group> remote-as <number>
config bgp peer-group <peer-group> [route-reflector-client | no-route-
reflector-client]
config bgp peer-group <peer-group> weight <number>
config bgp peer-group <peer-group> source-interface [any | vlan <vlan>]
config bgp peer-group <peer-group> timer keep-alive <number> hold-time
<number>
config bgp peer-group <peer-group> nlri-filter [in | out] [none |
<access profile>]
config bgp peer-group <peer-group> as-path-filter [in | out] [none |
<access profile>]
config bgp peer-group <peer-group> route-map-filter [in | out] [none |
<route map>]
config bgp peer-group <peer-group> [send-communities | dont-send-
communities]
config bgp peer-group <peer-group> soft-reset {input | output}
config bgp peer-group <peer-group> password <password>
```



```

config bgp peer-group <peer-group> [next-hop-self | no-next-hop-self]
[enable | disable] bgp peer-group <peer-group> soft-in-reset
[enable | disable] bgp peer-group <peer-group>

```

When you modify the parameters, the changes will be applied to all neighbors in the peer group. Modifying the following parameters will automatically disable and enable the neighbors before the changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

To create a new neighbor and include them as a member of the peer group, use the following command:

```
create bgp neighbor <ip address> peer-group <peer-group> {multi-hop}
```

This command creates the new neighbor as part of the peer group, and the neighbor inherits all existing parameters from the peer group. This command requires the Peer Group to have remote AS configured.

To add an existing neighbor to a peer group, use the following command:

```
config bgp neighbor [<ip address>| all] peer-group <peer-group>
{acquire-all}
```

If you do not specify `acquire-all`, then only the mandatory parameters are inherited from the peer group. If you specify `acquire-all`, then all the parameters of the Peer Group are inherited. This command will disable the neighbor before adding the neighbor to the peer group.

You can display existing peer groups using the command:

```
show bgp peer-group {detail | <peer-group> {detail}}
```

If you specify `detail` the parameters of the neighbors in the Peer Group that are different from the Peer Group are displayed.

BGP Route Selection

BGP will select routes based on the following precedence (from highest to lowest):

- Weight
- Local Preference
- Shortest length (shortest AS path)
- Lowest origin code
- Lowest MED
- Route from external peer
- Lowest cost to Next Hop
- Lowest RouterId

New Features in ExtremeWare 6.1

BGP MD5 Authentication

A new CLI command allows users to configure MD5 authentication between BGP neighbors. The maximum length of the password string is 31 characters.

To configure BGP MD5 authentication:

```
configure bgp neighbor <ip address> password <password>
```

To un-configure BGP MD5 authentication:

```
configure bgp neighbor <ip address> password none
```

To show BGP MD5 authentication configuration:

```
show bgp neighbor detail
```

BGP Password Encryption

The neighbor password for BGP is now encrypted in upload/download configuration.

BGP Route Flap Dampening

BGP route flap dampening and associated commands are not yet supported in ExtremeWare.

IP Multicast Routing and Snooping

Static Rendezvous Points RPs

ExtremeWare now allows you to override the PIM bootstrap message that selects a dynamic RP so that you can define a static RP in your network. To define a static RP, use the following command:

```
config pim crp static <rp address>
```



Note: If you configure a static RP in your network, configure the static RP on all switches in that network.

PIM Mode Translation

An Extreme switch functioning as a PMBR (PIM Multicast Border Router) will now integrate PIM-SM and PIM-DM traffic separated by the PMBR.

When forwarding PIM-DM traffic into a PIM-SM network, the PMBR will notify the rendezvous point (RP) that the PIM-DM network exists. The PMBR will then forward PIM-DM multicast packets to the RP, which will then forward the packets to those routers that have joined the multicast group.

The switch will also forward PIM-SM traffic to a PIM-DM network. The PMBR will send a join message to the RP and the PMBR will then broadcast traffic from the RP into the PIM-DM network.

There are no new commands that need to be entered to enable PIM-SM to PIM-DM functionality. By having both the DM mode interface and SM mode interface on the same router, the PMBR functionality will be automatically enabled.

IP Multicast Cache Display

The `show ipmc cache` command now displays a legend with a summary of each entry in the table (8984).

IGMP Snooping

IGMP leave Message

IGMP Snooping now supports the IGMP leave message. Previously, when a port sent an IGMP leave message, the router would send a query to determine which ports wished to remain in the multicast group. If other members of the VLAN wished to remain in the multicast group, the router would ignore the leave message.

Now, when a port sends an IGMP leave message, the switch will remove the IGMP snooping entry after 10 seconds. The router will still send a query to determine which ports wish to remain in the multicast group. If other members of the VLAN wish to remain in the multicast group, the router would ignore the leave message, but the port will be removed from the IGMP snooping table.

If the last port within a VLAN sends an IGMP leave message, the router will not receive any responses to the query, and the router will immediately remove the VLAN from its multicast group.

IGMP Display

The `show igmp snooping` command can now be displayed with a summary or detail view (9062).

IPX

IPX Routing Design Restrictions Lifted

Several design restrictions for routing IPX traffic have been removed. For all "i" chipset products, ExtremeWare now supports separate routing interfaces for IP and IPX traffic on the same VLAN, load sharing of IPX routed traffic, and supports 802.1Q tagged packets on a routed IPX VLAN..

IP and IPX on the Same VLAN

ExtremeWare now supports IP and IPX routing within the same VLAN. This feature does not require any special configuration changes.

Tagged IPX VLAN

Previously, IPX routing could only be performed on untagged traffic (e.g traffic without 802.1Q encapsulation). ExtremeWare now supports tagged 802.1Q traffic on an IPX VLAN that is performing routing. Tagging is most commonly used to create VLANs that span multiple switches. Using VLAN tags, multiple VLANs can span multiple switches using one or more trunks. In a port-based VLAN, each VLAN requires its own pair of trunk ports. Another benefit of tagged VLANs is the ability to have a port be a member of multiple VLANs. This is particularly useful if you have a device (such as a server) that must belong to multiple VLANs. A single port can be a member of only one port-based VLAN. All additional VLAN memberships for that port must be configured with tags.

To configure a tagged IPX VLAN, assign a tag to the VLAN in the usual way using the following command:

```
config vlan <name> tag <vlanid>
```

The valid range is from 1 to 4095.

To assign tagged ports to the VLAN in the usual way, use the following command:

```
config vlan <name> add port <portlist> {tagged | untagged}
{nobroadcast}
```

To display your VLAN settings, use the following command:

New Features in ExtremeWare 6.1

```
show vlan {<name>} {detail}
```

These commands are unchanged, but they now support tagged IPX VLANs.

IPX load sharing

Previously a VLAN involved with IPX routing could not also be a member of a load-shared group. ExtremeWare now supports IPX load sharing on all products that use the “i” chipset. There is no additional configuration requirement to support this function, simply configure load sharing as you would normally. Please refer to Chapters 4 or 5 in the *ExtremeWare Software User Guide v6.1*

Security and Access Policies

ICMP ACL Precedence

You can now assign precedence values to access lists for ICMP traffic. The precedence number is optional; access list entries that contain a precedence number are evaluated from highest to lowest. Precedence numbers range from 1 to 25,600, with the *number 1 having the highest precedence*. Assigning precedence allows the switch to resolve conflicts between ICMP rules.

Access List Display

The `show access-list` command now displays a single line summary of each access list, instead of the multiple lines previously displayed.

IPX Routing Access Policies

Routing Access Policies now support IPX RIP, SAP and IPX node rules. Routing Access Policies consist of access rules, and are used to perform packet filtering and forwarding decisions on incoming traffic. Each IPX RIP or SAP packet arriving on an ingress port is compared to each access profile rule in sequential order, and is either forwarded or dropped. To create IPX access lists, use the following command in its new form:

To create an IPX access profile, use the following commands:

```
create access-profile <access_profile> type [ipaddress | ipx-node |  
ipx-net | ipx-sap | as-path | bgp-community]
```

To configure an IPX net, node or sap access profile, use the respective following command:

```
config access-profile <access_profile> [add | delete] {seq-number} ipx-  
net <ipx_net_id_in_hex> <ipx_net_id_mask_in_hex>
```

```
config access-profile <access_profile> [add | delete] {seq-number} ipx-  
node <ipx_net_id_in_hex> <ipx_net_id_mask_in_hex>  
<ipx_node_id_in_mac_address_format>
```

```
config access-profile <access_profile> [add | delete] {seq-number} ipx-  
sap <ipx_sap_type_in_hex> <ipx_name_string>
```

To assign IPX access lists as either import or export filters to RIP or SAP, use the following commands:

```
config ipxrip vlan [<vlan name> | all] import-filter [<access_profile>  
| none]
```

```
config ipxrip vlan [<vlan name> | all] export-filter [<access_profile>  
| none]
```

```

config ipxsap vlan [<vlan name> | all] import-filter [<access_profile>
| none]
config ipxsap vlan [<vlan name> | all] export-filter [<access_profile>
| none]

```

To view your access list configuration, use the following command:

```
show access-profile <access_profile>
```

BGP and OSPF Route Map Support for Tagging and DSB Accounting

This release adds enhancements to the route map support for BGP and OSPF tagging. Enhancements have been made to Match on Tags and Set on Tags, Accounting Indices, Cost, and Cost Type.

Match Tag---This can be used in the applied route map when redistributing OSPF routes from the Kernel routing table to BGP.

Set Tag---This can be used in the applied route map to set the tag value of the exported route when redistributing routes into OSPF from the Kernel routing table. It can also be used in the applied route map to set the tag value in the route that is being added when adding BGP routes to the routing table.

Set Accounting---This can be used in the route map when BGP, OSPF, RIP, static, or direct routes are added to the Kernel route table.

Set Cost---This can be used in the applied route map to set the cost of the exported route when redistributing routes into OSPF from the Kernel routing table. This can also be used in the applied route map to set the cost of the route when BGP and static routes are added to the Kernel routing table.

Use the following command to enable tagging:

```

config route-map <route-map> <sequence number> [add | delete] match
[nlri-list <access-profile> | as-path [access-profile <access-profile>
| <as no>] | community [access-profile <access-profile> | <as no> :
<number> | number <community> | no-advertise | no-export | no-export-
subconfed] | next-hop <ip address> | med <number> | origin [igp | egp |
incomplete] | tag <number>]

```

Use the following command to enable accounting:

```

config route-map <route-map> <sequence number> [add | delete] set [as-
path <as no> | community [[access-profile <access-profile> | <as no> :
<number> | number <community> | no-advertise | no-export | no-export-
subconfed] | remove | [add | delete] [access-profile <access-profile>
| <as no> : <number> | number <community> | no-advertise | no-export |
no-export-subconfed] | ] | next-hop <ip address> | med <number> |
local-preference <number> | weight <number> | origin [igp | egp |
incomplete] | tag <number> | accounting index <number> value <number> |
cost <number> | cost-type [ase-type-1 | ase-type-2 ]]

```

Server Load Balancing

Health check definitions

For reference, the following health checks are available on all Server Load Balancing, Web Cache Redirection and Policy-based Routing functions. SLB functions will test individual servers. Web

Cache Redirection and Policy-based routing functions will test the next hops in accordance with the flow-redirection rules.

Layer 3 Ping Check—The default health checking is a simple ping check, where the switch sends an ICMP ping packet to the configured server or next hop. If 3 replies are lost, the server or next hop is set to “down” and flows are not redirected to it. The ping check is the only health checking that will work with a wildcard as the IP-Port.

Layer 4 Port Check—The switch will attempt to establish a TCP connection to the server or next hop.



Note: When using Web Cache Redirection or Policy Based Routing, the Layer 4 port must be defined in the flow and open on the next hop in order for the health check to succeed.

Layer 7 HTTP Check—The HTTP health check will download a specific page from the server or next hop configured for the flow. The switch will then search the page for a specific text string in the first 500 bytes. If the text string is found, the check passes. As an alternative you can configure the check to accept any data from the downloaded page.

Layer 7 FTP Check—The FTP health check establishes an FTP connection between the switch and the server or next hop. The switch will attempt to login using the name and password supplied during the configuration. The check will succeed when the switch successfully logs into the next hop.

Layer 7 NNTP Check—The NNTP health check connects to the server or next hop, establishes a connection, and attaches to a user defined newsgroup.

Layer 7 POP3, SMTP, and Telnet Check—These health checks attach to the server or next hop using the specified protocol and log in. After successful login the next hop is marked as “up”.

GoGo Mode Health Checking

ExtremeWare now supports health checking on servers participating in SLB GoGo Mode. You can configure multiple health checks (ping-check, tcp-port-checks and service-checks) simultaneously on a given GoGo mode grouping. A physical port in a GoGo mode grouping will be considered available for GoGo traffic only if all configured health checks pass.

Use the following commands to enable GoGo mode health checking:

```
enable slb gogo-mode master ping-check {ipaddress}
enable/ slb gogo-mode master tcp-port-check [port | all]
enable slb gogo-mode master service-check [http | ftp | telnet | smtp |
nntp | pop3 | all | tcpport]
```

Use the following commands to disable GoGo mode health checking:

```
disable slb gogo-mode master ping-check
disable slb gogo-mode master tcp-port-check [port | all]
disable slb gogo-mode master service-check [http | ftp | telnet | smtp
| nntp | pop3 | all | tcpport]
unconfig slb gogo-mode master health-check
```

This command disables and deletes all ping-check, tcp-port-check, and service-check configurations for this GoGo mode grouping. The GoGo mode grouping itself is not affected.

```
unconfig slb gogo-mode master service-check [http | ftp | telnet | smtp
| nntp | pop3 | all | tcpport]
```

This command disables and deletes the service check configuration. If the associated TCP port has not been used for any tcp-port-check configuration, the TCP port will be deleted as well.

Use the following commands to configure GoGo mode health checking:

```

config slb gogo-mode master ping-check frequency seconds timeout
seconds

config slb gogo-mode master health-check ipaddress

config slb gogo-mode master tcp-port-check [add | delete] port

config slb gogo-mode master tcp-port-check timer port frequency seconds
timeout seconds

config slb gogo-mode master service-check http {l4-port port} {url url
match-string [match_string | any-content]}

config slb gogo-mode master service-check ftp {l4-port port} {userid
userid | password {encrypted} password}

config slb gogo-mode master service-check telnet {l4-port port} {userid
userid| password {encrypted} password}

config slb gogo-mode master service-check smtp {l4-port port}
{dns_domain}

config slb gogo-mode master service-check nntp {l4-port port}
{newsgroup}

config slb gogo-mode master service-check pop3 {l4-port port} userid
userid password {encrypted} {password}

config slb gogo-mode master service-check timer [http | ftp | telnet |
smtp | nntp | pop3 | tcpport] frequency seconds timeout seconds

```

Use the following command to view your GoGo mode health checking configuration:

```
show slb gogo-mode {master} {configuration}
```

SLB Global Connection Timeout

A new command has been added to SLB transparent and translational modes to allow for the configuration of the global connection timeout period. This helps to avoid cases where connections would be closed when the TCP "FIN" and "ACK" timeout was too short (9487/9613).

To configure the global connection timeout period where seconds can be from 1 to 180, use the following command:

```
config slb global connection-timeout <seconds>
```

Note that the default value is 1 second. In addition, the timeout should be set as low as possible to avoid stale connections staying in the table.

Combined SLB and ESRP Failover

SLB and ESRP can be combined to provide a very high availability topology. Two commands are added to assist mapping an ESRP configured VLAN to the SLB failover unit number and to display the current SLB/ESRP configuration.

```

config slb esrp vlan <vlan name> [ add | delete ] unit [ 1 - 16 ]

show slb esrp

```

New Features in ExtremeWare 6.1

SLB Pool and VIP Statistics

Additional commands are added to display the statistic of SLB pool members and SLB VIPs.

```
show slb stats pool
show slb stats pool <poolname>
show slb stats vip
show slb stats vip <vipname>
```

SLB Pool Member Configuration

Three commands are added to configure the ratio and priority of an existing pool member and to display the current SLB pool statistics.

```
config slb pool <poolname> member <ipaddress : port> [ ratio <ratio> |
priority <priority> ]
```

SLB Proxy Client Persistence

Three commands are added to configure client persistence. The command is needed when a remote service provider uses multiple NAT address ranges to translate their internal client IP addresses and when client persistence is needed.

```
enable slb proxy-client-persistent
disable slb proxy-client-persistent
config slb proxy-client-persistent [ add | delete ] <ipaddress / mask>
```

Web Cache Redirection/Policy Based Routing

Health Checks

Several additional health checks are now supported for the flows that are defined under Web Cache Redirection and Policy Based Routing. The operation and definition of these health checks is identical to those used for Server Load Balancing. For a complete definition of these health checks, please see the 'Health Check Definitions' in this document under the Server Load Balancing section.



Note: Health checking works on the ports configured by the flow that they are associated with. For example, if a flow is configured to redirect on port 80 (traditionally HTTP) but FTP is configured as the service check, the switch will try to open an FTP session on port 80. The health check will fail if the protocol will not work on the configured flow.

Ping Check: The ping check is the only health checking that will work with a wildcard as the Layer 4 IP-Port. To configure a ping check for a defined flow, use the following command:

```
config <flow> service-check ping
```

Layer 4 Port Check: The port has to be defined and open on the next hop in order for the health check to succeed. To configure a Layer 4 health check for a defined flow, use the following command:

```
config <flow> service-check L4-port
```

HTTP Check: To configure an HTTP health check for a defined flow, use the following command:

```
config <flow> service-check http url "/test.htm" match-string "pass"
```


In this example the switch will connect to the cache and download the page **test.htm** in the root WWW directory and search the page for the word “**pass**” in the first 1000 bytes. Note that the quotation marks are necessary for the switch to recognize the web page and the string.

FTP Check: To configure an FTP health check for a defined flow, use the following command:

```
config <flow> service-check ftp user <user> <password>
```

NNTP Check: To configure an NNTP health check for a defined flow, use the command:

```
config <flow> service-check nntp <newsgroup>
```

POP3, SMTP and Telnet Checks: To configure a POP3, SMTP or Telnet health check for a defined flow, use the following command:

```
config <flow> service-check <pop3|smtp|telnet> user <user> <password>
```

Configuring health check timeouts and frequencies borrows from the Server Load Balancing command:

```
conf slb global service-check frequency <seconds> timeout <seconds>
```

Support for 'any' Layer 4 Flows

Policy-based routing and Web Cache Redirection now support an 'any' option for the Layer 4 protocol type which allows the redirection of TCP, UDP and other traffic types with the exception of ICMP traffic. To configure this capability, use the 'any' option in the new syntax for flow re-direction:

```
create flow-redirection <flow_rule_name> [tcp | udp | any] destination
[<ip_address>/<mask> | any] ip-port [<L4_port> | any] source
[<ip_address>/<mask> | any]
```

Policy-Based Routing with Route Load-Sharing

Policy based routing is used to alter the normally calculated next hop route which is based on the route table. This same alteration can also load-share across multiple routers. It implies a set of rules or policies that take precedence over information in the route table. These policies can perform a "flow-redirection" to different next-hop addresses based on the following criteria:

- IP source address and mask
- IP destination address and mask
- Layer 4 destination port

In the event that the next-hop address (or addresses) becomes unavailable, the switch will route the traffic normally. Several rules can be defined, the precedence of rules is determined by "best match" of the rule to the packet. If no rule is satisfied, no redirection occurs.

There are two types of commands to setup policy-based routing. One to configure the redirection rule(s) and one to configure the next hop IP address(es) to use:

```
create flow-redirection <flow_rule_name> [tcp | udp | any] destination
[<ip_address>/<mask> | any] ip-port [<L4_port> | any] source
[<ip_address>/<mask> | any]

config flow-redirection <flow_rule_name> [add | delete] next-hop
<ip_address>
```

If multiple next-hop addresses are defined, traffic satisfying the rule will be load-shared across the next hop addresses based on destination IP address. If next hop address(es) fail (do not respond to

New Features in ExtremeWare 6.1

ICMP pings), the switch will resume normal routing. Using policy-based routing has no impact on switch performance.

To show configuration and status of flow redirection rules, use the following command:

```
show flow-redirection [<flow_rule_name | <cr>]
```

SNMP

MIB Support

ExtremeWare now supports the following MIBs:

- *Ping (according to guidelines in RFC 2925)—Allows you to ping Extreme switches from your network management system (NMS). Also allows you to monitor pings originating from both your NMS and the console.
- BGP (according to guidelines in RFC 1657)—Allows you to monitor (but not configure) BGP routed traffic from your NMS.
- OSPF (according to guidelines in RFC 1850)—Allows you to monitor (but not configure) OSPF routed traffic from your NMS.
- CPU activity (private Extreme MIB)—Allows you to monitor CPU activity from your NMS.
- *VLAN configuration (private Extreme MIB)—Allows you to configure several additional VLAN parameters from your NMS, as follows:
 - Add loopback and rate-limited ports to a VLAN
 - Associate active or passive OSPF areas with a VLAN
 - Enable IP forwarding on a VLAN
 - Ignore STP on a VLAN

In addition, you can add per port QoS using your NMS.

These MIBs must be compiled in your network management system.

*SNMP ifDescription Enhancements

The ifDescription now has more detailed port information. For example, on a modular switch, ifDescription will display the module on which a port resides.

*SNMP ifTable Enhancement

The ifTable now indicates that you have enabled loopback mode on a VLAN.

SNMP ifMib Enhancements

The ifMib now has the ability to display slot/port for physical ports and VLAN name for VLANs index.

SNMP ifType Enhancements

The ifType in the If-MIB table now displays "softwareLoopback" for VLANs configured on the switch for which loopback-mode is enabled (10077).

***SNMP Trap Receiver Changes**

With ExtremeWare 6.1.6 and above, you can specify the source IP address for the SNMP trap to be sent by the switch to the network management station (10788).

***RADIUS and TACACS Password Length**

You can have RADIUS and TACACS passwords of up to 63 characters.

OSPF Traps

ExtremeWare now supports all standard OSPF traps. These include the following:

- ospfIfStateChange
- ospfVirtIfStateChange
- ospfNbrStateChange
- ospfVirtNbrStateChange
- ospfIfConfigError
- ospfVirtIfConfigError
- ospfIfAuthFailure
- ospfVirtIfAuthFailure
- ospfIfRxBadPacket
- ospfVirtIfRxBadPacket
- ospfTxRetransmit
- ospfVirtIfTxRetransmit
- ospfOriginateLsa
- ospfMaxAgeLsa
- ospfLsdbOverflow
- ospfLsdbApproachingOverflow

***SNMP dot1dTpFdbTable Enhancements**

The dot1dTPFdb Table can now be enabled and disabled by the user.. The default setting is disabled.

To configure the dot1dTPFdb table, use the following commands:

```
<enable|disable> snmp dot1dTpFdbTable
```

To display configuration of the dot1dTpFdb table, use the following command:

```
show management
```

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 supercedes any values mentioned in the

Supported Limits

ExtremeWare Software User Guide.

METRIC	DESCRIPTION	LIMIT
Access Profiles	Used by SNMP, Telnet, SSH2, Vista Web interface, and Routing Access Policies	128
Access Profile entries	Used by SNMP, Telnet, SSH2, Vista Web interface, and Routing Access Policies	256
Access List rules	Maximum number of Access Lists in which all rules utilize all available options	worst case: 255
Telnet - number of sessions	Maximum number of simultaneous Telnet sessions	8
SSH2 - number of sessions	Maximum number of simultaneous SSH2 sessions	8
SNMP - Trap receivers	Maximum number of SNMP trap receiver stations supported.	16
Syslog servers	Maximum number of simultaneous syslog servers that are supported.	4
Jumbo Frame size	Maximum size supported for Jumbo frames, including the CRC.	9216
VLANs - Summit "i"-series and Alpine	Includes all VLANs plus sub VLANs, super VLANs, etc.	3000
VLANs - BlackDiamond switch	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
IP Router interfaces	Maximum number of VLANs performing IP routing - excludes SubVLANs	512
MAC-based VLANs - MAC addresses	Maximum number of MAC addresses that can be downloaded to the switch when using MAC-based VLANs	7000
Protocol-sensitive VLANs - active protocol filters	The number of simultaneously active protocol filters in the switch.	15 for "i" based switch products; 7 otherwise
GVRP - Max VLANs	The maximum number of VLANs that can be learned and supported via GVRP.	128
Spanning Tree - Max STPDs	Maximum number of Spanning Tree Domains	64
Spanning Tree - Maximum number of ports	Maximum number of ports that can participate in a single Spanning Tree Domain.	Same as the number of available physical ports on the switch
IP Static Routes	Maximum number of permanent IP routes.	1024
IP route sharing entries	Maximum number of IP routes used in route sharing calculations. This includes static routes and OSPF ECMP.	8
IP Static ARP entries	Maximum number of permanent IP static ARP entries supported.	512
Static IP ARP Proxy entries	Maximum number of permanent IP ARP proxy entries.	512
Static MAC FDB entries - Summit7i and Alpine 3808	Maximum number of permanent MAC entries configured into the FDB.	256
Static MAC FDB entries - BlackDiamond switch	Maximum number of permanent MAC entries configured into the FDB.	512
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
ESRP Route-track entries -	Maximum number of routes that can be tracked by ESRP.	256

Summit 7i and Alpine 3808		
ESRP Route-track entries – BlackDiamond switch	Maximum number of routes that can be tracked by ESRP.	1024
ESRP – number of instances	Maximum number of ESRP supported VLANs for a single switch.	64
ESRP – number of ESRP groups	Maximum number of ESRP groups within a broadcast domain	4
ESRP – number of VLANs in a single ESRP domain – Summit 7i and Alpine 3808	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 default; 2000 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 default; 2000 max
FDB – Maximum number of L2/L3 entries – MSM64i with “i” series I/O modules, Summit 7i, Alpine 3808/3804	Maximum number of MAC addresses for the MSM64i, ALPINE 3808, and Summit 7i.	256,000
FDB – Maximum number of L2/L3 entries – Summit 1i, Summit 5i, and Summit 48i	Maximum number of MAC addresses for the Summit 1i, Summit 5i, and Summit 48i	128,000
FDB – Maximum number of L2/L3 entries for non-“i” series BlackDiamond I/O modules.	Maximum number of MAC addresses for the G4X, G6X, F32T, and F32F.	32,000
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
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
OSPF areas	As an ABR, how many OSPF areas are supported within the same switch	8
OSPF routes - BlackDiamond, Summit 7i, and Alpine 3808/3804	Recommended maximum number of routes contained in an OSPF LSDB for a "real" network.	100,000
OSPF routes – Summit 1i, Summit 5i, and Summit 48i	Recommended maximum number of routes contained in an OSPF LSDB for a "real" network.	30,000
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
BGP routes – BlackDiamond, Summit 7i and Alpine 3808/3804	Maximum number of routes contained in the BGP route table	1,000,000
BGP routes – Summit 1i, Summit 5i, and Summit 48i	Maximum number of routes contained in the BGP route table	250,000
BGP peers	Maximum number of BGP peers on a single router	128
BGP Peer Groups	Maximum number of BGP peer groups on a single router	16

Clarifications, Known Behaviors, and Problems

Policy Based Routing	Maximum number of policy based routes that can be stored on a switch	64
WCR - Max number of 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)
WCR – Max number of entries	Maximum number of active entries for any WCR rules. For example, one /16 rule can take all of the available entries.	64,000
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
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

Clarifications, Known Behaviors, and Problems

This section describes items needing further clarification, behaviors that may not be intuitive, and known problems.

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 Chapter 20 of the ExtremeWare Software User Guide v6.1.

Setting Autonegotiation 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 autonegotiation off must still include specifying the duplex mode. For example the command:

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

will turn autonegotiation off if port 4 is a gigabit port.

Flow Control

Flow Control is supported on gigabit ports only and is enabled or disabled as a part of auto-negotiation. If auto-negotiation is configured off, then flow-control is disabled. Status can be checked with the 'show port config' command under the column for flow control (2815).

Config Sys-Recovery Level Command

The `config 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).

Xmodem Downloads

Though not performed under normal circumstances, there are two ways to perform an Xmodem download of an ExtremeWare image. The first method is through the BootRom menu. The second is through CLI after the switch has booted. Listed below are issues associated with Xmodem download.

Extreme Switch Platform	Xmodem download through BootRom	Xmodem download through CLI
All Summit switches	No issues	Not Operational – Do not use (3662)
BlackDiamond switch	Remove 2 nd MSM first (see below)	MSM in slot “A” must be master (see below)

Xmodem Download Through BootRom on the BlackDiamond Switch

Though not performed under normal circumstances, if it is necessary to Xmodem download an image to an MSM using the BootRom menu; remove the second MSM from the BlackDiamond switch if present prior to beginning the operation (4936).

Xmodem Download Through CLI on the BlackDiamond Switch

To perform an Xmodem download using CLI to a BlackDiamond switch with two MSMs, be sure the MSM in slot A is “master” as indicated by the master LED showing green on the MSM module. You can then perform the Xmodem download through the upper serial port on the MSM in slot A. (4710, 4848).

Show Memory Output

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

EDP Packet Length

The length field in the Ethernet MAC Header of an EDP frame indicates that the length is 316 bytes even though the length of an EDP frame is 338 bytes. Some hosts and network equipment report an error when receiving an EDP frame and may correct the packet (truncate it) before forwarding the

Clarifications, Known Behaviors, and Problems

packet. In the case where the MAC Header is modified, Extreme switches may report an EDP “PDU length>packet length” message to the log (7830).

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).

System Related – BlackDiamond Switch

Using 110v Power on a BlackDiamond Switch

The BlackDiamond switch requires 220-volt power for correct operation. If 110-volt power is supplied, not all the I/O modules of the BlackDiamond switch may power up. The MSM will perform power calculations and will power up the maximum number of I/O modules from left (slot1) to right (slot 8). A module can be skipped if that module is not within the power budget, but the subsequent module is. 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 blade for the BlackDiamond switch 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 <DEBUG:KERN> killPacket: HW pqmWaitRx failed
04/13/1999 17:18.46 <DEBUG: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 timeframe. 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
```

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

The LEDs for the redundant ports on port 49 and port 50 do not turn amber when forwarding packets (5387).

Load sharing and the Gigabit Redundant ports are not current supported (9458/10716).

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).

Command Line Interface (CLI)

Don't Use the Encrypted Option When Creating an Account

There is an option available in the CLI for encrypting a password when creating a user account. Do not use this option. It is for use only in conjunction when uploading and downloading an ASCII configuration file to the switch so passwords are not indicated in clear text within the configuration file (4229, 4719).

Clarifications, Known Behaviors, and Problems

“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 may 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 Telnetting

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 on a Dual MSM System

The show port commands will only display the statistics and configuration information of the "mgmt" port on the Master MSM (7129).

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).

AutoNegotiation 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.

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 'unconfig switch all' to restore normal operation (2644, 4935).

GVRP/GARP

GVRP is currently not operational in EW 6.1.x software.

GVRP/GARP and Load-Sharing

Ports that are joined to VLANs using GARP/GVRP may not be part of a load-shared port group. Do not configure the GVRP protocol on ports involved in load sharing (3444).

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).

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 v6.0 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 `link` will be distributed on only 1 of the other active load share links in the round robin group. The traffic is not distributed evenly between the remaining ports (6977).

Port Based Load Sharing on Summit7i

Port-based load sharing on the Summit7i requires ingress ports to be on the same side of the switch 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 Blade Load Sharing

The I/O module configured to contain the “master” port must be physically present in a cross-blade load sharing group when the system is rebooted (8589).

When the module containing the Master port of a cross-module load share group is removed, load sharing will not be operational for any configured algorithms. The module containing the master port of a load share group must always be present in the system or replaced when removed. If servicing a switch with this configuration, load sharing should be disabled before the module is removed (13594).

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).

Spanning Tree

STP not Supported with ESRP

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

Mirroring

It is currently recommended that port mirroring not be used on ExtremeWare 6.1. There are several problems with mirroring that include inaccurate packet formats and system stability when mirroring is enabled and configured on a switch. These issues will be addressed in an upcoming EW release.

Mirroring Combined with Load Sharing

The following limitations apply when doing mirroring that also involves load-sharing ports:

- Mirroring VLANs or mirroring a VLAN on a specific port is known to cause behavioral problems when used in combination with load sharing. If enabled, load sharing will only make use of the master port and will not fail-over correctly. Deleting the mirror entry will restore normal operation (3735).
- If the master port of a load-shared port group is down, mirroring will not provide the traffic for the load-shared port group (4486).

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 Timeout period (260 sec.) (3534).

Mirroring Bandwidth

Performing mirroring on gigabit ports running at line-rate will reduce the traffic throughput by approximately thirty percent (4151).

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).

Mirroring and Download Configuration

Uploading a configuration file with mirroring parameters will not download the mirroring configuration to the switch because the tagged/untagged argument is missing (10429).

Mirroring and Tagged Packets on the BlackDiamond

On non-“i” series modules, the untagged parameter in port mirroring will capture both tagged and untagged packets (10643).

Mirroring and Destination MAC Address

On some mirrored packets, users may see the destination MAC address transmitted as an all “0”’s address for addresses destined to the switch instead of the MAC address of the switch (10080).

Mirroring on the BlackDiamond

When disabling mirroring on the BlackDiamond on a port that has multiple VLAN configured, the following message is printed to the syslog (10061):

```
<WARN:KERN> pty0: Block de68 on slot 2 is already free Index=238
Owner=0
```

QoS

Bandwidth Settings and their impact

Bandwidth settings applied to QoS Profiles that are used for ingress or egress traffic are expressed as a percentage of bandwidth. QoS Profile bandwidth settings are in turn applied to queues on physical ports. The actual impact of the bandwidth setting is determined by the port speed (10, 100 or 1000 Mbps) and by the actual granularity capabilities of the switch.

Maximum bandwidth settings

The maximum bandwidth percentage settings determine the port bandwidth available to each queue. Use the following table to determine the actual maximum bandwidth associated with each setting. If the maximum percentage bandwidth configured does not match one of the settings listed below, it will be rounded up to the next setting.

Max BW setting (%)	Max Bandwidth@ 10Mbps	Max Bandwidth@ 100Mbps	Max Bandwidth@ 1000Mbps
2%	200 Kbps	2 Mbps	20 Mbps
3%	310 Kbps	3.1 Mbps	30 Mbps
5%	490 Kbps	4.9 Mbps	50 Mbps
7%	690 Kbps	6.9 Mbps	69 Mbps
8%	790 Kbps	7.9 Mbps	79 Mbps
10%	960 Kbps	9.6 Mbps	96 Mbps
11%	1.12 Mbps	11.2 Mbps	112 Mbps
15%	1.5 Mbps	15 Mbps	150 Mbps
20%	1.9 Mbps	19 Mbps	190 Mbps
25%	2.5 Mbps	25 Mbps	250 Mbps
30%	3.3 Mbps	33Mbps	330 Mbps
35%	3.5 Mbps	35 Mbps	350 Mbps
40%	4.2 Mbps	42 Mbps	420 Mbps
50%	5 Mbps	50 Mbps	500 Mbps
60%	5.7 Mbps	57 Mbps	570 Mbps

65%	6.5 Mbps	65 Mbps	650 Mbps
70%	7.3 Mbps	73 Mbps	730 Mbps
80%	7.9 Mbps	79 Mbps	790 Mbps
95%	9.5 Mbps	95 Mbps	950 Mbps
100%	10 Mbps	100 Mbps	1000 Mbps

Minimum bandwidth settings

The minimum bandwidth percentage settings determine the reserved port bandwidth available to each queue. Use the following table to determine the actual reserved bandwidth associated with each setting. If the reserved percentage bandwidth configured does not match one of the settings listed below, it will be rounded up to the next setting. If the actual bandwidth utilized is below the minimum bandwidth within a queue, it is available for usage by other queues on that physical port.

Min BW setting (%)	Min Bandwidth@ 10Mbps	Min Bandwidth@ 100Mbps	Min Bandwidth@ 1000Mbps
4%	420 Kbps	4.2 Mbps	42 Mbps
6%	570 Kbps	5.7 Mbps	57 Mbps
8%	750 Kbps	7.5 Mbps	75 Mbps
9%	930 Kbps	9.3 Mbps	93 Mbps
10%	1 Mbps	10 Mbps	100 Mbps
20%	1.87 Mbps	18.7 Mbps	187 Mbps
25%	2.63 Mbps	26.3 Mbps	263 Mbps
35%	3.4 Mbps	34 Mbps	340 Mbps
50%	4.9 Mbps	49 Mbps	490 Mbps
60%	6.3 Mbps	63 Mbps	630 Mbps
80%	7.9 Mbps	79 Mbps	790 Mbps
89%	9.4 Mbps	94 Mbps	940 Mbps

QoS Profile minimum Bandwidth should not exceed 90% totals

The sum of the minimum bandwidth values for the applied QoS profiles should be kept to less than 90% to avoid any incidental starving of traffic. If the minimum bandwidth settings exceed 90% it is possible under a sustained situation of over-subscription, that a lower priority queue could become “starved” and not transmit traffic (4735).

Clarifications, Known Behaviors, and Problems

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/UDP port information is not included after the first fragment (for subsequent fragments).

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 1000Mbps to 100Mbps as the bandwidth settings will not be accurate when configured in 100Mbps mode.

ESRP

ESRP Port Count

The *ExtremeWare Software User Guide 6.1* incorrectly states on page 10-3 that a load-sharing port group is considered a single port when determining the ESRP master. ESRP actually considers all ports in a load-sharing port group when determining the ESRP master.

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 v6.1. These include route tracking and the ability to modify the election algorithm.

Mixing Clients and Routers on an ESRP-Enabled VLAN

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

ESRP and Host Attached Ports

Any ESRP VLANs that share ESRP host attached ports must be in different ESRP Groups.

ESRP host attached ports are not supported in subVLANs or domain-member VLANs (10728). In addition, it is recommended that the use of ESRP host attached ports be limited on normal ESRP VLANs, superVLANs, and domain-master VLANs to limit the amount of processing placed on the system CPU to manage these ports. ESRP and 3000 VLAN Configurations

When ESRP is configured with thousands of VLANs and multiple active ports in these VLANs (i.e., 8,000+ active ports), a save config command can cause the ESRP slave router to transition from Slave to Master and immediately return to the correct state when the save is complete (10860).

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 the "Save" Command

When executing a "save configuration" command on an ESRP slave router, ESRP can transition between Slave to Master status as the save compression is being completed (11724/11741).

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).

Clarifications, Known Behaviors, and Problems

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 CLI Show VLAN Stats Command

The CLI “show vlan stats <vlan_name>” command is not supported on multinetted VLANs.

RIP Routing

RIP V2 Authentication

The authentication feature of RIPv2 is not supported.

Disabling RIP

When RIP is enabled and disabled on a router, the following warning message is printed to the log. This has no effect on router functionality (8395).

```
04/11/2000 18:12.51 <WARN:RIP > ripTask: select return error
S_iosLib_INVALID_FILE_DESCRIPTOR
```

Routing with OSPF

Set the RouterID

It is recommended that you manually set the routerID of the switches participating in OSPF instead of having the switch automatically choose its routerID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older link state database being used. The CLI command is: ‘config ospf routerid <address>’ in which the address is provided in dotted decimal notation. Each switch must have a unique routerID (3823).

OSPF Default Cost

The default OSPF cost for a Gigabit Ethernet port in the metric table has been changed to 4. The *ExtremeWare Software User Guide* incorrectly states that the cost is 1. To modify the OSPF metric table, use the following command (7140):

```
config ospf metric-table 10M <cost> | 100M <cost> 1G <cost>
```

To view the ospf metric-table, use the following command:

```
show 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 Settings

The maximum value that can be set for the IGMP Query Interval is 429,496,729 seconds. The values that can be set for Query Response Interval and the Last Member Query Interval are between 1 and 25 seconds (9069).

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. This can be checked 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."
```

Server Load Balancing

Default Ping Health Checking

For Transparent and Translational modes, the L3 PING health check 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

We have observed the following issues:

1. Under very high sustained loads a Web Cache Redirect health check may fail and a cache servers 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
```

2. To use more than one next-hop server or router, the IP port must be set to any.
3. The IPFDB table will timeout before the IPARP table on the ports connected to the cache servers. To work around this configure the switch to have a higher IPFDB timeout than the IPARP timeout.
4. An ICMP PING health check 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 be utilized by using the proxy ARP functionality in the switch with minimal configuration changes to clients or servers.

WCR and SLB Not Supported on the Same Switch

Web Cache Redirection and Server Load Balancing are not supported simultaneously on the same switch due to possible overlapping of IP and IP port criteria that can be configured for each feature.

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
- 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 a 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.

Example 1:

Destination IP Address	Destination IP Port	Source IP Address	Priority Selection
------------------------	---------------------	-------------------	--------------------

Clarifications, Known Behaviors, and Problems

192.0.0.0/8	80	ANY	1
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.

Example 2:

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

WEB Management - VISTA

WEB Server Busy

In the event that multiple network managers are accessing the same switch you may experience a "Web:server busy" error message. A work-around is to log out and log in again via your web browser (1558).

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).

Configuring VLANs Through VISTA on GVRP Ports

Vista will allow the configuration of VLANs to ports already learned through GVRP, but does not override the GVRP settings. The configuration is silently accepted but no permanent change takes place. In general, explicit VLAN configuration should not be attempted for ports running GVRP (3461).

Default QoS Profile does not Appear

In the configuration of QoS using Vista, if the user does not have a user-configured QoS Profile, the default profile in use (qp1) will not appear in the "QoS Profile" column of the port configuration screen. An empty cell will be displayed instead of qp1 (2843).

Log Entry Order

If the log of a switch contains the maximum number of entries (999), Vista may not display the log entries in the correct order. Look at the log entry timestamp to determine the correct order (4712).

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).

If Radius is enabled on a switch, Vista is supported in read-only mode only (8887). The statistics menu can be viewed but the configuration menu will not be accessible.

Vista and Management Port

Vista cannot configure or be used to view the configuration of the “mgmt” port on the Summit 7i or BlackDiamond switch (7148).

SNMP

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).

Control of UDP Port used in Sending Traps

The control of what UDP port number to be used when sending SNMP traps can be done through the appropriate attributes in RFC 2021. It cannot be currently controlled through CLI and is not stored as part of a configuration (4914).

Bridge MIB Attributes

Unsupported counters

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

dot1dBasePortTable Display

The dot1dBasePortTable only displays the first 4 slots in a BlackDiamond switch (6918). The ifTable correctly displays all slots.

SNMP Timeout Setting

SNMP management stations may need to set the SNMP timeout value to 10 seconds as some large configuration operations take longer to perform, especially on the BlackDiamond switch (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 timeout and subsequently fail with the default timeout setting. It is suggested that the default timeout value be increased from 5 seconds to 60 seconds to decrease the frequency of such timeouts 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).

Issues Resolved

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).

DLCS

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

Virtual Chassis

The Virtual Chassis is not supported in ExtremeWare v6.0 or higher.

Issues Resolved

Below is a historical listing of issues resolved by release. If you wish to see issues resolved going back farther, 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>.

Issues Resolved from v6.1.8b7

The following issues were found in ExtremeWare v6.1.8b7 and resolved in this release (ExtremeWare v6.1.8b12).

General

- When upgrading to ExtremeWare 6.1.8b7, it is required that you upgrade directly from v6.1.7b9 (or below) to v6.1.8b12 or re-download the desired configuration to the switch after upgrading to v6.1.8b12 from v6.1.8b7. Upgrades directly from v6.1.8b7 to v6.1.8b12 may result in configuration parameters being inconsistent with user-defined options after a reboot.
- The System Health Check CPU diag failure messages were being printed to the system log with modules starting at slot number 0 versus slot number 1 (12533). All messages now correspond to the physical slot location within the BlackDiamond Chassis.
- The System Health Check error reporting now prints 2 different messages for backplane health check failures – one for corrupted packets and another for missed packets (12559). Previous to this release, a single message was printed for either a corrupted or missed health check packet.
- Modifications were made to the reporting of packet corruption messages to indicate whether the port was an internal or external port failure in the system (12465). The following messages will be printed to the system log upon detection of an error:

```
<CRIT:SYST> ERROR: Checksum Error on external port Slot 1, port 4,  
Previous=0xA, Current=0xB
```

```
<CRIT:SYST> ERROR: Checksum Error on Internal port Slot 1, port 8,  
Previous=0xA, Current=0xB
```

- Packets received by the CPU are now checked for corruption and reported in the system log (12531).
- The System Health Check feature has been improved to avoid conditions where false failures could occur when the system CPU would become busy. This would result in a log message and

- the missed packet counter being incremented in the CLI `show diag` output (12844/12842/12570).
- An oversubscribed backplane link on the BlackDiamond could result in a false health check failure message in the log and the missed packet counter being incremented in the CLI `show diag` output (12843).
 - The temperature overheat threshold has been reduced to 55 degrees Celsius to log and send an SNMP trap for high temperature conditions sooner than the current 65 degree Celsius threshold (12723/13303).
 - When configuring the FDB aging timer to the maximum value of 3600 seconds, the FDB would age entries at a time increment of 600 seconds per 15-second interval, resulting in a flush of the FDB at 90 seconds. Any value below 3600 seconds would be aged correctly at 1 second for each 1 second of elapsed time (12631).
 - When upgrading from a previous version of EW to EW6.1.8b7, the CLI `show vlan` command would display 2 active ports for the "mgmt" vlan (12439).
 - Repeated link transitions on a GBIC based Gigabit port with auto-negotiation disabled could cause a switch to become unresponsive to console and telnet access and become non-operational (11834).
 - Repeated attempts to run the CLI `show tech` command in a SSH enabled session while simultaneously running a CLI `show tech` command from a telnet session would cause the SSH session to become inactive. This would result in the switch becoming unresponsive and could only be recovered with a reboot of the system (this was observed only with the Windows based F-Secure SSH client) (12132/13349).
 - A software exception could occur in the "tRootTask" after a reboot of power loss on a system. Protection was put into the software to protect against this event (12924).
 - The `enable dot1d replacement ports <all | port number>` CLI command would automatically enable DLCS (13410).
 - When using the clear debug trace command, the "*" would not appear at the CLI prompt indicating that the system configuration was modified (11095).

BlackDiamond

- When upgrading from a version of EW to 6.1.8b7, permanent FDB entries would be assigned to port 2:66 rather than the user configured port. This would cause packets from that entry to be discarded (12540).
- The following message would be printed to the system log in rare circumstances when software was immediately trying to access a removed I/O module. This message was cosmetic for the module being removed (12698):


```
WARNING: Slot 8 Port 3 reg read bad MAC = 4
```
- When using the `show ports [slot number:port number, slot number:port number] stats` command, displaying ports across I/O modules 1-4 and 5-8 in the `show port <slot:port, slot:port> stats` command would result in the BlackDiamond to become slow and unresponsive until the port list was displayed (13446).
- When large number of 10/100 ports (greater than 140 ports) connected to multiple F48Ti module would simultaneously become inactive on a BlackDiamond (peer connections would simultaneously go down), the switch would become unresponsive and the slave MSM in a dual MSM system would force the system to reboot (12712).

Issues Resolved

Alpine

- A software exception would occur or the system would become unresponsive when using a saved pre-6.1.8 based configuration with jumbo frames enabled, upgrading to 6.1.8, and enabling sharing on an installed I/O module (12451/12452).
- When using the IP deny all ACL rule and saving and rebooting the system, the ACLs created prior to the IP deny all rule would not be applied correctly to the system. Removing and re-adding the affected ACLs would re-apply them (12789).

Port Mirroring

- Uploading a configuration file and re-downloading the same configuration to a switch would result in the following command to be download to the switch “enable mirroring to port 1:64” (13553).

Spanning Tree

- Enabling STP before enabling load sharing would cause spanning tree BPDUs to not be transmitted across a load share group. This would cause spanning tree to fail and a subsequent loop in the network. This would also occur when load sharing parameters were changed after Spanning Tree was already running and converged on the network (12694).

Load Sharing

- When using OSPF and load sharing, the OSPF interface cost would be modified when the master port was disconnected and the slave port would become active (12725).

General IP

- When disabling iparp checking and uploading the configuration via TFTP, a “#” symbol would follow the iparp checking configuration line in the uploaded config (12447).
- When a host sent a DHCP requests identified as having IEEE 802 network (hardware type 6) to a DHCP server, the receiving switch configured as a DHCP/BOOTP relay would properly forward the DHCP request to the DHCP Server. However, the switch would not correctly forward the server’s DHCP Offer to the requesting host and would change the packet’s HW type to a hardware type 1 and recalculate the checksum (12756).
- A software exception could occur in the “tNetTask” (mostly observed in WCR configurations) due to improper memory buffer management. Protection was put into software to protect against this event (12926).
- When using permanent FDB entries, deleting the specific entry would result in the IPFDB not clearing the learned information for that entry. Recreating the permanent FDB entry and clearing the IPFDB would result in the use count to become a negative “-1” (13282).

ESRP

- When configuring multiple VLANs tagged on a HA port in ESRP, the ESRP MAC address would not be removed when disabling ESRP (12419).

IP Multicast

- When 2 multicast streams were generated to the same destination from different source addresses, the second stream would be handled by the CPU forcing slow path packet forwarding for that stream (12532).

- Multicast addresses were not being handled correctly in the FDB in large multicast environments causing connectivity and performance problems (12542).
- Cisco HSRP packets destined to 224.0.0.2 were not being correctly learned in the IGMP snooping table as an "All groups" entry (12927).
- Having a non "i"-series module in a Black Diamond with an active port would cause multicast packets to loopback to the interface that sent out the packet when IGMP snooping was disabled. Enable IGMP snooping and disabling it would force the correct operation. The problem only exists if a Black Diamond was reset with IGMP snooping disabled (12837).

BGP

- When using a route map in redistribution and disabling BGP, saving the config, and rebooting the system, the route map could be deleted via CLI although it was still bound to the redistributed protocol (12461).

OSPF

- When using virtual link simple passwords, if a newly configured key was shorter than the old key and the first characters of the old key was the same as the new key, the password would not be authenticated and the virtual link would not be established to the backbone. For example, if the old key was "abcdef" and the new key was created as "abc", authentication would fail (11894).

RIP

- When a router was configured with multiple VLANs with different lengths for the subnet mask (24 bits, 30 bits etc), in some cases, the VLANs and their subnet mask would become mismatched causing the wrong network number and broadcast address to be used for routing decisions. When this event occurred, the following message would also be seen in the system log (12955):

```
<DEBUG:RIP > sendto: S_errno_ENETUNREACH dst 10.4.63.47
<INFO:RIP > ripSupply: Error 51 sending msg
```

DVMRP

- A software exception could occur in DVMRP when internal software would remove a route with the cache entry still in the table (12012).

PIM

- In PIM-SM environments, when a device was attached to the configured PIM router and would join the IP multicast group to receive the stream, multicast packets could leak to other ports in the same VLAN (12640).
- In a PIM environment with over 600 multicast streams, if PIM is disabled on the RP, the switch would become unresponsive and would have to be restarted to recover (12585/12841).

Bi-Directional Rate Shaping

- On switches configured with rate shaping and with multiple MAC addresses associated with IP FDB entries, upon aging the MAC addresses per the user-specified MAC aging-timer, the system could become unresponsive as it was flushing the MAC entries in the database. Subsequently, a message would be printed in the log indicating that the fdbAgeTask had consumed a high amount of CPU utilization (12369).

Issues Resolved

QoS

- The following message would be printed to the system log after configuring per port QoS profiles to ports and then adding them to a VLAN (12583):

```
<WARN:SNMP> SNMP PORT QOS Trying to insert duplicate instance 2
```

Access Control Lists

- Under certain conditions, when an ACL was created and TFTP downloaded to a switch with an overlapping precedence to an existing rule, the rule would be accepted into the ACL table and non-operational even though it existed in the show acl command. The following message would also be printed to the system log (13525):

```
Duplicate precedence XX: Rule name permit_name
```

SLB

- Using SLB health checking on a large number of servers (100-600 L3/L4 Health Checks) would result in degraded switch performance and SLB instability (12905).
- SLB running in conjunction with ESRP had connectivity problems whenever traffic destined for the active VIP passed through the Slave ESRP switch. The problem would be observed by the VIP refusing connections while still responding to ping requests, and no IPFDB entry being installed on the ESRP Slave switch for the IP address of the SLB VIP (10277).
- The default gateway configured on an SLB enabled switch would override the wildcard VIP configured to send traffic to the SLB servers in the pool. The traffic from the configured wildcard VIP to the SLB nodes would then be transmitted to the default gateway. With a default gateway configured on the switch, the connection would be closed immediately. Without the default gateway configured, the packets would be forwarded to the SLB servers (12214).
- The configure slb global connection-timeout <value> command would not be uploaded and subsequently downloaded via TFTP configuration upload/download (12412).
- The SLB node ping-check frequency can now be set to the available values of 1 – 60 seconds. Previously, the minimum configurable value was 5 seconds (13175).
- With SLB enabled, ICMP packets containing an IP fragment with protocol type set to TCP were not forwarded between hosts (12280).

WCR

- A software exception could occur in the “tNetTask” in WCR configurations due to improper memory buffer management (12926).

SNMP

- The capability to enable/disable the ignore-bpdu using SNMP is now supported (11913). Note an updated MIB is required for this capability.
- VLANs created with a loopback port and their corresponding RIF entries are now identified in the ifTable as a type 24 (softwareLoopback) interface (12589).
- Modifying the ignoreSTP flag in the extremeVlanIfTable would result in the following message(s) to be printed to the system log depending on enabling or disabling the functionality (12298):

```
SWITCH# Enabled Ignore Spanning Tree on vlan MacVlanDiscover  
SWITCH# Disabled Ignore Spanning Tree on vlan MacVlanDiscover
```

- Setting the defaultTTL to a value greater than 255 would result in the value being set to "0" which limited communication to the interface with the local segment (11572).
- The extremeVlanEncapsIfType would return a value of 2 as opposed to the correct value of 1 (12467).
- A SNMP get next query on the ospfNbrIPAddressTable would cause a system to become inaccessible. On the BlackDiamond, the switch slave MSM would subsequently reboot the Master MSM after the Master would become inaccessible (12702).
- Creating an entry in the extremePortLoadshare2 table and setting that entry to active would cause a software exception (12478).

Issues Resolved from v6.1.7b7

The following issues were found in ExtremeWare v6.1.7b7 and resolved in ExtremeWare v6.1.8b7.

General

- Modifications made to the power supply status checks on the Summit 5i/5iT to prevent false power supply failure messages (10869).
- A new facility has been enabled on all inferno-based products to examine packets on the system and report corruption in the packets fields if observed by the switch (10953).
- On switches with inferno series 10/100 ports, packet bursts greater than 128 packets between FDX and HDX ports would result in dropped packets. This would only occur when specific applications would burst packets in burst sizes greater than could be handled where the buffer would be overrun (11190).
- The LX70 GBIC ID would be read as "LX" in all CLI show port commands in EW 6.1.7b7 (11203).
- A link filter mechanism has been added to EW to protect against repetitive link state transitions that may occur as a result of end stations or other connections that are unable to maintain link (11526).
- When the system log was configured to display INFO messages, it would also display DEBUG level messages (11069).
- Typing show port 0 utilization would cause the console to become locked and require a reboot to regain console access (11868).
- As a user-level administrator, the show tech-support command would display the switch configuration. This has been changed to not allow the configuration to be printed to the system log for security purposes (11188).
- A user-level administrator would be allowed to save the configuration when using the quit command to exit from the login prompt (10134). The user would be correctly rejected from attempting a save command when issuing the command directly at the login prompt.
- VLANs configured with special characters using the " " parameters would not correctly set all feature related parameters when the configuration was uploaded via TFTP and subsequently downloaded back to the switch (10618).
- The CLI disable telnet command was not maintained across resets or in the TFTP upload/download configuration command (12176).

Issues Resolved

- When creating protocols for protocol based VLANs, the “LLC” type could not be set to 0x0000 or 0xffff. The switch now accepts these values as valid LLC parameters (12186).
- OSPF multicast “hello” packets Cisco EIGRP packets were not being forwarded by an L2 Extreme device if the qos profile for the VLAN passing the EIGRP packets was set to a qos profile other than QP1 (12179/12206/12171).
- When DLCS functionality was enabled on a switch, BootP Relay packets would not be forwarded from the client to the requested server (12255).

BlackDiamond

- Changes have been made to power supply (PSU) detection to allow the system to power up with a faulty power supply. Modifications made to an earlier release of EW 6.1.5b20 would attempt to detect a PSU based on a signal detection mechanism. Power supplies that are unable to correctly signal software would consequently not allow the system to power up. A message will be posted to the log if a user has a PSU that cannot be correctly detected although the system will be allowed to power up (11536).
- The BlackDiamond software power budget values have been modified to correctly enable 8 G8Ti modules to be powered up 220V power supplies (11974).
- A software exception would occur when modifying the IP protocol ethertype value (i.e., “configure protocol IP add etype 8035”) on a BlackDiamond with at least one non-“I” I/O module installed (11686).

Alpine

- On an Alpine3804 switch, the ACL flow rule hit counter (“show access-list-monitor”) would continue to increment even when there was no traffic present for that rule (11591).

STP

- Spanning Tree configured with Load Sharing would not properly restore a load share trunk back to forwarding state from disabled state when a load share group loses all active ports and recovers those active ports (11198).
- In an protocol sensitive VLAN environment, STP configs would not be able to add ports to the VLANs when the configuration would be uploaded and subsequently re-downloaded to the switch when the protocol sensitive VLANs contained the same ports as the STP protected VLANs (10924).

Port Mirroring

- A software exception would occur when upgrading to EW6.1.7 from EW6.1.5 and attempting a show mirroring command (11725).

BGP

- A software exception would occur when a new network was added to a peer group causing a route table update (11926).
- The CLI show bgp neighbor x.x.x.x transmitted-routes all command did not display aggregated routes transmitted to a neighbor (11538).
- A software exception could occur when displaying the IP route table at the same instant that the route table was being updated (11789).

ESRP

- The ESRP port restart command would not save across switch reboots (11394).
- The ESRP election algorithm setting would not be saved across reboots (11823).
- Unknown unicast packets would not be forwarded in domain-member VLANs upon domain-master failover (11352/10769).
- When a Summit stackable was configured with 4 ESRP groups on a single tagged port and the configuration was saved and the switch rebooted, the ESRP MAC address would be installed correctly on the first 3 groups only. This also occurred with a TFTP download of the configuration file and a subsequent reboot of the switch (11752). The workaround was to disable ESRP and re-enable it on the affected VLAN.
- In a split ESRP configuration where the protected VLANs were load shared between the ESRP routers or the ESRP interface was down for that VLAN, DHCP/BootP Relay responses would not reach the requesting client if the response was originated from a VLAN that was on the same subnet as that “down” ESRP IP interface (12170/12352).

OSPF

- The “noadvert” parameter in OSPF address range configurations would not be correctly uploaded via TFTP, causing the address range information to be lost upon a subsequent download of the config file (10962).
- The OSPF “wait time” value would not be set back to default when OSPF was unconfigured from an interface (10795). All other values would be correctly set back to the default values.
- Removing a configured OSPF inter-area filter would not be applied until OSPF was disabled and subsequently re-enabled on the router (11086).
- The CLI `ospf show lsdb` command would fail to display the 2nd screen page of the LSAs if there were many external LSAs installed in the routing table (11581). This was only a display issue and the LSAs would be correctly installed in the LSDB.

RIP

- The “routetimeout” parameter would not be preserved when the configuration file was uploaded via TFTP and re-downloaded to the switch (10968).

PIM

- A problem was fixed with memory management in PIM-SM configurations where multicast frames were incorrectly using memory allocations of the router (11714).

General IP

- With multinetting enabled, if a host port in a multinetted VLAN was moved from one port to another without the active link of the original port location going down (i.e., through a repeater or another switch), that host would be incorrectly mapped to the original port in the internal software tables resulting in a lack of connectivity to the multinetted interface (11549).
- With BootP relay configured, the switch would forward a BootP message to the client with a broadcast MAC address as opposed to the correct unicast address causing the DHCP Inform to fail (11348).
- ICMP Requests with the “from” flag set would not be sent with the requested source IP address (11637).

Issues Resolved

- The enable ipforwarding return message has been modified to indicate a concise error explanation when enabling IP forwarding on a VLAN without a configured IP interface (10670).
- A BootP obtained default route could not be deleted by the administrator (11569). The following message would be printed to the system log when this was attempted:

```
ipForwardEntrySet: ipForwardType failed to destroy (2)
```

- ICMP redirects are now sent for packets other than ICMP packets (10368).

IP TOS

- The diffserv replacement feature would not set all 6 bits in the code point from 1 to 0 (10877).

QoS

- QoS parameters (QoS configuration and profiles) would not be correctly applied to load sharing ports if load sharing was enabled prior to the QoS configuration changes being applied (11938). If the QoS configuration parameters were applied prior to load sharing being configured and enabled, the QoS parameters would correctly take effect.
- QoS configuration changes for minimum and maximum bandwidth value modifications when applied to physical ports (i.e., “configure qp1 minbw 0% maxbw 0% priority medium 2:9”) would not be saved across resets (10923).

Access Control Lists

- When a switch was configured with IP, TCP, or UDP ACLs, the following message would be printed to the system log when a module was removed from the BlackDiamond (11758):

```
<CRIT:HW> Twister access timeout slot=4
```

SLB

- The CLI command `enable slb node [<ipaddress>:<port>]` and `disable slb node [<ipaddress>:<port>]` command now allow the user the ability to enable/disable a node by specifying the L4 port name or number (11109).

SSH

- When SSH was enabled on a switch, no SSH sessions would be allowed if the maximum of 8 telnet sessions were already open with the switch (12000).
- When using Microsoft, secureCRT, or F-Secure-based SSH clients, the SSH sessions will remain open if 8 SSH sessions are initiated and idletimeouts are then enabled. After the configured 20minute idletimeouts session termination, the clients would correctly drop the sessions but the switch would keep the sockets in the incorrect state (10989).
- When logging into a switch with SSH clients, usernames that contained the “@” symbol would be rejected by the switch (11817).

TACACS

- TACACS configuration parameters could not be unconfigured using the CLI `unconfigure tacacs` command (11999).

RADIUS

- On switches with Radius enabled, sockets would be opened and not closed correctly, causing Radius to not be able to correctly authenticate until the switch was rebooted (11039).
- Many thousands of Radius logins could cause Radius sockets to not be released completely, causing Radius to not be able to correctly authenticate (10702).

SNMP

- When using SNMP to set the minimum bandwidth for a queue on a port (PerPortQosTable), if the parameter is set to 90 or greater the switch did not return an error (over the 90% limit for the MinBw setting) (11548).
- The dot1dTPFdbTable has been modified to only return MAC address information as opposed to MAC and IP address information (10014).
- When the dot1dTpFdbTable was enabled and ESRP was configured and running on a switch, a MIB walk of the dot1dTpFdbTable with multiple entries would cause ESRP instability (11702). This has been addressed so that the MIB can be accessed without causing ESRP to transition from Slave to Master status and vice-versa.
- The value returned for the object dot1dStpPortPriority would be reported as twice the correct value returned by the CLI (11673).
- When using a Basic Layer3 license on the Summit stackables and Alpine switches, the mib2.ip.ipForwarding.0 variable would report that IP forwarding was disabled even when it was enabled (10147/11400). The CLI would correctly display the configuration of the IP forwarding parameter.
- Modifying the “minbuf” parameter via the SNMP QosProfileTable resulted in an invalid entry to be set for the “minbuf” setting (11763)
- The enable and disable parameters for the dot1dTpFdbTable would not be correctly uploaded via TFTP, causing the user configuration parameters to not be maintained upon a subsequent download of the configuration file (12316). The command would be correctly saved across resets.

Issues Resolved from v6.1.7b5

The following issues were found in ExtremeWare v6.1.7b5 and resolved in ExtremeWare v6.1.7b7.

General

- When upgrading to ExtremeWare 6.1.7b5, the configuration for enable/disable commands for Radius, TACACS, Radius/TACACS accounting, and the SNMP trap receiver parameters would be modified in the upgrade process. Re-enabling these parameters or re-entering the SNMP trap receiver information and saving the configuration to flash would result in the correct configuration to be maintained (11447). However, in order to avoid having to modify these parameters when upgrading to ExtremeWare v6.1.7b7, it is required that you upgrade directly from v6.1.6b19 to v6.1.7b7 or re-download the desired configuration to the switch after upgrading to v6.1.7b7 from v6.1.7b5. Upgrades directly from v6.1.7b5 to v6.1.7b7 may result in the above configuration parameters being modified after a reboot.

IP Multicast

- The CLI show igmp snooping <vlan_name> command would only display the IGMP membership entries for the 1st port in the VLAN. Any additional ports in the VLAN would

Issues Resolved

display the memberships for the first group only (11455). The `show igmp snooping` command (without the VLAN name parameter) would correctly display the receivers on each port.

Issues Resolved from v6.1.6b19

The following issues were found in ExtremeWare v6.1.6b19 and resolved in ExtremeWare v6.1.7b5.

General

- The switch would respond to a ping when the router interface was down [i.e., had no active ports] (10954).
- Changes made to protect against inconsistent MAC addresses between Summit series stackable switches and Inferno series stackable switches (10928).

Alpine

- When a slot was configured for an FM24Fi (configure slot <slot number> FM24F, the TFTP uploaded configuration. Would display the slot as "new-type-43" which is an incorrect value for this I/O module (11956). The file could be modified to change this parameter to "FM24F" as a workaround.

Bi-Directional Rate Shaping

- The loopback ports were incorrectly counted as active ports in ESRP VLANs (10742).

Access Control Lists

- If an I/O module was not installed in slot 1 of a 3804 or 3808 system, the ACL hit counter would not increment (11103).
- If the precedence number of a TCP permit-established rule was automatically re-arranged by the system because a new rule was inserted by the user, the permit-established rule would no longer take effect. If the configuration is saved and the system rebooted, the rules will take effect as configured (10685).
- The precedence parameter for an ICMP rule was not being uploaded via a TFTP upload of the configuration file, resulting in the subsequent download of the configuration to not take effect (11104).

ESRP

- Modifications made to the ESRP Awareness function to ensure the switch correctly points to the Master ESRP router after multiple ESRP Master router changes (11092/10969). Learning the MAC entries on the ESRP Aware switch has been delayed by 2 seconds to ensure the ESRP Aware switch is pointing to the ESRP Master switch.

IPX

- The server type for a static IPX service was changed from a hex value to decimal value when uploading the configuration file via TFTP. As an example, this would cause a server type of 0078 to be listed as 120 in the server table after a subsequent download of the configuration file (11029).

Spanning Tree

- The spanning tree port info always displayed the port as being in “forwarding” state when using the “show stpd <spanning tree name> ports <port list>” command when the port was not active (10801).

OSPF

- When using OSPF export direct to export ESRP VLANs, ESRP state changes from Master to Slave and back to Master would cause OSPF to not export the directly attached interface to neighboring routers (10367).
- The address range “noadvert” command was misspelled via a TFTP upload configuration resulting in the subsequent download of the configuration to not take effect (10962).
- The “enable ospf export vip” command was not being uploaded via a TFTP upload of the configuration file resulting in the subsequent download of the configuration to not take effect (10963).
- The “configure ospf spf-hold-time” command was not being uploaded via a TFTP upload of the configuration file resulting in the subsequent download of the configuration to not take effect (11014).

BGP

- After learning and installing more than 97,000 routes, a router running multihop EBGP could stay in pend-start state when losing and recovering the link to a neighboring router (10959).

SNMP

- The dot1dTpFdbTable can now be enabled by the user. Please see new features section for description of configuration parameters (9394/9543/10014/10955).

Vista

- A software exception would occur when a banner was configured with extended ASCII characters [characters with an index in the ASCII table greater than 128] (11075).

Issues Resolved from v6.1.5b20

The following issues were found in ExtremeWare v6.1.5b20 and resolved in ExtremeWare v6.1.6b19.

General

- The LEDs on the Summit 48i Gigabit redundant PHY’s did not properly indicate activity between the primary and redundant physical interfaces (9239).
- The Summit 48i would not correctly bring link up between the primary Gigabit ethernet port when the link was physically moved from the primary port to the redundant port and back (10949).
- LX-70 GBICs are now recognized as “LX70” in the CLI show port commands and Vista screens. Note that for SNMP, the RFC for the ifMauMib does not differentiate between “LX” and “LX70” so when viewed from the MIB, the GBIC type will always appear “LX” (10591/10864).
- On the BlackDiamond and Alpine switches, when transmitting packets in bursts sizes greater than 160 packets per burst from a Gigabit port and a 10/100 port, users could experience packet loss. The buffers for the 10/100 ports on these 2 platforms have been adjusted for improved performance and decreased packet loss (10807).

Issues Resolved

- Configured dot1p to QoS mapping would not be preserved in a TFTP upload/download of the configuration file (9796).
- The show port utilization command could display a higher value than 100% in the Peak RX Bandwidth % column (9985).
- When using the save configuration command, anything but a lower case "y" would be interpreted as a "n". A capital "Y" is now accepted (10021).
- The Summit 1iSX/1iT now ship as Full Layer3 capable per the part number ordered (10054).
- A user level administrator could execute a save config if q "quit" was performed and a subsequent save config was attempted (10133).
- The show edp command would display the incorrect port number for load-shared links. The command now correctly displays the load-share master port (10569).
- Jumbo frame configuration parameters were not save across resets on 1000BaseX and 100/1000BaseT ports (10492).
- The following messages would be printed to the syslog and has been removed (10513):
PORT: ERROR: Slot 3 port 2 lane 1 spurious int 8

BlackDiamond

- In BlackDiamonds with mixed non-"i" and "i" series I/O modules, streams from non-"i" series I/O modules would not be forwarded to "i" series I/O modules when the forwarding entries for those streams would be re-added to the table (10020). The Layer 2 forwarding database (FDB) in this instance would appear to be correct and clearing the Layer 2 fdb for the specific entry would resolve the forwarding problem.
- In BlackDiamonds with mixed non-"i" and "i" series I/O modules, the internal table entries in the IP forwarding database (IPFDB) would force some entries to go through the CPU, resulting in a high tNetTask utilization and degraded network performance and loss of connectivity between hosts (10930).
- On "i" series I/O modules in a BlackDiamond, creating a permanent FDB entry between 2 hosts already learned dynamically by the FDB would cause a loss of communication between those hosts (10480).
- Enabling and disabling G1 support now prompts the user with a warning that the system must be rebooted (10095).
- In BlackDiamonds, a software exception could occur in the tBGTask if there were many L2 FDB entries and IGMP snooping was enabled on the switch (10112).
- The show port utilization screen displayed manually configured link speed of a 10/100 port as "AUTO" (9383).
- The following message would be printed to the syslog for the first care in the chassis upon bootup. This message has now been removed (9813):

<WARN:KERN> Cannot send packet out slot 1. Card not present.
- The system health check configuration parameters were not uploaded as part of a TFTP upload for an offline configuration (9868).
- When the system-health checker was enabled "default" VLAN had a user-configured tag, the following message being printed to the log for all inserted I/O modules (9908):

<CRIT:SYST> Packets are missed or corrupted between MSM-A and slot 4.
The problem needs to be fixed immediately

- A static FDB entry on a BlackDiamond with learning disabled will cause the BlackDiamond to not initiate ARP requests for the client configured as the static FDB entry. Communication between hosts is not affected (10315).
- Modifications were made to the fan checking of MSM64i's to ensure false fan failure messages would not be printed to the syslog (10652).

ALPINE

- Jumbo frame configuration would not be maintained across resets (9829).

VLANs

- The "R" flag in the CLI `show vlan` command was being used for both the "SubVLAN IP Range Configured" and "IPX RIP Enabled" options (10593).
- The "subvlan address range" option was not being uploaded via TFTP and subsequently re-downloaded to the switch correctly (10672).

Port Mirroring

- In some environments, enabling mirroring on a port before the port became active would result in a software exception. This would happen when ports were configured to point to the mirror port and the mirror port became active and was enhanced further in this release (9553).

General IP

- Pinging the switch's local interface with a large packet size (i.e., 8000 Byte packet) could cause a software exception (10384).
- Improvements have been made to the general IP forwarding database functionality for optimization (10761).
- When receiving IP ARP packets prior to having the route for the network installed, the following message would be posted to the syslog (10873/11093):

```
<WARN:IPHS> IP FDB entry not added as no route is available
```

- The `enable igmp` command did not enable IGMP on all of the VLANs if the "MGMT" port had an IP address configured (9811).
- The switch would allow users to enter duplicate proxy arp entries (9904).
- On switches with multiple MAC addresses associated with IP FDB entries, upon aging the MAC addresses per the user-specified MAC aging-timer, the system could become unresponsive as it was flushing the MAC entries in the database. Subsequently, a message would be printed in the log indicating that the `fdbAgeTask` had consumed a high amount of CPU utilization (9953).
- The following message could be printed to the syslog when BootP reply was forwarded to the BootP relay agent on the switch (10539):

```
<WARN:IPHS> IP FDB entry not added as no route is available
```

- The CLI `clear counters` command would not clear the "bad protocol" field when issuing a `show ipstats` command (10470).
- With IP forwarding enabled, when a host was moved from an active connection via another switch or repeater to a new directed-attached port within the same VLAN, the IPFDB would not

Issues Resolved

correctly update the new location of the host unless the FDB was cleared by the user (11618). The traffic was correctly forwarded within hardware and this was a cosmetic port display issue.

General Routing

- A route with a metric greater than 65535 would not be installed in the routing table (9423).

IP Multicast

- Whenever a link was disabled or disconnected on a VLAN with IGMP snooping entries, the IGMP snooping table would be re-initialized and could cause the multicast stream to appear to have jitter (10391).

OSPF

- Type 5 LSAs would not contain the correct forwarding address if multiple default or static routes existed and the entry in the LSA was manually removed. This issue would also occur with Type 7 LSAs in an NSSA environment (9425).
- A default route with a lower cost is not selected among multiple default routes in an NSSA configuration (9143).
- A type-7 LSA for a default route may be removed from and not re-installed in the routing table if the router interface is moved between the NSSA and the backbone area and back to the NSSA (9426).
- A type-7 LSA for a default route is removed from the routing table and not re-installed when a new router interface is added to the NSSA (9427).
- A type-7 LSA for a default route is flushed from the routing table in the NSSA routers when the originating router receives a new ASBR router LSA (9433).
- A type-7 LSA for a default route is not re-installed in the routing table of an NSSA ABR LSDB after the default route is deleted from the neighboring router (9454).
- OSPF hello packets were not being forwarded by intermediate routers when the master load share port was removed between the intermediate router and the router generating the hello packets (10568).
- The initial wait time was not updated with a user configured Router Dead Interval below 40 seconds (9882).
- If a Router Dead Interval below 40 seconds was configured on a router, the neighbor's interface information would display a negative integer for the dead time interval (9909).

BGP

- On a BlackDiamond with a full internet routing table, typing `show bgp neighbor <ip address> transmitted-routes` would cause a software exception (9952).
- The BGP MD5 Password was displayed in clear text in debug messages, Peer group definition, and uploaded peer group configuration (10214).
- The following messages would be printed to the syslog when configured route-maps were being utilized by the system(10895):

```
<DEBUG:SYST> New callback = 0x8b23ff10 handle=2 vftp=0  
<DEBUG:SYST> bindToRtMap(1, 0x0, 3010106, 4124)
```
- When the output route-map contains no ipaddress-based match statement, the show command did not apply the route-map did not show up in the display (10865).

RIP

- Host routes were not being exported correctly from a RIP network (10308).
- The RIP VLAN cost configuration option was not uploaded correctly via TFTP and would result in the value being set to default upon a re-download of the config file (10894).
- When modifying the RIP configuration commands with the “vlan all” option on systems with a “MGMT” port with no IP address configured, the option would not take effect on the VLANs configured on the system (10894).

DVMRP

- A software exception could occur in the tDvMrpTask if a multicast stream was using multiple egress VLANs while those VLANs were being pruned/updated (10324).
- The DVMRP cost configuration option was not uploaded correctly via TFTP and would result in the value being set to default upon a re-download of the config file (10811).

PIM

- The PIM CBSR configuration information was not uploaded correctly via TFTP and would result in the value being set to default upon a re-download of the config file (10867).
- After a multicast client would leave a VLAN, the last-hop switch would continue sending (*,G) to the upstream RP so that in the upstream switch a (*,G) entry (Z) was installed in the egress vlan as well as a (S,G) entry with (Z) causing multicast traffic flow problems (10856).

ESRP

- Protocol based VLANs and host attached clients in an ESRP configuration did not remove the ESRP MAC address from 10/100 ports on a BlackDiamond (9831).
- When active ports in a VLAN on an ESRP Master switch would go to 0, the ESRP Master would not properly transition to non-Master state. The selection algorithms would work correctly if ports transitioned above a 0 port count (9993).
- Debug messages would be printed to the syslog if ESRP was configured in conjunction with VLAN Aggregation (9984).
- Deleting a tracked ping from an ESRP VLAN would print the following message to the syslog (9994):

```
<WARN:SYS > HC: serversAliveTask: select problem (-1).
```
- An ESRP configuration with thousands of active interfaces (16,000+) would cause a failover attempt between 2 ESRP routers to take a large amount of CPU cycles, causing the system to appear unresponsive during this time (10164).
- If a critical task software exception occurs on ESRP switches, ESRP will no longer send out ESRP hello PDUs forcing the slave ESRP switch to immediately take over master responsibilities (9855).

Access Control Lists

- When creating a access-list with rate shaping on a BlackDiamond, the first port on an F48Ti configured as a loopback port caused traffic within the same VLAN (on the same module) to not be forwarded correctly (10651).
- The show access-list-monitor command would only display 45 entries (10250).
- Access control lists were not functional on a Summit 5iSX/5iT platform (10338).

Issues Resolved

- The hit counter for ACLs would not increment in a BlackDiamond if there was not an I/O module inserted in slot 1 even though the rule was being applied correctly (10337).

Server Load Balancing

- With SLB configured on a switch, the 3DNS iQuery packets (which manage global load balancing) could cause a software exception (9876).

Web Cache Redirection

- When using traffic generation equipment to test the performance parameters of WCR, the switch forwarding database would not update the entry for the simulated cache server port causing the traffic to go through the software. If a traffic flow was configured between the origin server and the cache server, the traffic would be forwarded at wire-speed rates (9263).

SNMP

- The ifMauType now returns "1000BaseSXF" for the BlackDiamond WDMi module (9481/10681).
- Optimizations were made to enhance the lookup speed of the OspfMib (9972).
- VLAN protocol information was missing when querying the MIB. VLAN protocol information was correctly identified in the CLI. This was order dependent on how the VLAN was created and would only be seen in some systems (10501).
- The SMMi returned a "0" value for the extremeSlotModuleSerialNumber (10620).

DLCS

- DLCS bindings were not displayed in the CLI `show dlcs` command (10512).

TACACS+

- The config tacacs server and shared secret parameters were not maintained in a TFTP upload and subsequent download of a configuration file (10227).
- TACACS+ configured to authenticate using the Microsoft server NTTacPlus 2.02 build 2 (Jan 2,2000) gives Invalid Authen/Start packet error and could result in a software exception (10366/10528).

VISTA

- The Route Report Interval and Route Timeout Interval on the DVMRP page could be configured above the allowed values causing the CLI to display a negative value for these options (8619).
- A user was able to incorrectly set a 100/1000 BaseT port to 10Mbps although the configuration would not be updated in the switch software (9873).
- The unconfigure option for PIM VLANs would not restore default values for the CRP timer (9933/10662).
- The BlackDiamond G8Ti I/O modules can now be added via the web interface (10056).

Issues Resolved from v6.1.4b20

The following issues were found in ExtremeWare v6.1.4b20 and resolved in ExtremeWare v6.1.5b20.

General

- Enabling ports on an I/O module while that I/O module was removed from a chassis which ports were disabled on while inserted would cause the I/O module LEDs to appear as if the board was disabled even though it was correctly forwarding traffic. This existed on the BlackDiamond and Alpine (8683).
- The time stamps in the system log were not always being displayed in the correct sequential order (8864).
- The disable port command could re-enable a port if attempted 2 times on the same port (8884).
- Spanning tree parameters would not be maintained across a save and reboot in s0 for the Summit 7i only and any user-defined spanning tree instances for all switches (8989).
- A telnet session will no longer be terminated upon the receipt of a null ASCII character (9053).
- 3 consecutive failed login attempts would display an invalid user name in the system log as the failed user (9591):

```
<WARN:SYST> User account instance 65535 out of range
<WARN:USER> Login failed for user admin44176 through console
<WARN:USER> Login failed for user admin44174 through console
<WARN:USER> Login failed for user admin44172 through console
<INFO:SYST> User   P aH k   logged out from console
```

- The Summit 5iTx switch would report the UTP port type as “SX” in the CLI show port commands (9370).
- If characters were sent to the console connection with a terminal or terminal sever after a reboot of a switch and when the operating system starts to load, the system could reboot or go to the Bootrom menu. This was dependent on the character that was being transmitted to the console at the time (9721).
- The “disable sharing” command would not display which slot:port loadsharing was disabled on in the system log of the BlackDiamond and ALPINE (9264).
- On a Summit 7iTx, ports 1, 5, 9, 13, 17, 21, and 25 could not receive packets when autonegotiation was disabled on those ports, the configuration saved, and the switch rebooted. Re-enabling auto-negotiation on those ports would not resolve the problem and an “unconfig switch all” would have to be performed to allow the ports to become operational (9302).

BlackDiamond

- Under certain circumstances, if a G6X is installed in a single MSM chassis with that MSM in slot B, the G6X would not be able to forward routed traffic between ports. If the MSM was in slot A or if a dual MSM system was used (with either slot as Master), this problem did not occur (8740).
- A software exception could occur if MSMs are being hot-swapped with an active management port connection (8923).
- When using jumbo frames, the BlackDiamond would only forward frames that are 4 bytes less than the configured jumbo frame MTU size when going across I/O modules. Within the same I/O module, frames with the configured system MTU size would be forwarded correctly (9431).
- Non-existent ports could be displayed in the CLI “show vlan” command when configuring a VLAN with ports that span multiple I/O modules (9628).
- A software exception could occur in certain instances while hot-swapping an I/O module when a packet that was destined for the switch processor was being processed on an incoming port (9310/9510/9584).

Issues Resolved

- A software exception could occur if the slave MSM was the only MSM with an active “mgmt” port link (9483).
- The BlackDiamond would intermittently report false power supply failure and recovery messages although the power supplies were operational (9240).

Alpine

- A software exception could occur on an ALPINE if a command used many arguments – for example, “config default delete port 2:1, 2:2, 2:3, 2:4, 2:5, 2:6, ...” (9063/8697).
- Enabling a load share group across FM32-Ti modules would display an error indicating the port types in the group are not similar (9524).

VLANs

- The following error message would be printed to the console if a user created 3000 VLANs and configured a single port to be added as a tagged port to each VLAN (9328):

```
<WARN:KERN> Unable to allocate otp block for index 1091 slot 1
```

ESRP

- ESRP “Awareness” on a configuration with ESRP groups was not working correctly and flushing the FDB on ESRP VLANs (8907).
- If the ESRP VLAN was configured for the priority-mac-only election algorithm, the active ports would not be counted correctly and the ESRP state would come up as “neutral” (9157).
- ESRP transition changes are printed to the log without configuring debug-trace parameters (9491).
- The “clear counters” command now clears ESRP counters (9588).
- The correct number of active ports was not being reported in ESRP if the VLAN was first created, ESRP enabled for that VLAN, and then ports added to the VLAN (9006).

OSPF

- A software exception would occur when enabling OSPF export of static routes and then configuring identical static routes with different subnet masks (7838).
- If the “mgmt” port had an IP address assigned to it, a user would not be able to configure OSPF authentication for area 0.0.0.0 (8875).
- Debug messages will not be printed to the log for interfaces that do not have OSPF configured on a router that is running OSPF (9076).
- When ase-summary is configured on an ASBR, a type 5 blackhole route is installed in the routing table to remove aggregated external routes within the ase-summary network which are originated by that ASBR (9361).

RIP

- The rip export direct values for cost and tag are reversed when performing a TFTP upload of the configuration and a subsequent download of that configuration (9616).
- The “enable rip export direct cost <number> tag <number>”, “enable rip export vip cost <number> tag <number>”, and “enable rip originate default” commands were not being restored

when using TFTP upload and a subsequent download of the configuration (9451). These functions did restore correctly across a save and reboot.

- In a mixed ESRP/OSPF/RIP configuration, under certain conditions with OSPF exporting RIP routes, the better RIP advertised route would not be used (9269).

PIM

- In a PIM-SM configuration, if an IP multicast sender transmits 2 streams with more than 150 destinations, a software exception could occur (8999).
- When a physical link between an upstream switch (which is also the BSR and RP) and a downstream switch was taken down, the volume of "join" messages being sent from the downstream switch could cause routing protocol instability while the neighbors were attempting to become operational (9457).
- When a router received a PIM prune packet reflected back to a port that initiated the packet, the router would eventually become unresponsive (9290).
- In PIM-SM, when the least cost backwards route to the source was returning in the table, the last hop switch should have pruned the (*,g) and (s,g) to the upstream switch on the older backward route (9267).

General Routing

- When using the "show iproute permanent sorted" command, the switch console could become unresponsive if a large number of routes existed in the routing table (9635).

General IP

- Traceroute results could indicate a negative number for round-trip times (8159).
- IRDP would not be enabled unless explicitly enabled for specific VLANs on systems with a "mgmt" port (9306).
- The "enable ipforwarding broadcast" global command would not enable IP broadcast forwarding on all VLANs after a VLAN interface that failed to accept the command due to a configuration error (9261).
- A software exception could occur when using the "show iparp" command if a static IP ARP entry was created and the corresponding router interface was deleted (9042).
- Adding a VLAN without an IP address and BootP enabled would cause other VLANs in the multinetting group to not forward traffic correctly (9286).
- When using proxy ARP, the "always" flag had to be used to enable proxy ARP functionality across VLANs/networks. The proxy ARP "always" flag is no longer necessary to use proxy ARP functionality across VLANs/networks (9384).
- A software exception could occur if a malformed IP packet with the 2 bytes of 0's was placed in front of the "0800" IP protocol field (9567).

IP Multicast

- The following debug message would be printed to the system log when multiple IP multicast senders were present on the network (9583):

```
<WARN:KERN> tNetTask: Block 4018 on slot 6 is already free
<WARN:KERN> tBgTask: Block 4018 on slot 4 is already free
```

Issues Resolved

IPX

- A software exception could occur when a Fluke LanMeter generated a unicast GNS request with a server available message (9718).

QoS

- The “minbuf” parameter for QoS profiles can no longer be modified to a value greater than 100% (6837/9274).
- When configuring MAC QoS for first generation I/O modules, the QoS profiles would not be assigned in the correct order to the 4 hardware queues on first generation modules (8881).
- The Peak BW setting on a general QoS profile was always set back to default after a save and reboot (9609).

Bi-directional Rate Shaping

- If an FDB entry for a rate-limiting port had already been learned and a port was added or deleted from the VLAN, the fdb would have to be cleared before the rate-limiting port entry could be re-learned (8429).
- Spanning Tree would not protect against network loops on ports within the same VLAN configured as rate-shape ports (8649).
- A software exception would occur when downloading a configuration to a switch with Spanning Tree enabled on a VLAN configured for rate shaping (8657).
- If a VLAN created with a rate shaped port was deleted, that port could not be re-added to another port (8824).
- A software exception could occur when configuring an VLAN with a tag that was identical to the configured loopback-id for that VLAN (9810).
- A VLAN configured for rate-shaping would display the VLAN/router interface as administratively up if a loopback port was configured with no active ports (9528).

Server Load Balancing

- A software exception could occur if a service health check was configured on a VIP for FTP and then later configured for another health check (i.e., HTTP) (8898).
- A new command has been added to SLB to allow for the configuration of the global connection timeout period. This helps to avoid cases where connections would be closed when the TCP “FIN” and “ACK” timeout was too short (9487/9613). Please see the “New Features” section for command syntax.
- The username for an SLB health check could conflict with the configured switch user accounts. The username can now be any valid alphanumeric name (9539).
- SLB client persistence settings were not functioning after a save and reboot even though the configuration parameters were visible from the CLI (9139).
- Memory management for SLB was not being handled correctly. This could be viewed using the show memory command in the “SLB Other” allocations column (9409).

Policy-Based Routing and Web Cache Redirection

- The name of configured flow policies was not being displayed in the “show flow-redirect” command when using specific destination IP addresses (9047).

- Flow policies configured with a destination IP port of “UDP” were not being properly redirected (9047).
- When a flow policy is enabled and subsequently disabled, it does not properly redirect when re-enabled (9255).

SNMP

- A new trap has been added for EDP in the Extreme private MIB to indicate the addition or deletion of an EDP neighbor (7126).
- The ifMauTable and ifMauAutoNegTable are now supported in read-write (8878).
- The ifMau MIB displayed a redundant physical port for each port on all Summit stackables (9257).
- When more than 30 VLANs were configured on a switch, SNMP can time out when accessing the switch. This affects EEM as well as off-the-shelf SNMP applications (9002).
- When 3000 VLANs were configured on a switch, multiple MIBwalks using SNMP get and get-next requests would cause ESRP transitions to occur (9280/9448/9450/9818).
- An SNMP “bulkget” request would cause high CPU utilization for the SNMP task and cause the system to not process routing protocols without timing out (9325). This could affect ESRP and OSPF among other functions.
- When executing a “get” on the dot3Stats MIB, error counters from one index would be inadvertently reported for the next index as well (9004).
- The log severity level of an SNMP task utilization message has been reduced to “WARN” from “CRIT” (9251).

Access Control Lists

- When a flow rule was created that had an ICMP rule for the same network, the flow rule would not be applied when the FDB was cleared (8637).
- ICMP access list rules would interpret the value of “0” for type or code as a wildcard “ANY” entry. The value of “255” is now interpreted as “ANY”. For example, if an ICMP deny rule is created with type configured as “0”, all ICMP traffic with any other type will be blocked (8696/8317).
- The “show config” command was not displaying the correct QoS queue for ip permit rules even though the correct queue was being used and displayed in the “show access-list” command (9417).

Port Mirroring

When using port mirroring on an ALPINE, the switch’s MAC address would be transmitted in reverse order from the mirror port (9542).

DLCS

- The “show dlcs” command would show ports that did not exist on an F48Ti I/O module (8972).

Radius

- Repeated failed login attempts could cause all connections to be refused (8915).

Issues Resolved

Vista

- The Vista statistics page has been modified to display OSPF external routes correctly. Previous to this release, multiple entries for the same network would fill the Vista statistics page (7752).
- If two default routes were configured for a switch, Vista would only display the first default route 2 times (8549).
- Port settings can now be correctly configured through Vista (9236/9579).
- Web access is now disabled by default in the SSH2 based software (9853).
- When using Vista to view permanent entries in the FDB table, a “bad action type” message would be posted to the browser (9396).

Issues Resolved from v6.1.4b12

The following issues were found in ExtremeWare v6.1.4b12 and resolved in ExtremeWare v6.1.4b20.

General

- The following message has been removed from the log (9552):

```
Received broadcast router discovery solicitation from <IP_ADDRESS>
```
- The following message has been removed from the log (9568):

```
<WARN:PORT> ERROR: Blizzard link up, phy link down
```
- A console connection would be prematurely terminated when idle-timeouts were enabled and a user would login and logout of the switch multiple times. This would result in the user connection to close prior to the 20 minute idle-timeouts timer value (9340).

BlackDiamond

- Under certain configurations, a BlackDiamond could reflect inbound unicast or multicast packets back to the port it received the packet from and to other ports within the same VLAN. This could cause connected switches to update the wrong entries in their forwarding databases and affect network performance (9388).
- In certain conditions, when a slave MSM is removed and re-inserted, unknown unicast, multicast and broadcast packets would be transmitted multiple times out of all of the ports in the VLAN causing what appears to look like a broadcast storm (9534).
- Fan recovery messages were inadvertently posted to the syslog after the switch had been running for a period of time. This was due to an incorrect reading of the fan status and needed to be enhanced for all Extreme fan manufacturers (9241).

```
04/06/2000 10:03.57 <CRIT:SYST> Fan(3) is back to normal
```

Alpine

- A link configured for 10Mbps half-duplex was not setting the duplex mode correctly until the link was removed and re-inserted on the FM-32T I/O modules (8496).

VLAN Aggregation

- On the BlackDiamond, when the only port (or last port) in a sub-VLAN was removed or moved to another sub-VLAN, inbound traffic requiring the switch processor to update internal tables

would not be processed correctly. This would result in such behavior as ARP timeouts and loss of connectivity with end-stations in those sub-VLANs (9509/9550).

SLB

- When configured for SLB, a software exception could occur when a client FTP control packet with an abnormally long "PORT" command was transmitted to the switch (9317).

Port Mirroring

- Enabling mirroring on a port before the port became active would result in a software exception. This would happen when ports were configured to point to the mirror port and the mirror port became active (9553).

Issues Resolved from v6.1.3b11

The following issues were found in ExtremeWare v6.1.3b11 and resolved in ExtremeWare v6.1.4b12.

General

- The forwarding database (FDB) entries for ingress rate shaped ports would not be updated correctly if the host was attached through another switch or repeater and was then moved to another physical port from the switch or repeater device (8650).
- For security reasons, the ExtremeWare copyright statement information will no longer appear when attempting a Telnet login (8666).
- The telnet access-profile configuration information was not operational when using a TFTP upload/download (8734).
- Disabling ports in a load share group would not display an error that the port is part of a load share group and would not provide the correct configuration status in "show port config" (8799).
- Protection put in to hide the telnet TCP port if telnet is disabled (8756).
- The following error message would be printed to the system log if a load share group was enabled, the ports were deleted from the VLAN, and the ports were re-added to the VLAN (8791):

```
ERROR: Port 1:2 is in load-sharing mode.
Perform all configuration changes to the logical link (1:1)
```

- On a Summit 7i, enabling ports with a link present on port 4 would result in port 2 also being presented as active and the LED on port 2 to also active (8570).
- A user could not login to the switch with SSH enabled with a password (8738).
- In certain cases, executing a "save" command via console could make the system unresponsive and a reboot would need to be performed to get access to the console (8710).

BlackDiamond

- In some cases, a hot insertion of an MSM could cause a system to become unresponsive or reboot after the MSM was seated. The following messages would also appear in the log (8712):

```
<WARN:SYST> Cannot send nmc message(3,8,0x0,0x0). Error=-1 (0xffffffff)
```

- In a 64 port Gigabit Ethernet system, some ports were unable to communicate with other ports in a full-mesh traffic pattern (8762).

Issues Resolved

- Fan recovery messages were inadvertently posted to the syslog after the switch had been running for a period of time. This was due to an incorrect reading of the fan status (8298).

```
04/06/2000 10:03.57 <CRIT:SYST> Fan(3) is back to normal
```

- With 2 MSM64i's in the system and the management port link connected and active, if syslog is enabled, removal of the slave MSM will cause a software exception (8777).
- During a TFTP download config, certain modules would not be ready and the port configuration would not be programmed correctly on those even though "show" commands displayed correct configuration. This would result in traffic not being forwarded correctly across those ports (8653).
- If IGMP snooping was disabled on a BlackDiamond, multicast packets within the same VLAN would not be flooded correctly across I/O modules but would be flooded within an I/O slot (8932).
- The link LED of a BlackDiamond Gigabit I/O module would not re-active under certain conditions if a remote port was disconnected and subsequently reconnected. The port would forward traffic correctly although the LED was not activated (8838/8856).
- If Spanning Tree was configured on an MSM64i and that MSM was subsequently moved to another active chassis as the Master MSM, displaying the Spanning Tree Domain info would show the Bridge ID to be the Bridge ID of the original chassis the MSM was configured in (8880).
- The "show fdb port <port number>" command was modified in EW 6.1.2b7 to display FDB entries associated with a particular port (8063).

VLANs

- The number of virtual ports have been increased to allow for a greater number of VLANs assigned to tagged a port(s) (8744).

Bi-directional Rate Shaping

- A software exception could occur after a system reboot if a configured loopback port had an active link attached (8866).

ESRP

- ESRP group information would not be uploaded and downloaded in a TFTP configuration file when ESRP was enabled (8514).

OSPF

- In a multi-ASBR environment, an ABR would not properly compute the SPF for all of the routes in the network, rendering the routing table as incomplete (8752).
- In a multi-ABR environment, an ABR would not properly compute the SPF for inter-area routes (8754).
- When a router was configured as an ABR but would lose link to one area, turning it into an internal router, the SPF would not be performed correctly for internal routes (8755).
- The OSPF export static cost could be set to a value greater than 65536 (8660).
- A default route is not correctly installed if a default external LSA with metric type 2 exists in the routing table (8702).
- The cost of an exported default route would be incorrectly doubled when forcing an SPF calculation to occur on the backbone area (8780).

RIP

- When configuring the RIP update interval timer values to 1 second, the router would send out the update intervals at 10 seconds. The default 30 second update interval timer would send out packets correctly at 30 seconds (8818).
- If a VLAN was not configured to be participating in RIP and RIP was globally enabled on the router, the show rip command would display that directly attached peer in the “show rip” display even though that VLAN was not configured with the RIP protocol (8813).

BGP

- If identical routes were learned through IBGP and EBGP, both routes would be displayed in the routing table. Only the EBGP route should be displayed (8760).
- With BGP synchronization enabled, a link flap with a large routing table could cause an exception error (8761).
- If a user repeatedly enables and disables BGP on a router or a peer, the system could become unresponsive (8767/8768).
- The TCP session with a BGP peer may not be closed if a user disables and enables BGP repeatedly (8766).
- With synchronization enabled, when EBGP routes were installed in the routing database and an IBGP neighbor was enabled, upon disabling the EBGP session, for each EBGP route removed from the routing table an IBGP route would be incorrectly installed (8773).
- BGP nexthop-self would not be correctly configured via TFTP if the configuration file was upload/downloaded to the router (8614).

IP

- The IP forwarding database would not correctly replace entries with a large number IP forwarding table entries (8758).

IP Multicast

- IP multicast entries for “i” series I/O modules would not be learned correctly with a large number of streams in multiple VLANs (8737).
- IP multicast packet loss could occur in a multi-router environment with multiple IP multicast streams (8652).
- The IP multicast forwarding database would not correctly replace entries with a large number IP multicast forwarding table entries (8759).

IGMP

- IGMP was not functioning correctly if IGMP snooping was disabled on an L2 switch. The L2 switch was not transmitting the group membership info to a directly attached L3 multicast router (8658).

QoS

- MAC based QoS is now operational in ExtremeWare (5509).
- Static FDB entries assigned to QoS profiles would result in that entry being a blackhole entry after a save and reboot (8842).

Issues Resolved

- On a BlackDiamond switch, first-generation I/O modules would not forward traffic if QP5 through QP8 were modified on the switch (8879).

SNMP

- When SNMP access profiles were created and then deleted, SNMP access would be blocked from the switch even though no access-profiles existed (7350).

Issues Resolved from v6.1.2b7

The following issues were found in ExtremeWare v6.1.2b7 and resolved in ExtremeWare v6.1.3b11. This list also includes issues that are fixed from v6.0.10b6 and not in ExtremeWare v6.1.2b7.

General

- SNTP could not resolve an IP address to hostname mapping using DNS (8456).
- The Summit 7i Power Status LED for PSU-A would not illuminate when connected to an active power source (8473).
- Additional modifications were made to telnet to protect against a software exception error under rare conditions (8483).
- When saving a configuration which had enabled, the switch would stop responding to user commands and the save would not take effect (8419).

BlackDiamond

- A TFTP downloaded config with G6X port configuration would not assign ports to the G6X module as stated in the configuration file (7468).
- The G6X module could not forward ICMP packets out of certain ports (8215).
- Enabling load sharing on a G6X after forwarding database entries had already been learned on those ports could result in uni-directional traffic flow across the load share group (8341).
- Show EDP would not properly display EDP neighbors on a MSM64i after a reboot was initiated (8472).
- Removing an I/O module with active ports could cause a software exception error (8509/8510).
- IP forwarding was not working correctly on the MSM64i between hosts on different VLANs (8568).
- MSM-B would not forward traffic for several minutes after a failover would occur from MSM-A (8608).

IP

- The traceroute "from" command designation was not going through the specified interface (8159).
- The traceroute command could display negative numbers for round-trip values (8159).
- Attempting to forward UDP traffic using a UDP-profile to a subnet with no matching route (i.e., destination VLAN is down) could cause a software exception to occur (8591).

Access Control Lists

- ICMP permit rules would not be displayed in the syslog if the "permit log" functionality was enabled (8397).

VLAN Aggregation

- If a secondary IP address was configured for a VLAN, an invalid router interface would remain in the 'show iproute' and 'show ipconfig' display of a switch if the VLAN was subsequently deleted without first removing the secondary IP address. A check now exists to ensure the secondary interface is first removed before deleting the VLAN (8458).

RIP

- Disabling RIP on a VLAN without a configured IP address could cause a software exception error (8445).

PIM-SM

- Running high volumes of multicast streams for extended periods of time could result in new entries not being added to the forwarding database (8414).

SLB

- The current connections value does not decrement when a new connection is established from a source address that already had a connection to a different destination (8483).
- When a SERVER vlan's slb-type was configured as "server", the IP address for the client originating the request was not added to IPFDB table when a telnet (or other TCP) session was initiated. This did not apply if the slb-type of the server's VLAN was configured as "none" (8494).
- The "disable slb" command did not remove slb-vip routes from the routing table (8576).
- The "unconfigure slb" command did not restore default values for pool member ratio, ping check for nodes, and http match string (8577).

SNMP

- Smarttraps were not sent out when a GBIC module was removed (8431).