




ExtremeWare Release Notes

Software Version 4.1.20b3



Contents

New Features in ExtremeWare v4.1	1
Released Software	1
Upgrading from v4.0 to v4.1	2
Upgrading to v4.1 from Releases <i>Prior</i> to 4.X	3
Summit Switches	3
Summit Hardware Requirements for Running ExtremeWare v4.X	3
Upgrading Summit switches	3
Downgrading Summit Switches	4
BlackDiamond	4
Upgrading the BlackDiamond	4
Downgrading the BlackDiamond	5
Upgrading BootRom	6
Upgrade Issues	6
Inserting Second v4.X-based MSM into v3.0.9 system	6
New RIP defaults	6
New Features in ExtremeWare V4.1	6
Improved Exception Handling	6
Enable and Disable Slots on Modular Switches	7
Automatic Module Disable on Hardware Failure	7
Enable and Disable IP ARP Refresh	7
SMON Support	8
Security	8
VLAN aggregation	8
VLAN aggregation properties	8
VLAN aggregation limitations	9
Security Option for inter sub-VLAN communication	9
VLAN aggregation CLI syntax	9


Configuring VLAN aggregation	9
Verifying the VLAN aggregation configuration	10
SubVLAN Address Range Checking.....	10
ESRP and VLAN aggregation	10
ESRP improvements	11
Tracking Another VLAN	11
ESRP tracking configuration	12
Tracking Routes	12
ESRP route tracking configuration	12
ESRP Interoperability	12
Layer 2 License now allows Layer 3 routing	12
RADIUS Client.....	12
RADIUS client configuration	13
RADIUS RFC 2138 attributes	13
RADIUS Server configuration example (Merit).....	13
Logging and Monitoring Configuration Changes	14
Intra-VLAN Broadcast Suppression	14
Renaming a VLAN.....	15
OSPF Passive Interface	15
vMANs - VPN services for Metropolitan Area Providers	15
Switch Access Security.....	16
Avoiding ‘smurf’ attacks	17
Vista Web Access	17
SNMP Access	17
Telnet access	17
Username and password database.....	17
Supported Limits	17
Clarifications, Known Behaviors and Problems	19
System related – all systems	19
Show Diagnostics	19
Setting Autonegotiation off on a Gigabit Port	19
Flow Control.....	19
Use of redundant phy ports when disabling Smart Redundancy.....	20
System Logging	20
Enabled IdleTimeouts and console connections	20
System related – Black Diamond	20
v4.X Diagnostics for BlackDiamond	20
Diagnostics runs on a module after clearing or unconfiguring a slot.....	20
Using 110v power on a BlackDiamond	20
Enabled IdleTimeouts and multiple BlackDiamond console connections	20
Modem port on MSMs.....	21
Hot Removal of an I/O module with traffic	21

Removal/Insertion of an I/O module	21
GVRP	21
Xmodem Downloads.....	21
Xmodem download through BootRom on the BlackDiamond	21
Xmodem download through CLI on BlackDiamond.....	21
Command Line Interface (CLI)	22
New option for disabling paging	22
Don't use the encrypted option when creating an account	22
Cosmetic PING errors	22
Cosmetic Configuration Download Warnings.....	22
"Interrupt messages lost" message	22
Console may appear locked after telnetting.....	22
Scrolling output using 'show config'	22
Serial and Telnet Configuration.....	23
Switching and VLANs	23
Default routes or static routes	23
ARP aging timer granularity.....	23
Modifying the protocol "IP"	23
Configuring a protocol filter with 'fff'	23
GVRP/GARP and Load-Sharing	23
GVRP statistics.....	23
VLAN Aggregation.....	23
Moving a sub-VLAN client.....	23
No Static ARP entries	24
VLAN aggregation and ESRP	24
VLAN Aggregation and QoS	24
Spanning Tree	24
STP not supported with ESRP	24
Last port, last slot on BlackDiamond.....	24
Mirroring.....	24
Mirroring combined with load sharing	24
Mirroring IP Multicast traffic	24
MAC-based Mirroring.....	24
Mirroring bandwidth.....	25
Port Mirroring and 802.1Q tags.....	25
Mirroring and flooding	25
Basic IP Routing	25
UDP forwarding vs. DHCP forwarding.....	25
It Takes Two VLANs to Route.....	25
IGMP& IGMP Snooping with IP unicast and multicast routing	25
A note on IP Route Sharing.....	25
RIP Routing.....	26
RIP V2 Authentication	26
Routing with OSPF.....	26
Set the RouterID	26
VLANs automatically joined to backbone area	26
Supported Maximums.....	26

Multinetting	26
Multinetting and IP Multicast routing.....	26
Multinetting and client default gateways.....	26
Port Tagging in Multinetting.....	26
ESRP	27
Mixing ESRP and non-ESRP VLANs.....	27
Mixing Clients and Routers on an ESRP-enabled VLAN.....	27
Maximum of Forty eight ESRP-protected VLANs.....	27
Ensure that EDP is enabled.....	27
Error message when enabling ESRP.....	27
Virtual Chassis and ESRP.....	27
IPX Routing	27
Load Sharing and IPX routing.....	27
Maximums.....	28
Tuning.....	28
Performance.....	28
Tagged VLANs and IPX.....	28
IP Multicast Routing	28
Cisco Interoperation.....	28
PIM and route sharing.....	28
IGMP Settings.....	28
QoS	28
QoS bandwidth Granularity.....	28
QoS Profile minimum Bandwidth should not exceed 90% totals.....	29
Adding L4 source port for IPQoS.....	29
IPQoS for multicast with multicast routing enabled.....	29
IPQoS for Blackhole Rule with TCP Specified.....	29
WEB Management - VISTA	29
WEB Server Busy.....	29
Closing Internet Explorer 4.0.....	29
Configuring VLANs through VISTA on GVRP ports.....	30
Default QoS Profile does not appear.....	30
Log entry order.....	30
Vista and RADIUS.....	30
SNMP/RMON	30
Enabling RMON.....	30
Eight entries for allowed managers.....	30
16 entries for SNMP trap receivers.....	30
Trap Receivers as broadcast entry.....	31
Control of UDP port used in sending traps.....	31
Bridge MIB Attributes.....	31
No entry for un-stuffed GBIC redundant slot.....	31
Indicated input voltage is always 110 volts for Summit switch.....	31
Redundant port fail-over traps.....	31
Cannot enable trap authentication from agent.....	31
SNMP Timeouts.....	31
Virtual Chassis	31
Maximum number of VLANs.....	31
Virtual Chassis and ESRP.....	31

Issues resolved from v4.1.19b2	32
General	32
General IP	32
BlackDiamond	32
Summit	33
EDP	33
ISQ	33
QoS	33
PIM-DM	33
Load Sharing	33
IP Multicast	33
RIP	33
DVMRP	34
ESRP	34
STP	34
SNMP	34
Vista	34
Issues resolved from v4.1.18b6	34
General	35
IP	35
PIM	35
SNMP	35
Issues resolved from v4.1.17b6	35
General	35
BlackDiamond	35
ESRP	36
OSPF	36
RIP	36
IP	36
VLAN Aggregation	37
QoS	37
Load-Sharing	37
SNMP	37
Issues resolved from v4.1.17b5	37
VLAN Aggregation	37
Issues resolved from v4.1.16b3	37
General	37
ESRP	38
IP	38
PIM	38
Issues resolved from v4.1.15b4	38
General	38
OSPF	38
QoS	39
ESRP	39
IP	39
RIP	39
Issues resolved from v4.1.12b3	39
General	39

BlackDiamond	39
ESRP	40
OSPF	40
VLAN Aggregation	40
QoS	40
SNMP	40
Issues resolved from v4.1.11b2	40
General	41
Routing	41
IP	41
Issues resolved from v4.1.10b3	41
General	41
Routing	42
BlackDiamond	42
Radius	42
ESRP	42
OSPF	42
Issues resolved from v4.1.9b7	42
General	42
Routing	43
BlackDiamond	43
Vista	43
Issues resolved from v4.1.9b2	43
General	43
Virtual Chassis	43
Routing	43
BlackDiamond	43
ESRP	44
Issues resolved from v4.0.16	44
BlackDiamond	44
Routing	44
General	45
Vista	45
Virtual Chassis	45
SNMP	45



Release Notes for ExtremeWare v4.1.19b2

These release notes contain information on features and issues specific to ExtremeWare v4.1 not covered in the *ExtremeWare Software User Guide v4.0*. Numbers that appear in parenthesis are used for internal reference and can be ignored.

New Features in ExtremeWare v4.1

A listing of new features in ExtremeWare v4.1 is below. Further documentation on these features is contained later within this document.

- VLAN Aggregation – A feature useful to various Service Providers needing to conserve on IP address space.
- ESRP Tracking – Ability to “track” availability of another VLAN or specified route as part of the fail-over criteria. Also, the number of VLANs enabled with ESRP has been raised to 48.
- 1024 VLANs – Ability to scale to 1024 VLANs on the BlackDiamond.
- BlackDiamond FDB entries – In this release, 254 static FDB entries are supported in the BlackDiamond.
- Free Layer 3 switching capabilities – For Summit24 and Summit48 switches that were licensed with Layer 2 capabilities only, ExtremeWare v4.1 allows the use of “BasicL3” which allows wire-speed IP routing between IP subnets that are directly attached, defined through static routes or by using the RIP protocol.
- RADIUS Client – A security feature allowing centrally authorized connections via Telnet and console port to the switch. Used in conjunction with a RADIUS server.
- Configuration Logging – Ability to log configuration changes performed through any CLI session including what user performed the change and at what time.
- Intra-VLAN Broadcast Suppression – Ability to restrict the forwarding of broadcast packets from a port within a VLAN.
- Renaming a VLAN – Ability to modify the name of a VLAN once the VLAN has already been created.
- SubVLAN Address Range Checking – Ability to configure an accepted address range in a subVLAN to prohibit IP addresses outside of that range to be forwarded by the router.
- OSPF passive interface support – Ability to configure passive interfaces for OSPF. Passive interfaces do not send or receive OSPF packets.
- vMANs – A Layer-2 VLAN tunneling capability allowing providers to configure point-to-multipoint “pipes” for customers, independent of the customer’s own VLAN tagging methodology, while preserving integrity and security among separate customers.

Released Software

This release note covers the following software and hardware:

Extreme Switch	ExtremeWare Filename/Version	BootRom filename/version
BlackDiamond	bd4119b2.Z / v4.1.19b2	Bdboot_3_7.bin / v3.7 (no change)
Summit	S4119b2.Z / v4.1.19b2	Sboot_1_9.bin / v1.9 (changed*)

*Version 1.9 bootrom for Summit switches addresses a limited but potential issue when the switch is connected to an inadequately powered UPS system during a voltage drop or spike. Under these circumstances, when the switch attempts to reboot, it may not correctly detect its memory configuration.

This document contains the following sections:

- Upgrading to v4.1. If you received this release note with your switch, you may skip this section.
- New Features in ExtremeWare v4.1.
- Supported Limits Table.
- Switch Access Security – notes on defaults and options for restricting switch access.
- Clarifications, known behaviors and problems – with sub-sections by feature.
- Issues resolved from v4.1.18b6
- Issues resolved from v4.1.17b6
- Issues resolved from v4.1.17b5
- Issues resolved from v4.1.16b3
- Issues resolved from v4.1.15b4
- Issues resolved from v4.1.12b3
- Issues resolved from v4.1.11b2
- Issues resolved from v4.1.10b3
- Issues resolved from v4.1.9b7
- Issues resolved from v4.1.9b2
- Issues resolved from v4.0.X releases.
- A detailed Table of Contents.

Upgrading from v4.0 to v4.1

If you are currently running a release of ExtremeWare v4.0 on a Summit or BlackDiamond to upgrade, 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 v4.1 use the commands:

```
download image <ipaddress> <v4.1_filename> secondary
use image secondary
reboot
```

Upgrading to v4.1 from Releases *Prior to 4.X*

The sections below address procedures and issues associated with upgrading to ExtremeWare v4.1 and new BootRom on the Summit and BlackDiamond switch for customers running previous versions of ExtremeWare (e.g. v2.X for Summits and v3.X for BlackDiamond). If the switch you have received already contains these release notes, you may skip this entire section. If upgrading, be sure to read this entire section.

Follow these instructions completely in order to maintain the ability to upgrade and downgrade ExtremeWare images and system configurations. While running ExtremeWare v4.1, a saved configuration is not readable by previous versions of ExtremeWare. If you do not read and follow these instructions, you may impact operation of the switch. Upgrading to ExtremeWare v4.1 is supported only from version v2.X for Summit switches and from version v3.0.9 for the BlackDiamond switch. Upgrading directly from version 1.X on a Summit switch may not preserve configuration information and is not supported.

Two separate ExtremeWare v4.1 image files are associated with ExtremeWare v4.1, one for all Summit switches and the other for the BlackDiamond switch. The Summit image filename starts with an "s" and the BlackDiamond image filename starts with a "bd". Two new BootRom files are also provided for the Summit and BlackDiamond switch as part of the ExtremeWare v4.1 upgrade. BootRom filenames start with a "sboot" and "bdboot" for Summit and BlackDiamond switches, respectively.

Summit Switches

Below are instructions specific to upgrading to, and downgrading from, ExtremeWare v4.1 for Summit switches.

Summit Hardware Requirements for Running ExtremeWare v4.X

It is necessary to have 32MB of DRAM to support the additional features in ExtremeWare v4.1. This is not an issue for BlackDiamond, Summit48 and Summit24 switches. All currently shipping Summit switches contain 32MB, though some earlier models have 16MB.

To determine if the memory size of a Summit1, Summit2, Summit3 or Summit4 switch is sufficient, perform a "show mem" command using CLI. ExtremeWare v2.1.7 or later will display a line indicating "total DRAM size" which needs to indicate 32MB. With previous versions of ExtremeWare, if the sum of the "current free" bytes and the "current alloc" bytes are greater than 16,000,000, there is no need for a memory upgrade. Otherwise, please contact Extreme Networks Customer Support for further information on arranging a memory upgrade. Phone: (888) 257-3000 or (408) 579-2826.

Upgrading Summit switches

ExtremeWare v4.1 can read a stored configuration saved by ExtremeWare v2.X, however ExtremeWare v2.X is not capable of reading an ExtremeWare v4.1 configuration. To perform the upgrade, you must first be running a 2.X image. The procedures outlined below will preserve the ability to downgrade should it become necessary:

1. If not already the case, TFTP download a current v2.X image (e.g. 2.1.8) into the secondary image space. Sample command would be: "download image <ipaddress> v2_1_8.Z sec".
2. Ensure that the currently used configuration is stored in both the primary and secondary configuration spaces. Commands are "save primary" and "save secondary".

3. Configure the switch to use the primary image and the primary configuration. Commands are “use image primary” and “use config primary”.
4. Verify that all of the above procedures were completed successfully with the command “show switch”.
5. Upload the configuration of the switch to a TFTP server for safekeeping. Command is “upload config <ipaddress> <filename>”.
6. TFTP download Summit ExtremeWare v4.1 software to the primary image space. An example command is “download image <ipaddress> s41xxx.Z pri”.
7. Reboot the switch. The previous configuration of the switch will be preserved going from v2.X to ExtremeWare v4.1. Verify that the switch is operating as expected. After verification, you may configure features specific to ExtremeWare v4.1. 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 v2.x image is not required.

Downgrading Summit Switches

It is assumed that you have followed the upgrade instructions correctly and that a v2.X configuration has been preserved in the secondary configuration space. A Summit switch booting ExtremeWare v2.X will be unable to read a configuration saved under ExtremeWare v4.1 and will resort to factory defaults.

1. If, as per upgrade instructions, the secondary configuration was saved while using a v2.X image, configure the switch to use the secondary configuration. Command is: “use config secondary”. If there is no stored configuration saved under v2.X, you will need to re-configure or download a v2.x configuration file to the switch when running the v2.X image.
2. Use the v2.X 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 a v2.x-based configuration, 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.
5. Because v2.X is not capable of reading an ExtremeWare v4.1 configuration, as a precaution, overwrite any v4.1 configuration that still exists with the command “save primary” and “save secondary”.

BlackDiamond

Below are instructions specific to upgrading to and downgrading from ExtremeWare v4.1 for the BlackDiamond switch.

Upgrading the BlackDiamond

ExtremeWare v4.1 can read a configuration that was saved while running ExtremeWare v3.0.9, however ExtremeWare v3.0.9 is not capable of reading a configuration that was saved while running ExtremeWare v4.1. The procedures outlined below will preserve the ability to downgrade should it become necessary:

1. If not already the case, TFTP download a current 3.0.9 image (e.g. bd309b3.Z) into the secondary image space. Sample command would be: "download image <ipaddress> bd309b3.Z sec".
2. Ensure that the currently used configuration is stored in both the primary and secondary configuration spaces. Commands are "save primary" and "save secondary".
3. Configure the BlackDiamond switch to use the primary image and the primary configuration. Commands are "use image primary" and "use config primary".
4. Verify that all of the above procedures were completed successfully with the command "show switch".
5. Upload the configuration of the switch to a TFTP server for safekeeping. Command is "upload config <ipaddress> <filename>".
6. TFTP download the BlackDiamond ExtremeWare v4.1 software to the primary image space. An example command is "download image <ipaddress> bd41xxx.Z pri".
7. Reboot the switch. The previous configuration of the switch will be preserved going from v3.0.9 to ExtremeWare v4.1. Verify that the switch is operating as expected. After verification, you may configure features specific to ExtremeWare v4.1. Save the configuration to the primary space and do NOT save to the secondary configuration space until you are certain a downgrade to the v3.0.9 image is not required.

Downgrading the BlackDiamond

It is assumed that you have followed the upgrade instructions and that a v3.0.9 configuration has been preserved in the secondary configuration space. If this is not the case, you will need to contact Extreme Networks Customer Support for assistance prior to downgrading.

A BlackDiamond switch booting ExtremeWare 3.0.9 will be unable to read a configuration saved under ExtremeWare v4.1 and will exhibit an exception error during the boot process. The reported exception error from the console port looks similar to the example below:

```
Address store Exception
Exception Program Counter: 0x800b8600
Status Register: 0x3400ad01
Cause Register: 0x00000014
Access Address: 0x800b8243
Task: 0x874ffd80 "tRootTask"
```

If you experience this error, contact Extreme Networks Customer Support to assist you in resolving this issue (5156).

1. If, as per upgrade instructions, the secondary configuration was saved while using a v3.0.9 image, configure the switch to use the secondary configuration. Command is: "use config secondary". If there is no stored configuration that was saved under v3.0.9, you must issue the command "unconfig switch all" while running ExtremeWare v4.1.
2. Use the v3.0.9 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 had to perform an 'unconfig switch all' above, 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.

5. Because v3.0.9 is not capable of reading an ExtremeWare v4.1 configuration, as a precaution, overwrite any v4.1 configuration that still exists with the command "save primary" and "save secondary".

Upgrading BootRom

This release is also supplied with two new BootRom images for the Summit and BlackDiamond switches. The new BootRom releases are not critical to the upgrade or typical operational aspects of ExtremeWare, but it is recommended that you do upgrade the switches to obtain some new feature capabilities in BootRom should a support issue arise. Be sure to perform the BootRom upgrade *after* upgrading to ExtremeWare v4.1 using the command:

```
download bootrom [ <host_name> | <ip_addr> ]
```



Warning: Be sure to download the correct file for your hardware platform. There is a limited amount of checking that occurs. Downloading the incorrect BootRom can result in the switch being unable to boot. Also note that the bootrom is not updated between MSM32's using the sync command, the bootrom must be downloaded to each MSM via TFTP.

Upgrade Issues

Below are issues found during the upgrade process that may affect your migration to ExtremeWare v4.X.

Inserting Second v4.X-based MSM into v3.0.9 system

In the BlackDiamond, if the switch currently has a single MSM running ExtremeWare v3.0.9 and a 2nd MSM is placed into the switch that has ExtremeWare v4.X, the 2nd MSM intended for standby operation will not boot properly. To work-around this problem, upgrade the first MSM from v3.0.9 to v4.1 prior to inserting the 2nd MSM (3907, 5148).

New RIP defaults

ExtremeWare v4.X has been enhanced to better support route re-distribution with the following changes in the default behavior for RIP:

- rip export static is disabled by default
- rip aggregate is disabled by default.

If you had RIP enabled when using a previous version of ExtremeWare, both these attributes were enabled by default and will be preserved after an upgrade. If you plan on utilizing route re-distribution under v4.X, we recommend carefully reviewing your configuration settings for aggregation and exporting of static routes (3951).

New Features in ExtremeWare V4.1

Below are descriptions of features not contained in the ExtremeWare 4.0 User Guide.

Improved Exception Handling

If the system fails prior to booting up, it will automatically start the console and allow access to the system to view the logs or otherwise debug the failure. You can also configure the system to respond to software failures using the following command:

```
config sys-recovery-level [critical | all | none] [shutdown | reboot]
```

If you specify `critical`, the system will shut down or reboot if a critical task exception occurs. Critical tasks include the `tBGTask`, `tNetTask`, or `tESRPTask`.

If you specify `all`, the system will shut down or reboot if any task exception occurs.

If you specify `none`, no action is taken when a task exception occurs.

You cannot specify different actions for different task exceptions. The default setting is “none.”

The `config sys-recovery-level` command monitors 2 tasks for the “critical” level software exceptions – `tBGTask` and `tNetTask` (16002/12477).

Enable and Disable Slots on Modular Switches

You can now enable or disable a slot. When you disable a slot, the I/O module in the slot is brought down and the slot is powered down. The LEDs on the module go OFF. To disable a slot, use the following command:

```
disable slot <number | all>
```

If you specify `all`, all slots are disabled.

If you do not save the configuration before you reboot the switch, the slot will be enabled upon reboot. If you save the configuration, the slot will remain disabled after a reboot.

If there is no I/O card present in a slot when you disable the slot, the slot is still disabled. If you insert a module in a slot that has been disabled, the module does not come up and stays in the “disabled” state until you enable the slot.

To enable a slot, use the following command:

```
enable slot <number | all>
```

If you specify `all`, all slots are enabled (12108).

Automatic Module Disable on Hardware Failure

You can now configure a module to automatically be disabled upon receipt of a PQM failure message (12475). To configure the switch to automatically disable a module experiencing a hardware failure, use the following command:

```
config card-failure-action [none | disable]
```

Enable and Disable IP ARP Refresh

You can now disable IP ARP refresh (if IP forwarding is also disabled). The purpose of disabling ARP refresh is to reduce ARP traffic in a high node count layer 2 switching only environment (12411/15252).

To disable IP ARP refresh, use the following command:

```
disable iparp refresh
```

To enable IP ARP refresh, use the following command:

```
enable iparp refresh
```

SMON Support

RFC 2613 is now supported (13899).

Security

If the switch receives a packet with a source address that is the same as the switch's address, the switch logs a message and dumps the packet (16003).

VLAN aggregation

VLAN aggregation is an application aimed primarily at Service Providers. Its purpose is to increase the efficiency of IP address space usage by allowing clients within the same IP subnet to utilize separate broadcast domains while using the same default router. A "super-VLAN" is defined with the desired IP address but without any member ports unless running ESRP. The super-VLANs IP address is shared among the "sub-VLANs" as the default router address. Groups of clients are assigned to sub-VLANs that have no IP address but are members of the super-VLAN. Clients may be informally allocated any valid IP addresses within the subnet. Optionally, communication between sub-VLANs may be prevented for security purposes. As a result, sub-VLANs can be quite small but allow for growth without re-defining subnet boundaries.

The typical alternative would be to have a separate VLANs for each client group with differing default router addresses and larger subnet masks. This leads to more IP address space being unused.

Optionally, a "secondary" IP address may be defined for the Super VLAN. This IP address is **only** used for the purposes of responding to ICMP Pings to verify connectivity. It is not used for any other purpose.

VLAN aggregation properties

Utilizing VLAN aggregation is a very specific application and several properties apply to its operation:

- All broadcast and unknown traffic will remain local to the sub-VLAN and does not cross the sub-VLAN boundary. All traffic within the sub-VLAN is switched by the sub-VLAN, allowing traffic separation between sub-VLANs while utilizing the same default router address among the sub-VLANs.
- Hosts are located on the sub-VLAN. Each host can assume any IP address within the address range of the super-VLAN's router interface. Hosts on the sub-VLAN are expected to have the same network mask as the super-VLAN and have their default router set to the super-VLANs' IP address.
- All traffic (IP unicast and IP multicast) between sub-VLANs is routed through the super-VLAN. For example, no ICMP redirects will be generated for traffic between sub-VLANs since the super-VLAN is responsible for sub-VLAN routing. Unicast IP traffic across the sub-VLANs is facilitated by the automatic addition of an ARP entry (similar to a proxy ARP entry) when a sub-VLAN is added to a super-VLAN. This feature may be disabled for security purposes (see "security option" later in this section).
- IP Multicast traffic between sub-VLANs is routed when an IP multicast routing protocol is enabled on the super-VLAN.

VLAN aggregation limitations

The following limitations apply to a VLAN aggregation installation:

- No additional routers may be located in the sub-VLANs. This feature is only applicable for “leaves” of a network.
- A sub-VLAN cannot be a super-VLAN and visa versa.
- Sub-VLANs should not be assigned IP addresses.
- Typically, a super-VLAN would have no ports associated with it except in the case of running ESRP (see “ESRP and VLAN aggregation”).
- If a client is moved from one sub-VLAN to another, you must clear the IP ARP cache at the client or the switch in order to resume communications (4977).
- Multinetting is not supported in conjunction with VLAN Aggregation (5848).

Security Option for inter sub-VLAN communication

To facilitate communications between sub-VLANs, by default, an entry is automatically made into the IP ARP table of the super-VLAN that performs a proxy-arp function. This allows clients on one sub-VLAN to communicate with clients on another sub-VLAN. In certain circumstances, intra sub-VLAN communication may not be desired for security reasons. To prevent this communication, the automatic addition of the IP ARP entries may be disabled on the super-VLAN. This is done using the command:

```
[enable | disable] subvlan-proxy-arp vlan <super-VLAN name>
```

VLAN aggregation CLI syntax

The CLI commands specific to VLAN aggregation are:

Adding/Deleting a sub-VLAN:

```
config vlan <super-VLAN NAME> [add | delete] subvlan <Sub-VLAN NAME>
```

Adding/Deleting a secondary IP Address for responding to ICMP ping requests:

```
config vlan <super-VLAN NAME> [add | delete] secondary <IP Address>/32
```

To control making intra-subvlan proxy arp route entries :

```
[enable | disable] subvlan-proxy-arp vlan <super-VLAN name>
```

Configuring VLAN aggregation

The example below illustrates a VLAN aggregation configuration. One VLAN ‘vsuper’ is created as a super-VLAN and several sub-VLANs ‘vsub1’, ‘vsub2’, ‘vsub3’ are added to it.

Create and assign an IP address to a VLAN designated as the super-VLAN. This VLAN should have no member ports. Be sure to enable IP forwarding and any desired routing protocol on the switch.

```
create vlan vsuper
config vsuper ipaddress 192.201.3.1/24
enable ipforwarding
enable ospf
config ospf add vsuper
```

Now create and add ports to VLANs that are intended to be sub-VLANs (up to the maximum number of VLANs).

```
create vlan vsub1
con vsub1 add port 10-12
create vlan vsub2
config vsub2 add po 13-15
create vlan vsub3
config vsub3 add po 16-18
```

Configure the intended super-VLAN by adding the other VLANs as sub-VLANs using the command:

```
config vsuper add subvlan vsub1
config vsuper add subvlan vsub2
config vsuper add subvlan vsub3
```

Optionally, to prevent intra sub-VLAN communication use the command (default is enabled):

```
disable subvlan-proxy-arp <super-VLAN name>
```

Verifying the VLAN aggregation configuration

The following commands are useful to verify proper configuration.

'show vlan' – This command will indicate a sub-VLANs membership in a super VLAN and a super-VLANs associated sub-VLANs.

'show iparp' – After communication with a client on a sub-VLAN, this command will indicate an ARP entry with sub-VLAN information similar to that shown below:

```
* Summit48:1 # sh ipa vsuper
Destination      Mac                Age Flags  Vlan
192.201.3.100    00:e0:2b:58:d8:00  3        sub VLAN: vsub1 vsuper (4092)
```

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 may result in unexpected behavior of ARP within the superVLAN and associated subVLANs.

ESRP and VLAN aggregation

ESRP may be used to provide clients redundant default router protection when part of a VLAN aggregation configuration. ESRP is enabled on the super-VLAN only (not the sub-VLANs). The procedure is to simply add ports to the super-VLAN that is shared with the sub VLANs. To do so, the

super-VLAN should be configured with an 802.1Q tag and added as tagged with the sub-VLAN ports to avoid a protocol conflict. Lastly, enable ESRP on the super-VLAN as you would normally.

Below is an example of combining ESRP and VLAN aggregation for the super-VLAN 'vsuper' and two sub-VLANs 'v1sub' and 'v2sub' that have ports 1 and 2 as members respectively.

Create the VLANs and setup the super to sub-VLAN relationship

```
create vlan v1sub
create vlan v2sub
create vlan vsuper
config vsuper ipaddress 10.1.2.3/24
enable ipforwarding
enable ospf
config ospf add vsuper
config v1sub add port 1
config v2sub add port 2
config vsuper add subvlan v1sub
config vsuper add subvlan v2sub
```

Now perform the configurations necessary to turn on ESRP for the VLAN 'vsuper'

```
config vsuper tag 1234
config vsuper add port 1,2 tagged
enable esrp vlan vsuper
```

If you wish to verify your configuration the following commands are helpful:

'show vlan' - to verify super and sub-VLAN relationships, IP addresses and port membership
'show esrp' - to verify ESRP is enabled and operational

ESRP improvements

There is an additional feature allowing ESRP to also "track" another VLAN or specified route as criteria for fail-over.

Tracking Another VLAN

In a typical Layer 3 router redundancy configuration in which the ESRP switches are routing to a "cloud" or router backbone, part of the criteria for determining the master/slave fail-over can now be the VLAN that links the switch to the routed backbone.

Other factors being equal, if one or more links to the routed backbone fails for the master, ESRP fails-over to the switch that has the most active ports associated with the routed backbone. If there are no (zero) active ports associated with the tracked VLAN, ESRP forces the switch to remain in slave state since no backbone connectivity is available. When the tracking mechanism is used, the precedence that determines ESRP master vs. slave status is as follows:

1. If the priority is = 255, the switch will always become slave
2. Number of Active ports on the ESRP VLAN
3. Number of active ports in the tracked VLAN if tracking is enabled*
4. Configured ESRP priority (1-254)

5. MAC address

* If the number of active tracked ports goes to zero, the ESRP priority is automatically changed to 255.

ESRP tracking configuration

To add or delete a VLAN to be tracked when using ESRP, enter the CLI command:

```
config vlan <VLAN_NAME> [add | delete] track <Tracked VLAN_NAME>
```

The 'show esrp' command has been modified to reflect status of the tracked VLANs and ports associated with those VLANs.

If ESRP tracking is used, all switches running ESRP must be at ExtremeWare v4.1.x or above (5107).

When using ESRP tracking, the tracked-VLAN must first be deleted from the ESRP VLAN before disabling ESRP (5866).

Tracking Routes

ESRP can now also "track" specified routes in the route table as criteria for fail-over. If any of the configured routes are not reachable within the route table, the switch will automatically relinquish "master" status and remain in "standby" or "slave" status. In a typical Layer 3 router redundancy configuration in which the ESRP switches are routing to a "cloud" or router backbone, part of the criteria for determining the master/slave fail-over can now be route entries to the routed backbone.

ESRP route tracking configuration

To add or delete a VLAN to be tracked when using ESRP, enter the CLI command:

```
config vlan <VLAN_NAME> [add | delete] track-route <ipaddress/mask_length>
```

The 'show esrp' command has been modified to reflect status of the tracked routes.

If ESRP route tracking is used, all switches running ESRP must be at ExtremeWare v4.1.16b3 or above.

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, please confirm that the enabled configuration options are identically configured on the ESRP routers. These include route tracking.

Layer 2 License now allows Layer 3 routing

In ExtremeWare v4.1, any Summit24 or Summit48 that previously contained an "L2" license will now be capable of "Basic L3" capabilities. Basic Layer 3 capabilities allow the user to route IP traffic between directly attached IP subnets, using static routes and using the RIP routing protocol. The "Full L3" license is still required for enabling routing protocols including: OSPF, DVMRP or PIM/DM.

RADIUS Client

Remote Authentication Dial In User Service (RADIUS, RFC 2138) is a mechanism for authenticating and centrally administrating access to network nodes. The ExtremeWare RADIUS client implementation allows authentication for telnet, Vista or console access to the switch.

You may define a primary and secondary RADIUS server for the switch to contact. When a user attempts to login via telnet, http or console, the request will be relayed to the primary RADIUS server and then to the secondary RADIUS server if the primary does not respond. If the radius client is

enabled but access to the radius primary and secondary server fails, the switch will use its local database for authentication.

The privileges assigned to the user (admin vs. non-admin) at the radius server will take precedence over what is configured in the local switch database.

RADIUS client configuration

You may define primary and secondary server communication information, and for each radius server, the radius port number to use when talking to the radius server. The default port value is 1645. The 'client IP address' is defined as the IP address used by the server for communicating back to the switch.

To configure radius, use the following commands:

```
[enable | disable] radius
config radius [primary | secondary] server [<hostname> | <ipaddress>]
{<radius_port_number>} client-ip <ipaddress>
```

To verify your configuration information entered above, use the command:

```
show radius
```

To delete the communication information associated with a radius server, use the command:

```
unconfig radius { server [ primary | secondary ] }
```

RADIUS RFC 2138 attributes

The supported RADIUS RFC 2138 optional attributes supported in this release are:

- User-Name
- User-Password
- Service-Type
- Login-IP-Host

RADIUS Server configuration example (Merit)

Many implementations of RADIUS server use the publicly available "Merit AAA" server application <http://www.merit.edu/aaa>. © 1999 Merit Network, Inc.

Included below are excerpts from relevant portions of an example Merit RADIUS server implementation. The example shows excerpts from the "client" and "user" configuration files. The client configuration file defines the authorized source machine, source name and access level. The user configuration file defines username, password and service type information:

ClientCfg.txt

#Client Name [prefix]	Key	[type]	[version]	
#10.1.2.3:256	test	type = nas	v2	pxf
#pm1	%^\$%#*(&!(*&)+	type=nas		pm1.
#pm2	:-) :-(; ^) :-}!	type nas		pm2.
#merit.edu/homeless	hmoemreilte.ses			
#homeless	testing	type proxy	v1	
#xyz.merit.edu	moretesting	type=Ascend:NAS	v1	
#anyoldthing:1234	whoknows?	type=NAS+RAD_RFC+ACCT_RFC		

```

10.202.1.3      andrew-linux    type=nas
10.203.1.41    eric            type=nas
10.203.1.42    eric            type=nas
10.0.52.14     samf            type=nas

```

UserCfg.txt

```

user Password = ""
  Filter-Id = "unlim"

admin Password = "", Service-Type = Administrative
  Filter-Id = "unlim"

eric Password = "", Service-Type = Administrative
  Filter-Id = "unlim"

albert Password = "password", Service-Type = Administrative
  Filter-Id = "unlim"

samuel Password = "password", Service-Type = Administrative
  Filter-Id = "unlim"

```

Logging and Monitoring Configuration Changes

A new CLI command allows the recording of all configuration changes and their source to the system log that are performed through CLI via Telnet or the local console port. The log entry includes the user account name that performed the change and the source IP address of the client if Telnet was used. The logging applies only to commands that result in a configuration change. The CLI command is:

```
[enable | disable] cli-config-logging
```

Intra-VLAN Broadcast Suppression

A new CLI command allows a network manager to restrict the forwarding of broadcast, multicast, and unknown unicast MAC addresses to specific ports within a VLAN. The capability is aimed at unique situations requiring broadcast and multicast suppression within very large Layer 2 broadcast domains.

Note that this feature limits the ability of end-stations to receive packets destined to broadcast or multicast MAC addresses. Also, end-stations attached to such ports must transmit a packet prior to being able to receive unicast packets destined to the end-station. This may affect user connectivity if an application relies on broadcast or multicast packets and should be used only in environments where the effects have been evaluated (i.e., ARP requests will not be forwarded).

To enable broadcast suppression:

```
configure <VLAN_NAME> add ports <portlist> nobroadcast
```

Ports with nobroadcast configured will have an "*" next to the port number in the 'show vlan' display.

To disable broadcast suppression, the port needs to be removed from the VLAN and re-added to the same VLAN:

```
configure <VLAN_NAME> del ports <portlist>
```

```
configure <VLAN_NAME> add ports <portlist>
```

Renaming a VLAN

A new CLI command allows a network manager to rename an existing VLAN. Do not rename the default VLAN, it is not possible to restore the name 'default' and this VLAN must be preserved for use by the switch.

To modify a VLAN name:

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

OSPF Passive Interface

A new CLI command allows users to configure an OSPF interface as passive. Passive interfaces do not send or receive OSPF packets.

To configure an OSPF interface as a passive interface:

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

To reconfigure an OSPF interface as a normal interface:

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

To display passive interface configuration:

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

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 may 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.

To enable VMAN support:

```
enable jumbo
```

To disable VMAN support:

```
disable jumbo
```

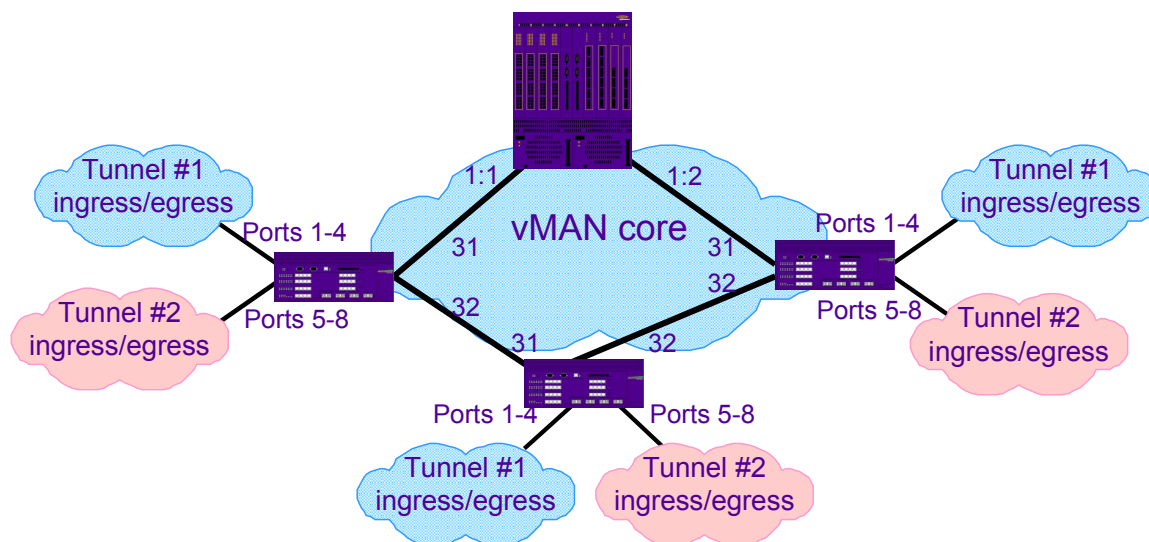
To display VMAN configuration:

```
show switch
```

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) and reboot the switch for the change to take effect (this command automatically updates all ports)
- 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 switch.



For all the switches that are shown, the configuration is:

```

config dot1q ethertype 9100
enable jumbo
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
config vlan tunnell1 tag 50
config vlan tunnell1 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

```

Note that the switch must be rebooted for the new MTU size to take effect. In addition, an uploaded config file will display the “enable jumbo” command as commented (with a “#” preceding the command). This is to protect against the download of this command without a subsequent reboot. The “jumbo” command is saved properly across resets.

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

Switch Access Security

Listed below is information relating to default behavior of the switch and configuration options in environments concerned about security.

Avoiding 'smurf' attacks

In order to avoid a denial-of-service type of attack referred to as a "smurf" or "fraggie" attack that can make many stations busy replying and handling bogus ICMP or UDP traffic, you can disable the IP forwarding of broadcast packets in the switch. Use the command 'disable ipforwarding broadcast <vlan>'. More information on this type of attack may be found on the internet at: <http://users.quadrunner.com/chuegen/smurf.txt> (4862).

Vista Web Access

ExtremeWare Vista Web Access is enabled by default and leverages the same username and password database as the CLI. You may disable Web access entirely or use a RADIUS server to authenticate users.

SNMP Access

SNMP access is enabled by default and utilizes the standard "public" and "private" read and read/write community strings. For security, you may disable SNMP entirely; you may identify specific authorized managers to manage the switch and of course modify the community strings from their defaults.

Telnet access

Telnet access is enabled by default and leverages the username and password database. You may disable Telnet access entirely or use a RADIUS server to authenticate users.

Username and password database

The default username and password database should be modified from its defaults or a RADIUS server should be used. The switch will require that at least one username with administrative capabilities exist.

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 *ExtremeWare Software User Guide*.

METRIC	DESCRIPTION	LIMIT
Access Profiles – Summit	Used by SNMP, Telnet, Vista Web interface, and Routing Access Policies	64
Access Profiles – BlackDiamond switch	Used by SNMP, Telnet, Vista Web interface, and Routing Access Policies	128
Access Profile entries - Summit	Used by SNMP, Telnet, Vista Web interface, and Routing Access Policies	128
Access Profile entries - BlackDiamond	Used by SNMP, Telnet, Vista Web interface, and Routing Access Policies	256
IP QoS rules	Maximum number of IPQoS rules configured on a single switch.	256
IP QoS rules for a single destination	Maximum number of IPQoS rules configured for a single destination	20
Telnet - number of sessions	Maximum number of simultaneous Telnet sessions	8

Supported Limits

SNMP - Trap receivers	Maximum number of SNMP trap receiver stations supported.	16
Syslog servers	Maximum number of simultaneous syslog servers that are supported.	1
Jumbo Frame size	These are not true Jumbo frames. Maximum size supported for Jumbo frames including the CRC and are used primarily for support of vMANs. Switches that use the "i" chipset support true Jumbo frames	1526 Bytes
VLANs – Summit	Includes all VLANs plus sub VLANs, super VLANs, etc.	256
VLANs – BlackDiamond switch	Includes all VLANs plus sub VLANs, super VLANs, etc.	1024
IP Router interfaces	Maximum number of VLANs performing IP routing - excludes SubVLANs	256
Protocol-sensitive VLANs – active protocol filters	The number if simultaneously active protocol filters in the switch.	7
GVRP - Max VLANs	The maximum number of VLANs that can be learned and supported via GVRP.	30
Spanning Tree - Max STPDs	Maximum number of Spanning Tree Domains	64
Spanning Tree – Maximum number of ports	Maximum number of ports that may participate in a single Spanning Tree Domain.	255
IP Static Routes	Maximum number of permanent IP routes.	128
IP route sharing entries	Maximum number of IP routes used in route sharing calculations. This includes static routes and OSPF ECMP.	4
IP Static ARP entries	Maximum number of permanent IP static ARP entries supported.	64
Static IP ARP Proxy entries	Maximum number of permanent IP ARP proxy entries.	64
Static MAC FDB entries – Summit series	Maximum number of permanent MAC entries configured into the FDB.	64
Static MAC FDB entries – BlackDiamond switch	Maximum number of permanent MAC entries configured into the FDB.	254
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.	4
ESRP VLAN-track entries	Maximum number of VLANs that can be tracked by ESRP.	1
ESRP – number of instances	Maximum number of ESRP supported VLANs for a single switch.	48
FDB – Maximum number of entries	Maximum number of MAC addresses.	32,000
Mirroring – Mirrored ports	Maximum number of ports that may be mirrored to the mirror port.	8
Mirroring – number of VLANs	Maximum number of VLANs that may be mirrored to the mirror port.	8
RIP-learned routes	Recommended maximum number of RIP routes supported without aggregation. This is the maximum number of routes supported in the routing table and should be considered in conjunction with other route entry types.	2,000
OSPF areas	As an ABR, how many OSPF areas are supported within the same switch	8

OSPF routes	Recommended maximum number of routes contained in an OSPF LSDB for a "real" network. This is the maximum number of routes supported in the routing table and should be considered in conjunction with other route entry types.	2,000
OSPF LSAs	Recommended maximum number of LSAs contained in the OSPF LSDB for a "real" network.	4,000
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	256
OSPF virtual links	Maximum number of OSPF virtual links supported.	32
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

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 17 of the User Guide.

Show Diagnostics

The Show Diagnostics command will display results from the last power-cycle reboot or "cold boot" but not a "warm boot" done through a management interface (3371).

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 may be checked with the 'show port config' command under the column for flow control (2815).

Use of redundant phy ports when disabling Smart Redundancy

For Summits equipped with supporting redundant gigabit ports (e.g. Summit24, Summit48, Summit2 and Summit3), if the SmartRedundancy feature is disabled and both primary and redundant ports are dual-homed to active gigabit ports, it is not possible to predict which port will become active, the first port completing initialization wins. Enabling SmartRedundancy will allow predictable port failover and fail-back behavior (5079).

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).

System related – Black Diamond

v4.X Diagnostics for BlackDiamond

ExtremeWare v4.0 features additional boot diagnostics for the BlackDiamond switch. If a problem is found with a port during the boot process, the port will be disabled and will not be operational. In addition, an error message is written to the system log, the master MSM will display an amber "sys" LED and the 'show port <portlist>' and 'show slot <slot>' will indicate the failed ports. Also, you will see an asterisk "*" before the CLI prompt indicating a configuration change has occurred due to the disabling of the port. The failed port may be re-enabled manually by issuing an 'enable port <portlist>' command, but the I/O module should be returned to Extreme Networks (4511).

Diagnostics runs on a module after clearing or unconfiguring a slot

For the BlackDiamond switch, if a 'clear slot <slot #>', or 'unconfig slot <slot #>' command is used with an I/O module in the slot, the module will run diagnostics. This is normal behavior. The module should complete diagnostics with default settings restored (3906).

Using 110v power on a BlackDiamond

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 may 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 “modem” on the MSM blade for the BlackDiamond does not allow any connectivity to the device at this time. The upper 9-pin console ports of both the primary and secondary MSM may be used as console or modem connections (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 may 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).

GVRP

GVRP does not advertise VLAN information after a system reboot on the BlackDiamond and must be manually enabled after each reboot to correctly advertise GVRP information (7856).

Xmodem Downloads

Though not performed under normal circumstances, there are two ways to perform a 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	Xmodem download through BootRom	Xmodem download through CLI
Summit	No issues	Not Operational – Do not use (3662)
BlackDiamond	Remove 2 nd MSM first (see below)	MSM in slot “A” must be master (see below)

Xmodem download through BootRom on the BlackDiamond

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 BlackDiamond

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 may then perform the Xmodem download through the upper serial port on the MSM in slot A. (4710, 4848).

Command Line Interface (CLI)

New option for disabling paging

This release contains a new option for disabling the normal paging that occurs when a large amount of screen output occurs or between the output pages of major categories of information (e.g. 'show vlan'). Disabling CLI paging can be useful when using scripts to capture output of the switch. By default CLI paging is enabled. The setting can be changed and stored as part of the switch's normal configuration. The syntax for the command is (4942):

```
{enable | disable} clipaging
```

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 that passwords are not indicated in clear text within the configuration file (4229, 4719).

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 when the ping is terminated by hitting a carriage return (5132).

When a ping is redirected, the statistics for the last packet received are 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 may 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 may be safely ignored (3427). The error message will resemble:

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

Console may appear 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).

Scrolling output using 'show config'

The 'show config' command will scroll the output on CLI without providing any page breaks. The intended purpose of show config is for the screen capture of configuration information (3790).

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).

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.

ARP aging timer granularity

The ARP aging granularity of the switch is set in minutes, using the command `'config iparp timeout <min.>'`. The resolution of the timer is plus or minus 60 seconds (4811).

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 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).

GVRP statistics

GVRP transmit statistics are inaccurate. The number shown will always be greater than the actual (4793, 4794, 4795).

VLAN Aggregation

Moving a sub-VLAN client

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

No Static ARP entries

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

VLAN aggregation and ESRP

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

VLAN Aggregation and QoS

Because sub-VLANs contain no IP address, the ISQ (Intra-Subnet QoS) feature cannot be enabled when using VLAN aggregation (5120).

On the BlackDiamond IPQoS does not work correctly when used in conjunction with VLAN aggregation across different I/O blades. Summits work correctly (5114, 5126).

Spanning Tree

STP not supported with ESRP

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

Last port, last slot on BlackDiamond

If an F32T or F32F I/O module is placed in slot 8, it is not possible to run Spanning Tree from the last port on the module. This limitation is independent of the I/O modules in slots 1-7. The limitation is due in part to an 802.1d Spanning Tree limitation of 255 ports per bridging device (3146).

Mirroring

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).

MAC-based Mirroring

MAC-based mirroring does not operate correctly in this release (3498, 4952, 4953).

Mirroring bandwidth

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

Port Mirroring and 802.1Q tags

When performing mirroring of a port, egress traffic from a port configured for tagged traffic will appear on the mirror port as un-tagged while ingress traffic to the port will accurately reflect the presence or lack of an 802.1Q tag (4939).

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).

Basic IP Routing

UDP forwarding vs. DHCP forwarding

For basic DHCP and BOOTP forwarding services which forward all requests to a single set of servers, use the 'bootprelay' service as described in the User Guide on page 10-11. For more complex implementations, use the UDP forwarding feature as described in the User Guide on page 10-12. Do not, however, use both of these features simultaneously (4877).

It Takes Two VLANs to Route

For proper behavior as a router, you must have at least two VLANs with associated IP addresses. Though it is possible to enable IP forwarding and a routing protocol with only one VLAN defined, the switch will not act as a router (2134).

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 may be checked using the 'show ipconfig' command (5112).

A note on IP Route Sharing

IP route sharing allows multiple equal cost routes to be used concurrently by the switch. In OSPF, this capability is referred to as "equal cost multi-path" (ECMP) routing. To use IP route sharing, simply enable it ('enable iproute sharing') and configure static routes or a routing protocol as you would normally. A maximum of five OSPF ECMP routes can be used for a given destination. RIP and DVMRP will only advertise a single route to a destination so usefulness of these protocols with route sharing is limited to being combined with OSPF and/or static routes.

Route sharing is typically useful only in instances where you are constrained for bandwidth, which is typically not the case with Extreme switches. Using route sharing also makes router trouble-shooting more difficult because of the difficulty in predicting the path that traffic will travel over. The route sharing algorithm for UDP and TCP traffic is based on the equivalent of a "session" between two IP end stations. For every new entry to the IP forwarding database, the switch will round robin among

the ECMP routes. ICMP traffic may follow a packet by packet round robin depending on the circumstances.

RIP Routing

RIP V2 Authentication

The authentication feature of RIPv2 is not supported.

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).

VLANs automatically joined to backbone area

When a VLAN is configured to run OSPF, unless otherwise configured, it will be automatically joined to the backbone area (0.0.0.0). This may be undesirable and cause a lack of connectivity depending upon system design. Be sure to configure the VLAN with the desired OSPF area with the command `config ospf <vlan name> area <area id>`. If this is the first instance of the OSPF area being used, create the area first with the command `create ospf area <area id>` (4000).

Supported Maximums

To ensure a stable system design, we suggest:

- Not having a single OSPF area exceeding 40 routers.
- An approximate maximum of 2000 routes of combined inter, intra and external routes.

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).

Port Tagging in Multinetting

Ports in Multinetted VLANS must be untagged (13389).

ESRP

Mixing ESRP and non-ESRP VLANs

VLANs that have ESRP enabled should not be combined with VLANs that do not have ESRP enabled on the same physical gigabit Ethernet port. For 10/100 Ethernet ports, please observe the requirements in the ExtremeWare User Guide starting on page 9-5 (5125).

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).

Maximum of Forty eight ESRP-protected VLANs

The maximum number of VLANs for which ESRP should be enabled on a Summit or BlackDiamond is 48. Previously, the stated limitation was 24.

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).

If you have downloaded a configuration to the BlackDiamond switch that was generated while running ExtremeWare v3.0.9, the EDP configuration will be the *opposite* of the intended configuration (4027).

Error message when enabling ESRP

When configuring ESRP, ExtremeWare will issue a warning if the VLAN being enabled for ESRP contains the same ports as a VLAN that does not have ESRP enabled. It will also issue a warning if the VLAN being enabled for ESRP shares the same block of 10/100 ports with another VLAN. This issue is explained in the ExtremeWare v4.0 User guide on p. 9-5. An example of the warning is:

```
WARNING: MAC address conflict between VLAN junk2 and VLAN Default.
```

The warning may be safely ignored only if you have correctly followed the configuration rules outlined in the User Guide and these release notes (5268).

Virtual Chassis and ESRP

Running ESRP on a switch connected to a Virtual Chassis is not supported (5105).

IPX Routing

Load Sharing and IPX routing

Load Sharing should not be used in conjunction with IPX routing. Specifically, a VLAN that is used for IPX routing should not have a port that is a member of a load-sharing group (4876).

Maximums

In this release, it is recommended to use IPX routing in environments having less than 2000 IPX RIP routes and less than 2000 IPX SAP routes. A maximum of 64 IPX RIP and 64 IPX SAP static entries may be configured (5261)

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 CLI commands: (4859)

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

Performance

IPX routing does not leverage the line-rate IP routing capabilities of the switch. As indicated in the User Guide, though dependent on many factors, in most circumstances performance is much less than line rate. BlackDiamond IPX routing performance is better than that of Summit. Layer 2 switching of IPX packets still occurs at full wire-speed.

Tagged VLANs and IPX

IPX routing is not supported on tagged VLANs (7061).

IP Multicast Routing

Listed below are issues specific to running IP Multicast routing using PIMv2 or DVMRP.

Cisco Interoperation

For proper cisco interoperation, you must run cisco IOS version 11.3 or better which supports PIM 2.0. Cisco customer support has also been recommending the use of PIM in favor of DVMRP whenever possible on Cisco routers (4669).

PIM and route sharing

In this release, PIM is not supported in conjunction with route sharing (4699).

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. This differs from the 4.0 Manual and the switch will return an error if these values are exceeded (7705).

QoS

QoS bandwidth Granularity

Though the CLI allows any number between 0 and 100 to be used as a percentage bandwidth, the granularity of the bandwidth control is done according to the following values: 0, 1, 2, 3, 4, 5, 10, 20,

30, 40, 50, 60, 70, 80 and 100 percent. If a percentage bandwidth other than these is given, the value is rounded up to the closest value in the list above (4713).

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).

Adding L4 source port for IPQoS

When configuring IPQoS that is specific to a particular layer 4 source port, it is necessary to also specify a source IP address and range. To indicate any source IP address, use the “wildcard” syntax of 0.0.0.0/0. For example, the CLI command of an IPQoS grouping for UDP port 80 traffic from any source IP address to a range of IP addresses with the first byte of “192” associated with QoS Profile “qp3” would be (4716):

```
config ipqos add udp 192.0.0.0/8 0.0.0.0/0 L4-srcport 80 qp3
```

IPQoS for multicast with multicast routing enabled

Assume that multicast routing is enabled and there is a configuration in IPQoS for a destination multicast address/range for multicast traffic that is sourced from a specific IP address. Take the following IPQoS example:

```
config ipqos add udp 225.1.2.3/32 192.1.2.3/32 qp3
```

In this situation, the granularity (or mask) of the source address is determined not by the IPQoS configuration but by the multicast route as determined by the routing protocol. In this example, the true mask of the source entry might be 192.1.2.0/24. If another station (say 192.1.2.4) transmits first to the multicast 225.1.2.3, all subsequent traffic from that subnet (e.g. 192.1.2.0/24) to the destination multicast address 225.1.2.3 will use the default QoS profile instead of qp3. As a work-around, enter an IPQoS configuration that matches the granularity of the multicast route (4715, 4580, 4581). From the above example this would change to:

```
config ipqos add udp 225.1.2.3/32 192.1.2.0/24 qp3
```

IPQoS for Blackhole Rule with TCP Specified

If a blackhole rule with TCP is specified, the blackhole rule will not take effect if additional rules are added using the protocol TCP. This rule will be applied only if it is created as an L3 rule as opposed to a TCP rule (6371).

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 may 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. The correct order may be checked by looking at the log entry timestamp (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 an 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).

SNMP/RMON

Enabling RMON

There is a CLI command for enabling and disabling RMON statistics collection. By default RMON is “disabled”, however the switch will still respond to RMON queries and sets for alarms and events. By enabling RMON, the switch will start the processes necessary for acquiring switch statistics (3453). The CLI syntax is:

```
[enable | disable] rmon
```

You may verify RMON settings with the command `'show management'`.

Eight entries for allowed managers

A maximum of eight addresses may be configured in the allowed manager’s list for SNMP. The command summary in the User Guide and Quick Reference Guide incorrectly indicate a maximum of 32. With no entries, any address may be used and normal SNMP security applies based on community string (3621, 3620).

16 entries for SNMP trap receivers

A maximum of sixteen entries may be made for SNMP trap receivers. The User Guide and Quick Reference Guide indicate a maximum of six trap receivers (4913).

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 may 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

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

No entry for un-stuffed GBIC redundant slot

If no GBIC is inserted into the redundant port for those switch models that have them, the switch will have no entry in the ifMauTable for that port. An empty primary GBIC port will be reported as "unknown MAU" (2423).

Indicated input voltage is always 110 volts for Summit switch

For Summit switches, the private MIB attribute 'extremeInputPowerVoltage' always indicates a voltage of 110 volts. It is not possible for the Summit switch to determine the voltage input for management purposes (4362).

Redundant port fail-over traps

When a redundant port fails-over a trap is sent by the switch. However, if the trap receiver's communication path is through the involved port, the trap may not be received because the trap is sent before the fail-over is complete (4692).

Cannot enable trap authentication from agent

A set command on the snmpEnableAuthenTrap object will return success, but will not take effect. This is confirmed by using the CLI command 'show management'. As a work-around, Trap authentication may be configured from CLI or Vista (4934 4929).

SNMP Timeouts

MIBWALKS or lookups of certain MIB tables will timeout when the switch is configured with 1000 VLANs or above. ExtremeWare SNMP performance has been enhanced to speed up EEM access even when 1000 VLANs are configured on the switch (9205).

Virtual Chassis

Maximum number of VLANs

The maximum number of VLANs that is supported for use within a Virtual Chassis stack is 24 (5127).

Virtual Chassis and ESRP

Running ESRP on a switch connected to a Virtual Chassis is not supported (5105).

Issues resolved from v4.1.19b2

The following issues were found in ExtremeWare v4.1.19b2 and resolved in this release (ExtremeWare v4.1.20b3).

General

- 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 (10858).
- 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).
- The temperature displayed was in degrees Centigrade. The display now matches other current ExtremeWare displays (13976).
- The show port <portlist> qosmonitor command ignored commas, displaying all ports between the ports specified. The command now recognizes commas (14571).
- The show port <portlist> qosmonitor command did not display the correct ports if the port numbers were entered in reverse order. The command now recognizes ports listed in reverse order (15259).
- When a Telnet connection to a switch was terminated without logging in at the admin prompt, a logout message would be printed to the syslog even though the user never logged in (13842).
- A software exception would occur when DNS name resolution was used for the timed download function. There were no issues if the system IP address was used (14419).
- The syslog priority would not be saved in the offline configuration file for TFTP upload/download (11209).
- A login name was incorrectly placed in the log after quitting a Telnet shell (12941).

General IP

- 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 (11547).
- 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 (9385).
- When you created a VLAN and assigned an IP address, the ifIndex of that VLAN in the ipAddr table did not match. When you created a second VLAN and gave it an IP address, the ifIndex of the first VLAN appeared correctly while the second did not match. IP addresses now appear correctly in the ifIndex (12136).

BlackDiamond

- When using a BlackDiamond with multiple slots, if the I/O module in the first available slot was hot-swapped followed by a hot-swap of the second module, IPARP entries would not be correctly installed in the IPARP table with a MAC address association for the remaining modules in the chassis (13632).

Summit

- On a Summit48, when a 10/100 port was administratively disabled and a new port became active, the disabled port would become active (13636).
- On a Summit48, the following message would be printed to the log when a primary and redundant Gigabit interface were active and the primary interface was receiving data traffic at wire-speed:

```
Recv buffers not empty on port 49/50 mask ff02ff message
```

Although the redundant link was being held in standby mode by software, it could be recognized as down, which would result in the log message (11067).

EDP

- EDP packets had the wrong Ethernet length value in the MAC header, causing some endstations to report errors upon receiving EDP packets due to the incorrect value in this field (7830).

ISQ

- When using ISQ IP QoS rules, ISQ rules would not be applied when the host device is a multi-protocol client and a non-IP protocol packet would be the first packet sent to the switch and learned by the FDB (11839).

QoS

- An IP FDB flush could cause connectivity issues with bi-directional traffic when Layer 3 IP QoS was applied to multiple IP FDB entries (9608).

PIM-DM

- Under certain conditions, a multicast "join" message would be sent from a switch with a TTL of "0" leading to multicast stream disruption (11493).

Load Sharing

- Multicast traffic would not properly be forwarded across a load share group when the master port in the group would become unavailable (12817).
- ARP responses would not be transmitted across a load share group when the master port in the group would become unavailable (7317).

IP Multicast

- Cisco HSRP packets destined to 224.0.0.2 were not being correctly learned in the IGMP snooping table as an "All groups" entry (12555/12927).

RIP

- Host routes were not being exported correctly from a RIP network (10309).
- When an access profile was configured to deny RIP routes, the following message would be printed to the system log even though the log was configured to display INFO and above (12403):

```
<DEBUG:RIP > Reject route X.X.X.X from Y.Y.Y.Y on rip1. Access violation
```

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 (11081).

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 (11454). 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.

STP

- 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 (11630).

SNMP

- When using a Basic Layer3 license, 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.
- 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 dot1dTPFdbTable has been modified to only return MAC address information as opposed to MAC and IP address information (14448).
- The SNMP trap destination port would be reset to default upon a TFTP upload and subsequent download of the configuration file (13860).
- A software exception would occur when modifying VLAN configuration data when multiple trap receivers were configured on the switch (10931/13670).
- When creating and deleting a VLAN, SNMP would become unresponsive when traversing the ifStackTable. If EEM 2.1 was used to perform a discovery function of a switch in this state, EEM would become unresponsive and would have to be restarted (10817).

Vista

- The Vista statistics page has been modified to display OSPF external routes correctly. Prior to this release, multiple entries for the same network would fill the Vista statistics page (7751).
- A software exception would occur when using Vista to access a switch that was using extended ASCII characters (i.e., European characters) for the system banner (11076).
- After changing the VLAN membership of a port, there were no longer any entries reported for this port in the ifStackTable. The problem remained after reboot (13644).

Issues resolved from v4.1.18b6

The following issues were found in ExtremeWare v4.1.18b6 and resolved in this release (ExtremeWare v4.1.19b2).

General

- Support added for the new OUI on all Extreme switches (9294).

IP

- Adding a VLAN without an IP address and BootP enabled would cause other VLANs in the multinetting group to not forward traffic correctly (9286).
- A software exception would occur if UDP forwarding was enabled and a route to the destination network for that rule was not available (9343).

PIM

- When a router received a PIM prune packet reflected back to a port that initiated the packet, the router would eventually become unresponsive (9290).

SNMP

- On the Summit switches, an SNMP get on the MAU MIB resulted in a response with a redundant physical interface for every port (9257).

Issues resolved from v4.1.17b6

The following issues were found in ExtremeWare v4.1.17b6 and resolved in ExtremeWare v4.1.18b6.

General

- Vista would not display the configured banner if it exceeded 1024 characters. The CLI would display the banner correctly (7964).
- The “show edp” command now displays both slot and port information of neighboring Extreme switches (8835). Note that Summit switches display a slot number of 1.
- SNTP time updates were not working when the SNTP server needed to be contacted by resolving the name via DNS (7418).
- Under certain circumstances, it is possible to run out of “pty’s” (telnet port) when using telnet to access the switch even though all telnet sessions have been closed (8587).
- The “show edp” command incorrectly displayed the EDP VLAN information on port 1 for all EDP neighbors. The correct port number is now displayed (7804).

BlackDiamond

- If Spanning Tree was configured on an MSM32 and the MSM32 was subsequently moved to a new chassis, the Spanning Tree Bridge ID would be incorrectly displayed as the Bridge ID from the original chassis where Spanning Tree parameters were configured (8882).
- The BlackDiamond would intermittently report false power supply failure and recovery messages even though the power supplies were operational (8911).
- The BlackDiamond will now generate a trap on a DC input failure (8048).
- Load sharing on a G6X I/O module was not working on all ports of a load-share group (8304). Load sharing is now operational on all available ports of a G6X with a maximum of 4 ports per load share group.

ESRP

- Adding ports to an ESRP VLAN on a G6X I/O module could result in an error when other ports on the same I/O module also belonged to ESRP protected VLANs (8099).

OSPF

- An ABR with multiple areas, including an NSSA attached area, would intermittently stop generating type-7 default LSAs to other areas (8202).
- An ABR with multiple normal areas, including an area that is connected as an NSSA, would continue to advertise itself as an ASBR even after the connection to the NSSA area was removed (8836).
- An ABR would stop generating type-3 default summary LSAs when changing from an ABR to internal router (8854).
- An ABR configured between the backbone area and a stub area would not withdraw the default route LSA after becoming an internal router to the stub area (8843).
- The secondary default route will not be used if the primary ABR cannot reach the Backbone Area in a multi-ABR environment (8175).
- An NSSA router will not use a default route provided by a neighboring ABR when the ABR exports a static default route (9022).
- The default LSA cost is not updated on an NSSA router with the stub-default-cost is modified on the NSSA ABR (9023).
- Debug messages will not be printed to the log for interfaces that do not have OSPF configured on a router that is running OSPF (8659).
- Type-4 LSAs are no longer advertised to an NSSA router (8189).
- When exporting static routes into OSPF with different mask lengths for the same IP addresses, the removal of a more specific mask for the same address would result in both routes not being exported to the neighboring routers (7834).

RIP

- On a BlackDiamond, configuring RIP with an import filter applied to all VLANs would result in a software exception (9219).
- Debug messages will not be printed to the log for interfaces that do not have RIP configured on a router that is running RIP (7486).
- Disabling RIP on a VLAN that had been added to RIP without an IP address would result in a software exception (8443).

IP

- Receiving an ARP packet with an illegal address of broadcast or multicast for the destination would result in the switch to stop bridging packets (8770).
- If a static ARP entry is created and the associated route is no longer available, the CLI "show iparp" command would result in a software exception (9042).
- A user can now configure an IP route with a gateway pointing back to the same subnet as the configured route. This allows for the redirection of host routes to another router on the same subnet (8654).
- IPFDB entries for address ranges that are learned via a routing protocol are now flushed so as not to result in a stale entry pointing to a previously used route (8500)

VLAN Aggregation

- Unconfiguring the IP address of a super-VLAN could result in a software exception in some cases (8217).

QoS

- Configuring more than 260 QoS rules would cause the serial interface to become unresponsive. The maximum number of QoS rules that can be configured is 256 and an error will be generated if this maximum is exceeded (7684).

Load-Sharing

- A Summit 1 with load-sharing group on ports 1 and 8 with Spanning Tree enabled would report the following message if the master and slave load share ports were continuously flapping (8693/8694).

```
<CRIT:BRDG> PANIC: updateBridgePtag passed zero.
```

SNMP

- Error counters on the "dot3stats" MIB would incorrectly report errors for the next port in the MIB as well as the port that was actually receiving the errors [i.e., when port 1001 was receiving the errors, port 1002 would also be updated with the identical error count] (9004).
- If the switch has many VLAN configured (greater than 256), SNMP will timeout when a query is made for an SNMP table that has VLAN information (9003).
- The log severity level of an SNMP task utilization message has been reduced to "WARN" from "CRIT" (7411).
- Incorrect SNMP authentication was not generating a message to the log or configured syslog server (7593).
- A VLAN Description field has been added as part of the ifTable in the MIB (8675).

Issues resolved from v4.1.17b5

The following issue was found in ExtremeWare v4.1.17b5 and resolved in ExtremeWare v4.1.17b6.

VLAN Aggregation

- When enabling subvlan-proxy-arp in VLAN Aggregation configurations, the switch would incorrectly respond to intra-subVLAN ARP requests between 2 hosts (8084).

Issues resolved from v4.1.16b3

The following issues were found in ExtremeWare v4.1.16b3 and resolved in ExtremeWare v4.1.17b5.

General

- A software exception could occur when using telnet to and from switches with the idletimeouts feature enabled (7768).

ESRP

- When setting the ESRP hello interval to 2 seconds, it was possible for ESRP to intermittently indicate a failover between the Master and Standby router (7618).
- Error checking has been added to not allow an ESRP enabled VLAN to be added as a subVLAN of another VLAN (7859).
- A software exception could occur if a port was deleted while the ESRP task was attempting to send a hello packet out of the recently deleted port (7729).

IP

- An IP ARP entry can no longer be learned in a network that does not match the configured subnet/mask of the VLAN and will not be entered in the IP ARP table (6941). There is a new field in the show iparp command that displays the rejected IP address count, last rejected IP address, and rejected VLAN interface.

PIM

- The IP multicast cache would not properly timeout if a prune and graft message sequence was not followed by additional control packets, resulting in the improper timeout of all subsequent multicast cache entries (7780).

Issues resolved from v4.1.15b4

The following issues were found in ExtremeWare v4.1.15b4 and resolved in ExtremeWare v4.1.16b3.

General

- The ignore-stp command was not properly ignoring the STP protocol and not blocking traffic for VLANs that were associated with a port that had spanning tree enabled (7180).
- The upload configuration command would not properly free memory when executed manually or through the timed upload configuration feature (6872).
- When using VLAN Aggregation, unknown traffic could leak between subVLANs in the same superVLAN when the ARP cache was flushed dynamically or manually (6393).
- The disable learning feature was not working when applied to L3 traffic (7265).
- Additional CLI error checking was added to prevent the system date from being configured to a 3 digit year or the system day to a "0" date (7318).
- When using the enable ignore-stp feature, upon a save and reboot of the switch, the CLI show vlan command could display an incorrect number of active ports in the "ports" field (7396).
- The enable ignore-stp command could cause a software exception when using port 8:32 on a BlackDiamond (7402).
- When using the "show fdb port <portlist>" command, a software exception could occur if the portlist was mis-typed when using the values 800 through 910 (7403).

OSPF

- Appendix E processing added to avoid the possibility of duplicate link IDs, type, and advertising router to be sent in Type 3 LSAs when different ABRs configured with the same address range transmit Type 3 LSAs into non-backbone areas (7199).

- A higher cost path could be selected between ASBR's in the external SPF calculation (7214).

QoS

- A software exception would occur in Egress QoSMode if a TCP rule was created and the ipfdb was cleared for the end nodes that would hit that rule (7256).
- With ISQ when an ipfdb entry would timeout, it would not properly timeout the L2 fdb entry which would affect the ability of the ipqos rule to take effect (7368).

ESRP

- ESRP failover time between ESRP routers has been improved to 4 to 6 seconds between a master and slave transition (7270). Note that this does not include any convergence that may need to occur with higher-layer routing protocols.

IP

- A dynamic IP ARP received by the switch could overwrite a user configured static IP ARP entry (7315).

RIP

- When more than 20 RIP v2 interfaces were configured on the router, no more than those 20 RIP interfaces would be active routing interfaces. This has been addressed and up to 256 interfaces can be configured and enabled (7319).

Issues resolved from v4.1.12b3

The following issues were found in ExtremeWare v4.1.12b3 and resolved in ExtremeWare v4.1.15b4.

General

- When more than one Spanning Tree Domain (STPD) was created on a BlackDiamond, the VLANs in STPD domains other than the s0 domain did not receive broadcast traffic after a switch reboot (6392).
- Multinetting would not populate the ARP table correctly when using tagged VLANs and ports (6868).

BlackDiamond

- A software modification has been made to prevent a condition where the MSM MAC address could be displayed as 00:00:00:00:00:00. (6300).
- A task exception would occur when enabling load sharing on ports 31 and 32 of an F32T I/O module (6376).
- G4X 4-port I/O modules would display ports 5 and 6 in the CLI "show vlan" command (6428).
- Deleting or adding large numbers of I/O ports on a BlackDiamond could sometimes produce the following benign message to the CLI:

```
BlackDiamond:6 # 0x874e5b90 (tExcTask): 41 messages from interrupt level lost.
```

- Upon the ARP table expiration of entries, the G6X might not correctly re-learn ARP entries (6404).

ESRP

- An issue was addressed when using the Cisco Catalyst 1924 attached to an ESRP-enabled BlackDiamond. In this configuration, the BlackDiamond MAC address and ESRP MAC address would cause the Cisco Catalyst to communicate to the incorrect MAC address (6972).
- A MAC address conflict problem was resolved that was associated with inserting new I/O modules and subsequently attempting to add ports to an ESRP VLAN (6979).

OSPF

- If an OSPF router which is the “actual next-hop” receives an external LSA with its own interface address as the forwarding address, it now discards the external LSA instead of adding an OSPF external route to its routing table for the directly connected network (6320).

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 (6857).

QoS

- In egress qosmode, the CLI console could lock if a user configured qosprofile was removed prior to deleting any associated IPQoS traffic-grouping rules that referenced that qosprofile (7020).
- IP QoS rules between the G6X I/O module and other I/O modules would not be applied correctly on all ports of the G6X (6403).
- If a “wildcard” (0.0.0.0/0) blackhole rule was configured on a BlackDiamond, more specific rules would not be applied correctly for inter-module traffic (7090).

SNMP

- CLI port naming now populates the ifalias attribute (5832).
- The ifName attribute now displays the module/port of an interface using a numeric/numeric convention (5832).
- SNMP access to a switch could be terminated if a default route and SNMP trap receiver station was configured, the configuration saved, and the system rebooted (6628).

Issues resolved from v4.1.11b2

The following issues were found in ExtremeWare v4.1.11b2 and resolved in ExtremeWare v4.1.12b3.

General

- The following debug message will no longer be logged for a VLAN with BootP disabled (6038).
Sep 28 14:15:45 anacreon IPRT: bootprelay bad op type 104
- The following debug message will no longer be logged when configuring VLAN aggregation (6234).
10/15/1999 11:38.06 <DEBUG:SYST> bridgeUpdateSuperPtagList vlan 1ffd 1
- When typing “show config” from the CLI, the switch now displays the config with page breaks. CLI paging must be enabled for this feature to take effect (6231).

Routing

- A software processing error would occur if an SNMP trap receiver was configured on a switch and a trap was sent to the trap receiver management station by way of a route that was defined to be “blackhole” through IP QoS (6107).

IP

- Improvements were made to prevent heavy ICMP traffic from affecting normal traffic handling performance (6197/6232).

Issues resolved from v4.1.10b3

The following issues were found in ExtremeWare v4.1.10b3 and resolved in ExtremeWare v4.1.11b2.

General

- STP priority and cost parameters were not maintained when uploading a configuration to a TFTP server (5491).
- Protection added for the “Octopus” denial of service attack. The switch IP stack would fail to communicate if an attack that opened hundreds of connections with the switch was generated (5526).
- Inserting 4 serial GBICs in a Summit48 could cause IP management (i.e., ping) and EEM to be slow or unresponsive (5542).
- In 4.1.10b3, the following message may have been inaccurately printed to the log file (5699).
12/16/2037 02:23.33 <CRIT:IPHS> Possible spoofing attack from 00:e0:2b:6f:3f:00, port 17
- Enabling sub-VLAN proxy in a super-VLAN configuration did not properly generate an ARP response from the super-VLAN interface (5812).
- IP “theft” checking added for super/sub-VLANs – once an ARP entry is associated with a sub-VLAN, that ARP entry cannot be re-learned from another sub-VLAN (5837).
- Previously, DHCP/BootP forwarding was not supported for VLAN aggregation (5863).
- The port label is now included in various ‘show port’ commands and displays the first 8 characters of the port label string (5833).
- A port affected by Electro-Static Discharge (ESD) would cause the switch to boot up very slowly (5873/5874).
- Spanning Tree can now be configured on load shared links (5964).

Routing

- If a default DVMRP route (0.0.0.0/0) is the first entry in the IPMC forwarding cache, the router would not properly add routes to leaf nodes that are not in the same path as the default route (5472).

BlackDiamond

- If a G6X is the first I/O module in a BlackDiamond chassis, clients on an F32 I/O module in the adjacent slot (G6X + 1) would be unable to ping their local router interfaces (5496).
- EDP messages are now displayed in slot:port format on the BlackDiamond (5834).

Radius

- When using Vista in conjunction with RADIUS and SecureID services, authorization requests would fail for Vista configuration pages (5698)

ESRP

- On a BlackDiamond, a router acting as an ESRP slave would not properly block its ESRP interfaces leading to a broadcast storm on the VLANs configured for ESRP (5824).
- ESRP and load sharing did not work properly (5775).
- Disabling ESRP caused tracked VLANs to be removed without notification to the user. There is now a warning message for the user that tracked VLANs must be removed before disabling ESRP (5866).

OSPF

- OSPF areas are now displayed in both dotted and 32-bit notation (5835).

Issues resolved from v4.1.9b7

The following issues were found in ExtremeWare v4.1.9b7 and resolved in ExtremeWare v4.1.10b3.

General

- Disabling a port did not remove the entries associated with that port from the forwarding database (FDB) and addresses could still be learned on that port although they could not receive traffic from any other stations (5117).
- With IRDP enabled on multiple VLANs, a switch could send a unicast advertisement with the wrong source IP address in the ICMP packet (5272).
- IP communication with the switch would be lost if invalid IP OPTIONS were sent as an “attack” to the switch (5352).
- Protection added for “LAND” and “La Tierra” denial of service attack (5403).
- Could not ping between two switches if multinetting was enabled between the two switches though through traffic worked correctly (5368).
- If a user configured the secondary DNS server entry, upon failure of the primary DNS server entry, the switch would not use the configured secondary server (5389).
- If a login banner is defined, the ExtremeWare copyright statement information will no longer appear when attempting a Telnet login (5489).

Routing

- Creating static routing loops between 2 interfaces would cause unknown destinations to duplicate traffic between the ports in the routing loop (5275).
- IPX xping would continue to display response received after links to the next hop destination were removed (5425).
- If a large LSA (packet size >1518Bytes) was flooded to a neighbor router, and no acknowledgement received for that LSA, retransmission of the LSA would cause a memory leak (5450).

BlackDiamond

- The permit function on VLAN-to-VLAN access policies was not working correctly on the BlackDiamond (5338).

Vista

- Vista did not allow a "User" (user privilege) to view port statistics on different slots of a BlackDiamond. It allowed a user to view the first selected slot only (5276).

Issues resolved from v4.1.9b2

The following issues were resolved in ExtremeWare v4.1.9b7 that could be found in 4.1.9b2.

General

- A Summit protected by an RPS will now write to the log in the event of a power failure in addition to giving information through the 'show switch' command and sending SNMP traps (5043).
- If the system were configured to mirror a VLAN, after some time the console port might lose responsiveness and after approximately 2 hours, a log entry indicating "ptagAdd: error – no free LIST tags" would appear (5191).

Virtual Chassis

- If the summits attached to a Virtual Chassis were configured to use QoS profiles other than defaults, it was possible that some EDP messages would get corrupted leading to instability. Typically this was accompanied by error messages in the log indicating "EDP checksum" and/or "PQM corruption" errors (5227).

Routing

- In instances where a routing loop occurs and ICMP traffic was involved, packet duplication could occur. This in turn would lead to many ICMP type 11 messages (TTL expired) being generated by the involved routers (5173).
- Route re-distribution from OSPF to RIP and visa-versa would not properly propagate information on a downed link (5229, 5233).

BlackDiamond

- Enabling a VLAN to ignore STP packets ('enable ignore-stp <vlan>') would previously not be saved across a reboot of the BlackDiamond (5178).

ESRP

- When using VLAN aggregation, adding a sub-VLAN to a super-VLAN after ESRP was previously enabled on the super-VLAN would result in a bridging loop (5196).
- When downloading a proper configuration involving ESRP and VLAN aggregation, error messages would be generated indicating that ESRP could not be enabled on sub-VLANs, even though this was not part of the configuration (5236).
- Previously, a VLAN that had been configured to be “tracked” by ESRP could not be deleted even when ESRP was disabled (5262).

Issues resolved from v4.0.16

The following issues were resolved in ExtremeWare v4.1.9b2 and above that could be found in v4.0.16.

BlackDiamond

- Pressing the reset button on the slave MSM during the initialization of either MSM caused error messages (5001).
- Removing slave MSM caused slot 4 to stop transmitting packets (5042).
- Removing and re-inserting slave MSM in slot B would not return to the full fabric bandwidth of the switch (5022)
- Hot removal of the slave MSM would cause packets sourced from slot #4 to not be transmitted. This issue has been resolved (5042).
- Typing a command that referenced a zero as a port number would cause an exception error in CLI (5092).
- If an F32T or F32 F was hot-swapped with a G4 I/O module, the log might contain several “port mapping failed” messages (5060).
- Using the “*” wildcard option when configuring multiple slots simultaneously would not work properly (5097).
- Under certain configuration circumstances involving an F32F in slot 8 configured with Spanning Tree, performing a MIB walk on the 802.1D Bridge MIB would result in an exception error on the BlackDiamond (5071)

Routing

- The configuration option for the exporting of static routes under OSPF based on versions of ExtremeWare prior to v 4.0 were not retained when upgrading to v4.0 (4943).
- NSSA internal routers did not contain the default route in the route table. (4932).
- Using RIP, the export static option would advertise the incorrect default route when configured to transmit RIP V1 (5138).
- Using UDP forwarding to filter and forward port 67 traffic to a VLAN instead of an IP address was not working properly (4941).
- Disabling and re-enabling the option for forwarding of directed IP multicasts ([enable | disable] ipforwarding broadcasts) could result in an exception error (5155).
- IPX global counters were not reflecting the per vlan counters (4465).
- The ‘show ipxroute’ command would sometimes give erroneous numerical values for RIP route distribution (5059).

General

- Several “debug” messages associated with IGMP, BOOTP messages have been addressed (4966, 4962).
- In several circumstances, the ISQ QoS feature would not operate correctly after a save and reboot (4979).
- The error message “user account instance 65535 out of range” appearing in the log typically due to an incorrect login attempt through the Vista web interface has been addressed (4839, 4924, 4925).
- Using the redundant-phy capability of the Summit 2,3,24 or 48 would fail-over correctly but may be delayed in a fail-back to the primary by several seconds if both links were connected to active gigabit ports (5079)
- The ‘show port info’ command would not properly display duplex settings if the settings changed dynamically. Also redundant port information on the Summit2 and Summit3 would not be shown (5088, 5090).
- Information relating to SmartRedundancy was not present in an uploaded configuration file or in the results of the ‘show config’ command (5078).

Vista

- Enabling IP forwarding through Vista was previously not possible (5000).
- The setting for exporting static routes through OSPF was not working (5091).

Virtual Chassis

- A problem seen with correctly forwarding IP multicast packets with IGMP Snooping enabled in conjunction with the Virtual Chassis has been resolved (5038).
- Problems could be seen with larger numbers of VLANs (16 or greater) under significant traffic load. Symptoms could include flooding of unknown addresses. (5041).

SNMP

- The dot3MauType would return a value of 0.0 for the Summit1 on ports one and eight (5063).
- The dot1dBasePort values for ports when in load-sharing mode would not be reported correctly (5073, 5074).
- A private MIB attribute indicating the configuration to use on reboot (ConfigToUseReboot) was not operating correctly (5068).