AlliedWare Plus Software Maintenance Release Note

Software Version 5.4.2-3.16 for SwitchBlade x8112, SwitchBlade x908, x900, x600, and x610 Series Switches

Introduction

This document lists the issues addressed and enhancements in AlliedWare Plus™ software maintenance version 5.4.2-3.16. This document applies to the following switches:

Models	Series	Release File	Date	GUI file
x600-24Ts, x600-24Ts/XP x600-24Ts-POE x600-24Ts-POE+ x600-48Ts, x600-48Ts/XP	x600	x600-5.4.2-3.16.rel	Jan 2017	gui_542_30.jar
x610-24Ts x610-24Ts-POE+ x610-24Ts/X x610-24Ts/X-POE+ x610-24SPs/X x610-48Ts x610-48Ts-POE+ x610-48Ts/X x610-48Ts/X-POE+	x610	x610-5.4.2-3.16.rel	Jan 2017	gui_542_30.jar
SwitchBlade x908	SwitchBlade x908	SBx908-5.4.2-3.16.rel	Jan 2017	gui_542_30.jar
x900-12XT/S x900-24 x900-24XT x900-24XT-N x900-24XS	x900	x900-5.4.2-3.16.rel	Jan 2017	gui_542_30.jar
SBx8112 with SBx81CFC400	SwitchBlade x8100	SBx81CFC400-5.4.2-3.16.rel	Jan 2017	

Caution: Information in this document is subject to change without notice and does not represent a commitment on the part of Allied Telesis, Inc. While every effort has been made to ensure that the information contained within this document and the features and changes described are accurate, Allied Telesis, Inc. can not accept any type of liability for errors in, or omissions arising from, the use of this information.

Contents

Installing and Enabling this Version

To use this version, your switch must already be running AlliedWare Plus. Contact your distributor or reseller for more information.

To install this version, perform the following:

1. Put the version file onto your TFTP server.
2. If necessary, delete or move files to create space in the switch's Flash memory for the new file.

   Note that you cannot delete the current boot file.

   To list files, use the command:

   ```
   awplus#dir
   ```

   To see the memory usage, use the command:

   ```
   awplus#show file systems
   ```

   To delete files, use the command:

   ```
   awplus#del <filename>
   ```
3. Copy the new release from your TFTP server onto the switch.

   To do this, enter Privileged Exec mode and use the command:

   ```
   awplus#copy tftp flash
   ```

   Follow the onscreen prompts to specify the server and file.
4. Set the switch to boot from the new release.

   Enter Global Configuration mode.

   On the x600 Series switches, use the command:

   ```
   awplus(config)#boot system x600-5.4.2-3.16.rel
   ```

   On the x610 Series switches, use the command:

   ```
   awplus(config)#boot system x610-5.4.2-3.16.rel
   ```

   On the x900 Series switches, use the command:

   ```
   awplus(config)#boot system x900-5.4.2-3.16.rel
   ```

   On the SwitchBlade x908, use the command:

   ```
   awplus(config)#boot system SBx908-5.4.2-3.16.rel
   ```

   On the SwitchBlade x8112 with controller card SBx81CFC400, use the command:

   ```
   awplus(config)#boot system SBx81CFC400-5.4.2-3.16.rel
   ```

   If desired, check the boot settings by entering Privileged Exec mode and using the following command:

   ```
   awplus#show boot
   ```
5. Reboot.

   To do this, enter Privileged Exec mode and use the command:

   ```
   awplus#reload
   ```

Installing the GUI to your switch using an SD card

1. **Download a GUI Java applet.**

   The GUI Java applet file is available in a compressed (zip) file with the AlliedWare Plus Operating System software from the Support area of the Allied Telesis Website: http://www.alliedtelesis.com. Download the Java applet file. This file will have a **.zip** file name extension. You need to extract the Java **.jar** file from the compressed **.zip** file. The version number of the software applet file (**.jar**) gives the earliest version of the software file (**.rel**) that the GUI can operate with.

2. **Copy the GUI Java applet .jar file to an SD card.**

   Insert the SD card in the SD slot on the front of your switch. Connect to the management port, then login to the switch.

   Copy the GUI Java applet to your switch, using the below commands:

   ```
   awplus# copy card:/<filename.jar> flash:/
   ```

   Where **<filename.jar>** is the GUI Java applet file you downloaded in Step 1.

3. **Assign IP addresses.**

   Use the following commands to assign the IP addresses for connecting to the Java applet.

   ```
   awplus# configure terminal
   awplus(config)# interface vlan1
   awplus(config-if)# ip address <address>/<prefix-length>
   ```

   Where **<***address***>** is the IP address that you will subsequently browse to when you connect to the GUI Java applet. For example, to give the switch an IP address of **192.168.2.6**, with a subnet mask of **255.255.255.0**, use the following command:

   ```
   awplus(config-if)# ip address 192.168.2.6/24
   ```

4. **Configure the gateway.**

   Configure your switch with a default gateway, if necessary, using these commands:

   ```
   awplus(config-if)# exit
   awplus(config)# ip route 0.0.0.0/0 <gateway-address>
   ```

   Where *<gateway-address>* is the IP address for your gateway device. Note that you do not need to define a default gateway if you browse to the switch from within its own subnet.

5. **Create a user account.**

   In order to log into the GUI, you must first create a user account. Use these commands to setup a user account:

   ```
   awplus(config)# username <username> privilege 15 password
      <password>
   awplus(config)# exit
   ```

   Note that you can create multiple users to log into the GUI. See the AlliedWare Plus Software Reference for information about the **username** command.

6. **Ensure HTTP service is enabled.**

   The HTTP service needs to be enabled on the switch before it accepts connections from a web browser. The HTTP service is enabled by default. However, if the HTTP service has been disabled, you must enable the HTTP service again. If the HTTP service is disabled, use the following command to enable it:

   ```
   awplus(config)# service http
   ```

   See the AlliedWare Plus Software Reference for information about the **`service http`** command.

7. **Log into the GUI.**

   Start a browser and enter the IP address you configured in Step 3 as the URL. You will be presented with a login screen after the GUI Java applet has started. Log in with the username and password that you defined in the earlier step, named **Create a user account.**

Note: Any configuration changes should be saved to ensure the device settings are retained.

Installing the GUI to your Switch via TFTP Server

1. **Download a GUI Java applet file from the support site.**

   The GUI Java applet file is available in a compressed (**.zip**) file with the AlliedWare Plus Operating System software from the Support area of the Allied Telesis Website: http://www.alliedtelesis.com. Download the Java applet file. This file will have a **.zip** file name extension. You need to extract the Java **.jar** file from the compressed **.zip** file. The version number of the software applet file (.jar) gives the earliest version of the software file (**.rel**) that the GUI can operate with.

2. **Copy the GUI applet.**

   Copy the GUI applet .jar file onto a TFTP server. Ensure this TFTP server is enabled and ready for the switch. Connect to the management port of the switch, then login to the switch. Do not connect to the management port of the TFTP server

3. **Assign the IP addresses.**

   Use the following commands to configure your switch with an appropriate IP address:

   ```
   awplus# configure terminal
   awplus(config)# interface vlan1
   awplus(config-if)# ip address 192.168.2.6/24
   ```

   Where ***<address>*** is the IP address that you will subsequently browse to when you connect to the GUI Java applet. For example, to give the switch an IP address of **192.168.2.6**, and a subnet mask of **255.255.255.0**, use the following command:

   ```
   awplus(config-if)# ip address 192.168.2.6/24
   ```

   Use the following commands to configure your switch with a default gateway:

   ```
   awplus(config-if)# exit
   awplus(config)# ip route 0.0.0.0/0 <gateway-address>
   ```

4. **Configure the default gateway.**

   In necessary, use the following commands to configure the default gateway.

   ```
   awplus(config-if)# exit
   awplus(config)# ip route 0.0.0.0/0 <gateway-address>
   ```

   Where ***<gateway-address>*** is the IP address for your gateway device. Note that you do not need to define a default gateway if you browse to the switch from within its own subnet.

5. **Copy the GUI Java applet to your switch.**

   Use the following commands to copy the GUI Java applet to your switch:

   ```
   awplus# copy tftp://<server-address>/<filename.jar>
      flash:/
   ```

   Where ***<server-address>*** is the IP address for the TFTP server, and where ***<filename.jar>*** is the GUI Java applet file you downloaded in Step 1.

6. **Create a user account.**

   In order to log into the GUI, you must first create a user account. Use the following commands to setup a user account:

   ```
   awplus(config)# username <username> privilege 15 password
      <password>
   awplus(config)# exit
   ```

   Note that you can create multiple users to log into the GUI. See the AlliedWare Plus Software Reference for information about the **username** command.

7. **Log into the GUI.**

   Start a browser then enter the IP address you configured in Step 3 as the URL. You will then be presented with a login screen after the GUI Java applet has started. You can then Log in with the username and password that you defined previously in Step 6.

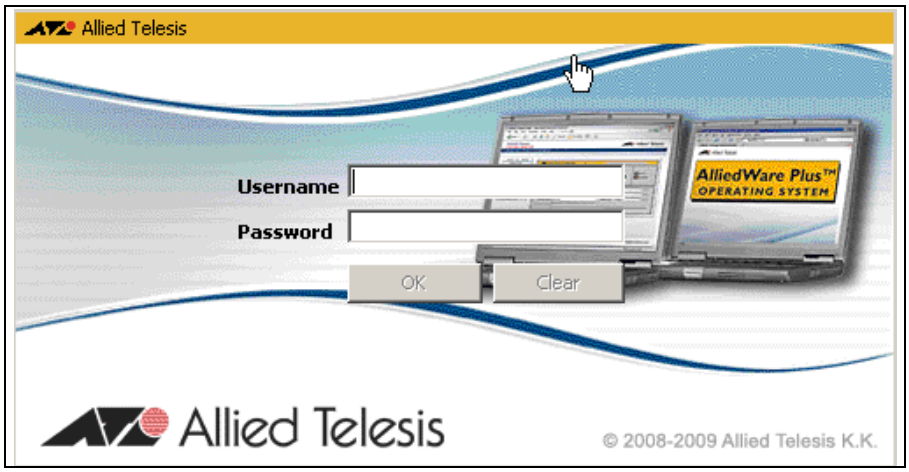

Note: Any configuration changes should be saved to ensure the device settings are retained.

For more information please refer to the 5.4.2 Software Reference available from the Support area of the Allied Telesis Website: http://www.alliedtelesis.com.

Features in 5.4.2-3.16

This release includes the enhancement:

CR	Module	Description
ER-1083	**AMF**	With this software update, the x600 series switch will run AMF Agent Client by default. This will allow the switch to be seen as an AMF Guest node by an AMF node that has the AMF Agent enabled with an "atmf-agentlink" connected to the switch.

Features in 5.4.2-3.14

This release includes the enhancements in the following table:

CR	Module	Description
CR00039430	**SNMP**	The switch now generates a log message when an SNMP request is received, if the request is rejected because it failed authentication (i.e. the event that generates an authentication failure trap). The log message contains the source IP address of the rejected SNMP request, so the user can identify the sender of the unauthenticated request.

CR00039698	**Layer 2 Switching**	A new global command has been added:
		`linkflap action shutdown`
		This command will disable any ports that flap more than 15 times in less than 15 seconds.
		To enable this action, use the command:
		`awplus(config)# linkflap action shutdown`
		To disable port flapping detection (default), use the command:
		`awplus(config)# no linkflap action`
		Note: Because there is already another port flapping detection process that detects very fast flapping, and creates a log message. This other flapping detection is not affected by this new command.

Features in 5.4.2-3.13

This release includes the following enhancements:

CR	Module	Description
CR00038231	**Link Aggregation (Channel Groups)**	x900/x908 Previously, traffic from multiple multicast groups that shared the same VLANs and ports over an aggregator would always flow down one port only. It is now possible to specify a platform command that allows traffic in such situations from different groups to travel down different port members of the LAG, thus increasing maximum possible bandwidth. A **new command** has been introduced: `platform l2-table mode (compact|entry-per-group)`

CR	Module	Description

CR00039055	**Layer 2 switching**	x900/x908
		In certain configurations of over subscription, there could be inefficient throughput. A new buffer-drop mode has been added that allows back pressure to build and work with flow control to improve throughput under certain configurations.
		The **new command** is:
		`platform buffer-drop-mode lossless (new behaviour)`
		`platform buffer-drop-mode tail-drop (default and existing behaviour)`

Features in 5.4.2-3.12

This release includes the following enhancements:

CR	Module	Description
CR00036459	**Pluggable Transceivers**	x908 and x900-24XT switches now support the SP10LR20/1 SFP+ pluggable.
CR00038231	**Link Aggregation (Channel Groups)**	Previously, traffic from multiple multicast groups that shared the same VLANs and ports over an aggregator would always flow down one port only. It is now possible to specify a platform command that allows traffic in such situations from different groups to travel down different port members of the LAG, thus increasing maximum possible bandwidth. A **new command** has been introduced: `platform l2-table mode (compact|entry-per-group)` (**compact** is the default option)

CR00038814	**Hot Swap**	Previously, when performing a regular hotswap, there was a very slight chance that an unrecoverable PCI error would occur, which generated various unnecessary log messages.
		This issue has been resolved.
		To manually **replace** the XEM, first enter the following command:
		`no power enable (stack-member <1-8>) bay <1-8>`
		This command powers off the desired XEM bay.
		Then insert the new XEM.
		Note: Inserting the new XEM into the desired XEM bay will power **on** the XEM by default.
		
		To **reset** a XEM remotely, enter the following commands:
		`no power enable (stack-member <1-8>) bay <1-8>`
		`power enable (stack-member <1-8>) bay <1-8>`
		Note: These commands reset the XEM, but they will not repair a hardware fault.

CR00036176	**Health Check Monitoring**	Hardware Health Monitoring (HHM) is a new feature used to identify faulty or problematic hardware at the customer site. It functions by monitoring particular interrupt messages from the hardware and printing out log messages when the interrupts occur, and the user may configure XEMs to automatically be powered down when the interrupts are detected.
		Currently, only two interrupts are monitored in this fashion, *route table parity* errors and *hyper-g alignment lock* error. Upon finding the messages in the log, the user will be alerted to the fact that their hardware may be faulty.
		To show the current count of the interrupts, use the commands:
		`show system hardware-errors routing-table`
		and
		`show system hardware-errors xbar`
		To configure XEMs to be powered down on detection of each interrupt, use the commands:
		`system hardware-errors routing-table power-off`
		and
		`system hardware-errors xbar power-off`
		For routing table errors, after 50 interrupts have been detected within 5 minutes, the logging will be automatically disabled. To re-enable logging, use the command:
		system hardware-errors routing-table restart-log
		If a XEM has been powered off by HHM, you can use the command:
		`power enable (stack-member <1-8>) bay <1-8>`
		to re-enable the XEM, however a XEM that has been powered-down due to these interrupts is likely to be faulty and should be replaced with new hardware.
		If you are seeing log messages that you are concerned about, please contact your local Allied Telesis distributor for assistance.

Features in 5.4.2-3.11

This release includes these enhancements:

CR	Module	Description
CR00037695	**System**	On SBx8100, the internalports table has been added to the file generated by the **`show tech-support`** command.
CR00038260	**Port Authentication**	AlliedWare Plus web authentication with promiscuous mode now supports sending a login page to IPv6 enabled Windows7.

Features in 5.4.2-3.8

This software version includes these enhancements:

CR	Module	Description
CR00034407	VLAN	There is a new command: (**port-vlan-forwarding-priority**) to set a relative priority among EPSR, Loop Protection and MAC Thrashing protection for setting ports to forwarding a VLAN. Now, when a protocol is set to have the highest priority over a data VLAN on a port, it will not allow other protocols to put that port-vlan into a forwarding state if the highest priority protocol blocked it. See "CR00034407: Prioritising EPSR or Loop Protection" on page 12.
CR00034649	SNMP	There is a new command to configure the delay after startup before sending SNMP traps. See "CR00034649: SNMP startup trap delay" on page 14.
CR00036671	Port Authentication	When a client with a local DHCP server is successfully authenticated by Web authentication, the Authentication Success page is now always displayed in the client's web browser.

CR00036877	DHCP	Previously, when DHCP was configured (**service dhcp-server** command), DHCP request packets received on all Layer 3 interfaces were processed by the DHCP server, sometimes resulting in a high CPU load.
		This behaviour has changed. Now, only DHCP request packets received on interfaces assigned addresses matching DHCP pool settings are processed by the DHCP server. Others are rejected.
		In some cases, this will require configuration changes. For instance, if the DHCP server is configured for DHCP relay packets, the DHCP pool network address does not match its own interface addresses. Therefore, it must add dummy pool data to specify the DHCP server's running interface.
		For instance, if the DHCP server is enabled on VLAN 1 (10.1.1.1/24) and lease address for 192.168.1.0/24 subnet (3rd party network):
		[DHCP server]—10.1.1.0/24—[DHCP Relay] —192.168.1.0/24—[Client]
		use configuration like this:
		`ip dhcp pool relay-pool`
		`network 192.168.1.0 255.255.255.0`
		`range 192.168.1.10 192.168.1.20`
		`!`
		`ip dhcp pool dummy-pool`
		`network 10.1.1.0 255.255.255.0`
		`!`
		`interface vlan1`
		`ip address 10.1.1.1/24`

CR00034407: Prioritising EPSR or Loop Protection

EPSR, Loop Protection and MAC Thrashing protection do not usually need to be configured on a switch, because they perform similar functions—each prevents network loops by blocking VLANs on some ports. However, if more than one of these three features is configured on a switch, you can now use a new command to prioritise either EPSR or Loop Protection when their effects on a port would conflict and override each other. Previously, each protocol

could set a port to forwarding for a VLAN, sometimes overriding the previous setting by another protocol to block the port. This could sometimes lead to unexpected storms.

Now, when a protocol is set to have the highest priority over a data VLAN on a port, it will not allow other protocols to put that port-vlan into a forwarding state if the highest priority protocol blocked it.

The priority mechanism is only used for blocking-to-forwarding transitions; protocols remain independent on the forwarding-to-blocking transitions.

Two new commands configure and display the settings for this feature:

- **port-vlan-forwarding-priority** command on page 13
- **show port-vlan-forwarding-priority** command on page 13

port-vlan-forwarding-priority

Use this command to determine how the switch will prioritise port settings, if more than one of EPSR, Loop Protection and MAC thrashing protection are configured. Use the **no** version of this command to restore the default priority (EPSR).

Syntax

```
port-vlan-forwarding-priority {epsr|loop-protection|none}

no port-vlan-forwarding-priority
```

Parameter	Meaning
epsr	Set EPSR as the highest priority protocol (default).
loop-protection	Set Loop Protection as the highest priority protocol.
none	Set the protocols to have equal priority. This was the previous behaviour before this new command was added, and allows the protocols to override each other to set a port to forwarding a VLAN.

Default By default, the highest priority protocol is EPSR

Mode Global Configuration

Example To prioritise EPSR over Loop Protection or MAC thrashing settings, so that loop protection or MAC thrashing protection cannot set a port to forward a VLAN if EPSR has set it to block, use the command:

```
awplus(config)# port-vlan-forwarding-priority epsr
```

show port-vlan-forwarding-priority

Use this command to display whether of EPSR or Loop Protection is set as the highest priority for determining whether a port forwards a VLAN, as set by the **port-vlan-forwarding-priority** command.

Syntax

```
show port-vlan-forwarding-priority
```

Mode Privileged Exec

Example To display the highest priority protocol, use the command:

```
awplus# show port-vlan-forwarding-priority
```

Example

This example describes the behaviour of an EPSR master in a simple two-node EPSR ring with the following configuration:

- Primary port = port1.0.1; secondary port = port1.0.2
- Control VLAN = vlan10, data VLANs = vlan20, vlan30
- MAC thrashing protection set to protect p1.0.1 and p1.0.2

Initially, the EPSR ring is complete, with port p1.0.2 blocking data VLANs vlan20 and vlan30 and some broadcast traffic flowing through. If the user removes vlan30 from EPSR, a storm is created on vlan30. MAC thrashing protection detects it and blocks vlan30. Then after the storm has stopped, MAC thrashing protection sets it to forwarding again and it keeps oscillating between forwarding and blocking. In the meantime, the user adds back vlan30 to EPSR as a data VLAN and EPSR blocks it on port1.0.2.

If the priority is set to none (**`port-vlan-forwarding-priority none`**), MAC thrashing protection notices that the storm has stopped again and decides to put vlan30 on port1.0.2 into forwarding state. This overrides what EPSR requires for this port-VLAN and creates a storm. This matches the old behavior before this feature was implemented.

If the priority is set to none (**`port-vlan-forwarding-priority epsr`**), MAC thrashing protection notices that the storm has stopped again and attempts to put vlan30 on port1.0.2 into forwarding state. The higher priority protocol (EPSR) is blocking the VLAN on this port, so it stays blocking and no storm occurs.

CR00034649: SNMP startup trap delay

To ensure that SNMP traps (notifications) generated during and shortly after startup can be transmitted, they are buffered and their transmission is delayed to allow the device to startup and bring interfaces up. By default, the delay is 30 seconds; you can now configure a longer startup delay time, using the new **`snmp-server startup-trap-delay`** command.

snmp-server startup-trap-delay

Use this command to set the time in seconds after startup before the switch sends any SNMP traps (notifications). Use the no version to restore the default delay (30s).

Syntax

```
snmp-server startup-trap-delay <30-600>

no snmp-server startup-trap-delay
```

Default By default, the stat-up trap delay is 30 seconds.

Mode Global Configuration

Example To delay the sending of SNMP traps till 60 seconds after startup, use the command:

```
awplus(config)# snmp-server startup-trap-delay 60
```

Features in 5.4.2-3.7

This software version includes this enhancement:

CR	Module	Description

CR00019336	**ARP**	Normally, it is invalid for an ARP request to resolve to a multicast MAC address. By default, ARP replies with multicast MAC addresses are not learnt.
		A new feature had been added that allows control over learning of dynamic ARPs that resolve to a multicast MAC address. For example, in some situations ARP-MAC disparity may need to be enabled to support some solutions, such as network load balancing in multicast mode.
		This behaviour can now be controlled by using the **arp-mac-disparity** command.
		To allow ARP replies quoting multicast MAC addresses to be accepted and learnt, use the command:
		`awplus(config-if)# arp-mac-disparity`
		To reset the switch to the default, that is to not dynamically learn ARPs that quote multicast MAC addresses, use the command:
		`awplus(config-if)# no arp-mac-disparity`
		On x900 Series, SBx908, and SBx8100 switches, if the ARP-MAC disparity feature is enabled, ARP entries that resolve to a multicast MAC address are flooded to the VLAN regardless of the port associated with the ARP entry.
		All x600 and x610 Series switches, if the ARP-MAC disparity feature is enabled, the switch still only sends traffic to a single port as specified by the ARP entry.

Features in 5.4.2-3.6

This software version introduces support for the SwitchBlade x8112 Advanced Layer 3+ Chassis Switch.

Support for the new SwitchBlade x8112 Advanced Layer 3+ Chassis Switch

The SwitchBlade x8112 is a 12-slot Advanced Layer 3+ chassis switch designed to deliver high availability, wire-speed performance, and a high port count. Allied Telesis advanced features make it the ideal solution for the modern enterprise network where resiliency, reliability, and high performance are the key requirements.

The Allied Telesis SwitchBlade x8112 is a high performing scalable solution, providing an extensive range of connectivity options. Dual control cards are partnered with 10 line card slots. Gigabit and 10 Gigabit line card options ensure a system capable of meeting the requirements of today's networks, and the flexibility to expand when required.

Figure 1: SBx8112 chassis with line and controller cards installed

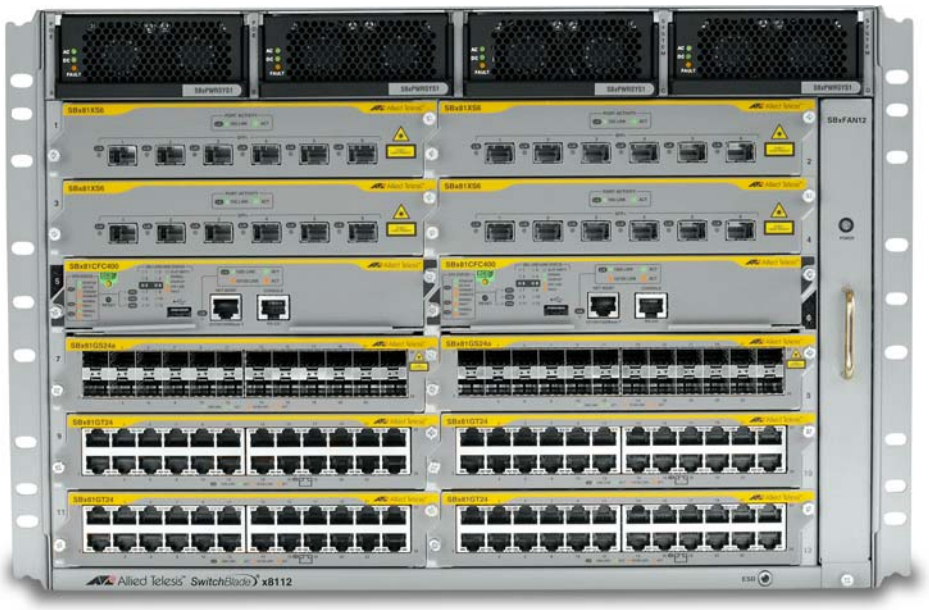

SwitchBlade x8112 is a 7-rack-unit modular chassis comprising:

- Controller Fabric Card (CFC) slots x2
- Line card slots x10
- System PSU bays x2
- Power over Ethernet (PoE) PSU bays x2
- Fan tray

SwitchBlade x8112 key features

- Dual CFCs are partnered with 10 line card slots.
- 80Gbps non-blocking throughput to each line card slot.

- Dual redundant CFCs inter-connect through redundant paths to the line cards over a passive backplane.
- Supports IEEE 802.3at PoE+ (30W); two additional PoE PSUs can be installed to maximize power available to connected devices.
- Designed to reduce power consumption (high efficiency power supplies and low power chip sets) and minimize hazardous waste.
- An ECO-Switch button on the front panel allows conservation of additional power by turning off all diagnostic LED indicators when they are not required.

For information about the AlliedWare Plus features on the SwitchBlade x8112, see the Software Reference.

Controller fabric card

- SBx81CFC400 controller fabric card provides 40Gbps bandwidth per line card slot (80Gbps with 2 control cards)
- 80Gbps x 10 line card slots = 800Gbps per system
- RS-232 and Ethernet management ports for out of band management
- USB connector for external file transfer and storage
- Centralised LEDs provide status of all 12 card slots, CFC, PSU, fan tray

Figure 2: SBx81CFC400

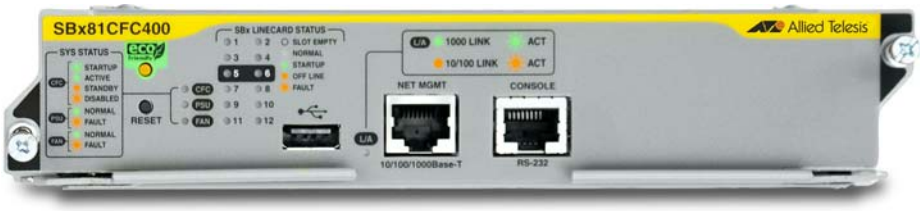

Line cards

- SBx81GT24 (24 x 10/100/1000T Line Card)

- SBx81GP24 (24 x 10/100/1000T PoE+ Line Card)
- SBx81GS24a (24 x 100/1000 SFP Line Card)

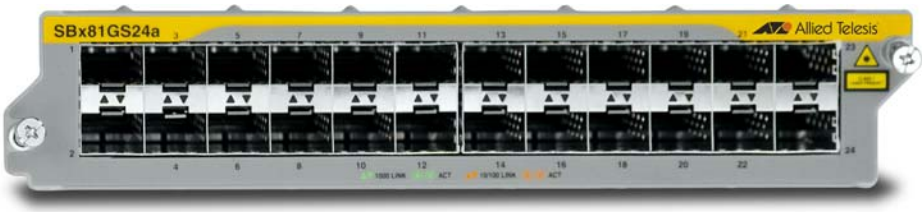

- SBx81XS6 (6 x 10Gbps (SFP+) Line Card)

PSUs and Fan tray

The switch uses high efficiency power supplies. Each power supply can be fed from a separate power source to increase reliability.

- SBxPWRSYS1 (1200W AC System Power Supply)

  The switch operates with one system PSU, and installing a second loadsharing PSU provides ultimate redundancy.

- SBxPWRPOE1 (1200W AC PoE Power Supply)

  Two PoE PSUs can be installed to maximize power available to connected devices.

The switch uses this fan tray:

- SBxFAN12 (Fan tray for 12 slot chassis)

The PSUs and the fan tray are all hot-swappable to maximise uptime during maintenance or reconfiguration.

For more information about the switch hardware, see the Hardware Reference.

Features in 5.4.2-2.5

This software version includes these new features:

- "Save power with the eco-friendly feature" on page 19
- "Static multicast for IPv6" on page 28

and the following enhancements:

CR	Module	Description
CR00035017	L2 Switching	The **snmp-server enable trap** command now has a new **thrash-limit** option, which enables or disables the sending of MAC address thrash limiting traps: `snmp-server enable trap [thrash-limit] [<other-options>]`
CR00035242	IGMP	SBx908, x900 only The theoretical maximum number of supported IGMP groups has been increased from 512 to 2048 respectively. The actual number of groups that can be learnt is hardware dependant.
CR00035687	QoS	Previously for the XEM-2XP, XEM-2XS, and XEM-2XT, the maximum number of hardware ACLs supported was limited to 128 in the default silicon profile, and to 1023 in the extended silicon profile. These limits have been increased—for **platform routingratio IPv4only** the maximum is now 8192; for **platform routingratio IPv4andIPv6** the maximum is now 4096.
CR00035694	L2 Switching	An SBx908 with only XEM-2XP, XEM-2XT, or XEM-2XS installed can now support up to 64K MAC addresses when configured for silicon profile extended. If other XEM types are also installed, this limit remains at 16K, as was previously supported.

Save power with the eco-friendly feature

This feature is supported on: x610 and XEMv2 installed in SBx908.

You can conserve power by enabling the eco-friendly LED (Light Emitting Diode) feature and the eco-friendly LPI (Low Power Idle) feature.

The new commands for configuring and monitoring this feature are:

- **ecofriendly led** command on page 21
- **ecofriendly lpi** command on page 21
- **show ecofriendly** command on page 23

Eco-friendly LED

You can conserve power by enabling the eco-friendly LED feature and the eco-friendly LPI (Low Power Idle) feature.

x600, x610: This feature disables power to the port LEDs, including the stack port status LEDs. Power to the system status, SD and stack management LEDs are not disabled.

x900, SBx908: This feature disables power to the port LEDs on XEMs installed in the switch, the stack ID LED on the XEM-STK, and all port LEDs on the

switch, except the eth0 port. On the switch, power to the fault, SD and PSU LEDs is not disabled, and in addition, on the SwitchBlade x908, power to the stacking status LED is not disabled. On the XEM-1XP and XEM-2XP power to the XFP LED is not disabled.

When the eco-friendly LED feature is enabled, a change of port status will not affect the display of the associated LED. When the eco-friendly feature is disabled and power is returned to port LEDs, the LEDs will correctly show the current state of the ports.

In a stack environment, enabling the eco-friendly LED feature on the stack master will apply the feature to every member of the stack.

The eco-friendly LED feature is disabled by default. To enable the feature, enter the commands:

```
awplus# configure terminal
awplus(config)# ecofriendly led
```

For an example of how to configure a trigger to enable the eco-friendly LED feature, see "Configuration: Turn Off Power to Port LEDs" on page 23.

Eco-friendly LPI

You can also conserve power by enabling the eco-friendly LPI feature with the **ecofriendly lpi** command on page 21.

x600, x610: This feature reduces the power supplied to the ports by the switch whenever the ports are idle and are connected to IEEE 802.3az Energy Efficient Ethernet (EEE) compliant host devices.

x900, SBx908: This feature reduces the power supplied to the ports on XEMs installed in the switch, except the eth0 port, whenever ports are idle and are connected to IEEE802.3az Energy Efficient Ethernet compliant host devices.

All ports configured for LPI must support LPI in hardware and must be configured to autonegotiate by default or by using the **speed** and **duplex** commands as needed.

LPI is a feature of the IEEE 802.3az Energy Efficient Ethernet (EEE) standard. LPI lowers power consumption of switch ports during periods of low link utilization when connected to IEEE 802.3az compliant host devices. If the switch is transmitting no data via the a switch port, the port can enter a sleep state, called Low Power Idle (LPI), to conserve power used by the switch.

The eco-friendly LPI feature is disabled by default. To enable the feature for a switch port, or for a range of switch ports in the example below, enter the commands:

```
awplus# configure terminal
awplus(config)# interface port1.0.2-port1.0.22
awplus(config-if)# ecofriendly lpi
```

To display the current eco-friendly LED and LPI configuration status of the switch, enter the command:

```
awplus# show ecofriendly
```

For an example of how to configure a trigger to enable the eco-friendly LPI feature, see "Configuration: Reduce Power Supplied to Ports" on page 26.

ecofriendly led

x600: Use this command to enable the eco-friendly LED feature which turns off power to the port LEDs, including the stack port status LEDs. Power to the system status, SD and stack management LEDs is not disabled.

x900: Use this command to enable the eco-friendly LED feature which turns off power to the port LEDs on XEMs installed in the switch, the stack ID LED on the XEM-STK and all port LEDs on the switch, except the eth0 port. On the switch, power to the fault, SD and PSU LEDs is not disabled, and in addition, on the SwitchBlade® x908 power to the stacking status LED is not disabled. On the XEM-1XP and XEM-2XP power to the XFP LED is not disabled.

Use the **no** variant of this command to disable the eco-friendly LED feature.

Syntax

```
ecofriendly led

no ecofriendly led
```

Default The eco-friendly LED feature is disabled by default.

Mode Global Configuration

Usage When the eco-friendly LED feature is enabled, a change in port status will not affect the display of the associated LED. When the eco-friendly LED feature is disabled and power is returned to port LEDs, the LEDs will correctly show the current state of the ports.

In a stack environment, enabling the eco-friendly LED feature on the stack master will apply the feature to every member of the stack.

For an example of how to configure a trigger to enable the eco-friendly LED feature, see "Configuration: Turn Off Power to Port LEDs" on page 23.

Examples To enable the eco-friendly LED feature which turns off power to all port LEDs, use the following commands:

```
awplus# configure terminal
awplus(config)# ecofriendly led
```

To disable the eco-friendly LED feature, use the following command:

```
awplus# configure terminal
awplus(config)# no ecofriendly led
```

Related Commands

ecofriendly lpi
show ecofriendly

ecofriendly lpi

Use this command to conserve power by enabling the eco-friendly LPI (Low Power Idle) feature. This feature reduces the power supplied to the ports in the base unit or in XEMs (except the eth0 port) by the switch, whenever the ports

are idle and are connected to IEEE 802.3az Energy Efficient Ethernet compliant host devices.

Use the **no** variant of this command to disable the eco-friendly LPI feature.

Syntax

```
ecofriendly lpi

no ecofriendly lpi
```

Default The eco-friendly LPI feature is disabled by default.

Mode Interface Configuration for a switch port, or Interface Configuration for a range of switch ports.

Usage For an example of how to configure a trigger to enable the eco-friendly LPI feature, see "Configuration: Reduce Power Supplied to Ports" on page 26.

All ports configured for LPI must support LPI in hardware and must be configured to autonegotiate by default or by using the **speed** and **duplex** commands as needed.

Examples To enable the eco-friendly LPI feature on a switch port, `port1.0.2`, use the following commands:

```
awplus# configure terminal
awplus(config)# interface port1.0.2
awplus(config-if)# ecofriendly lpi
```

To enable the eco-friendly LPI feature on a range of switch ports, `port1.0.2-port1.0.20`, use the following commands:

```
awplus# configure terminal
awplus(config)# interface port1.0.2-port1.0.20
awplus(config-if)# ecofriendly lpi
```

To disable the eco-friendly feature on `port1.0.2`, use the following commands:

```
awplus# configure terminal
awplus(config)# interface port1.0.2
awplus(config-if)# no ecofriendly lpi
```

To disable the eco-friendly feature on a range of switch ports, `port1.0.2-port1.0.20`, use the following commands:

```
awplus# configure terminal
awplus(config)# interface port1.0.2-port1.0.20
awplus(config-if)# no ecofriendly lpi
```

Related Commands duplex
ecofriendly led
show ecofriendly

show interface
speed

show ecofriendly

This command displays the switch's eco-friendly configuration status. Both **ecofriendly led** and **ecofriendly lpi** configuration status are shown in the output from this command.

Syntax `show ecofriendly`

Mode Privileged Exec and Global Configuration

Example To display the switch's eco-friendly configuration status, use the following command:

```
awplus# show ecofriendly
```

Output Figure 3: Example output from the **show ecofriendly** command

```
awplus#show ecofriendly
Front panel port LEDs          normal
Energy efficient ethernet
Port          Name             Configured  Status
port1.0.1     Port 1           lpi         lpi
port1.0.2                      lpi         lpi
port1.0.3                      lpi         lpi
port1.0.4                      off         off
port1.0.5                      lpi         off
port1.0.6     Port 6           off         off
port1.0.7                      off         -
port1.0.8                      off         -
port1.0.9                      off         -
port1.0.10                     off         -
```

Table 1: Parameters in the output of the **show ecofriendly** command

Parameter	Description
normal	The eco-friendly LED feature is disabled and port LEDs show the current state of the ports. This is the default setting.
off	The eco-friendly LED feature is enabled and power to the port LEDs is disabled.
Port	Displays the port number as assigned by the switch.
Name	Displays the port name if a name is configured for a port number.
Configured	The eco-friendly LPI feature is configured on the port. Either lpi or off is displayed.
Status	The eco-friendly LPI feature is active on the port. Either lpi or off is displayed. Ports that are not running show -.

Related Commands **ecofriendly led**
ecofriendly lpi

Configuration: Turn Off Power to Port LEDs

The following configuration allows you to conserve power by using the eco-friendly LED feature to turn off power to the port LEDs during non-work hours.

See the **ecofriendly led** command for a detailed command description and command examples.

- Trigger 6 activates at 5.30pm and runs a script called **LEDoff.scp**. This script adds commands to turn off power to all the port LEDs
- Trigger 7 activates at 8.30am and runs the script called **LEDon.scp**. This script removes the configuration specified by **LEDoff.scp**

1. Create the **LEDoff.scp** script

   Create a configuration script with the commands that are executed when the trigger conditions are met. You can either create the configuration script by using the CLI with the **edit** command or create a script on a PC and load it onto your device using the **copy (URL)** command. The configuration script for this example is:

```
!
enable
configure terminal
ecofriendly led
exit
exit
!
```

2. Create the **LEDon.scp** script

   Create a script to remove the configuration specified in the **LEDoff.scp** file. The configuration script for this example is:

```
!
enable
configure terminal
no ecofriendly led
exit
exit
!
```

3. Configure trigger 6

   To create trigger 6, use the commands:

```
awplus# configure terminal
awplus(config)# trigger 6
```

   Set the trigger to activate at 5:30pm, by using the command:

```
awplus(config-trigger)# type time 17:30
```

   Set the trigger to activate on Monday, Tuesday, Wednesday, Thursday and Friday:

```
awplus(config-trigger)# day mon tue wed thu fri
```

   Add the script **LEDoff.scp** to the trigger:

```
awplus(config-trigger)# script 1 LEDoff.scp
```

   Specify a helpful description, such as **Shutdown power to LEDs**. Use the command:

```
awplus(config-trigger)# description Shutdown power to LEDs
```

   Change to Global Configuration mode:

```
awplus(config-trigger)# exit
```

4. Configure trigger 7

   To create trigger 7, use the command:

```
awplus(config)# trigger 7
```

   Set the trigger to activate at 8.30am:

```
awplus(config-trigger)# type time 08:30
```

   Set the trigger to activate on Monday, Tuesday, Wednesday, Thursday and Friday:

```
awplus(config-trigger)# day mon tue wed thu fri
```

   Add the script **LEDon.scp** to the trigger:

```
awplus(config-trigger)# script 1 LEDon.scp
```

   Specify a helpful description, such as **Turn on power to LEDs**. Use the command:

```
awplus(config-trigger)# description Turn on power to LEDs
```

5. Verify the configuration

   To check the configuration of the triggers, use the commands:

```
awplus# show trigger 6
awplus# show trigger 7
```

Configuration: Reduce Power Supplied to Ports

The following configuration allows you to conserve power by using the eco-friendly LPI (Low Power Idle) feature to reduce power supplied to the ports during non-work hours.

See the **ecofriendly lpi** command for a detailed command description and command examples.

- Trigger 6 activates at 5.30pm and runs a script called **LPIon.scp**. This script adds commands to reduce power to all the ports.
- Trigger 7 activates at 8.30am and runs the script called **LPIoff.scp**. This script removes the configuration specified by **LPIon.scp**.

1. Create the **LPIon.scp** script

   Create a configuration script with the commands that are executed when the trigger conditions are met. You can either create the configuration script using the CLI with the **edit** command or create a script on a PC then load it onto your device using the **copy (URL)** command. The configuration script for this example is:

```
!
enable
configure terminal
ecofriendly lpi
exit
exit
!
```

2. Create the **LPIoff.scp** script

   Create a script to remove the configuration specified in the **LPIon.scp** file. The configuration script for this example is:

```
!
enable
configure terminal
no ecofriendly lpi
exit
exit
!
```

3. Configure trigger 6

   To create trigger 6, use the commands:

```
awplus# configure terminal
awplus(config)# trigger 6
```

   Set the trigger to activate at 5:30pm, by using the command:

```
awplus(config-trigger)# type time 17:30
```

   Set the trigger to activate on Monday, Tuesday, Wednesday, Thursday and Friday:

```
awplus(config-trigger)# day mon tue wed thu fri
```

   Add the script **LPIon.scp** to the trigger:

```
awplus(config-trigger)# script 1 LPIon.scp
```

   Specify a helpful description, such as **Turn on LPI**. Use the command:

```
awplus(config-trigger)# description Turn on LPI
```

   Change to Global Configuration mode:

```
awplus(config-trigger)# exit
```

4. Configure trigger 7

   To create trigger 7, use the command:

```
awplus(config)# trigger 7
```

   Set the trigger to activate at 8.30am:

```
awplus(config-trigger)# type time 08:30
```

   Set the trigger to activate on Monday, Tuesday, Wednesday, Thursday and Friday:

```
awplus(config-trigger)# day mon tue wed thu fri
```

   Add the script **LPIoff.scp** to the trigger:

```
awplus(config-trigger)# script 1 LPIoff.scp
```

   Specify a helpful description, such as **Turn off LPI**. Use the command:

```
awplus(config-trigger)# description Turn off LPI
```

5. Verify the configuration

   To check the configuration of the triggers, use the commands:

```
awplus# show trigger 6
awplus# show trigger 7
```

Static multicast for IPv6

This feature is supported on x600 and x610 Series switches.

The switch now supports static multicasting for IPv6. The new commands for configuring and monitoring this are:

clear ipv6 mroute

Use this command to delete entries from the IPv6 multicast routing table.

Syntax

```
clear ipv6 mroute {*|<ipv6-group-addr>
   [<ipv6-source-addr>]}
```

Parameter	Description
*	Deletes all IPv6 multicast routes.
<ipv6-group-address>	Group IPv6 address, in hexadecimal notation in the format `X.X::X.X`.
<ipv6-source-address>	Source IPv6 address, in hexadecimal notation in the format `X.X::X.X`.

Mode Privileged Exec

Usage When this command is used, the Multicast Routing Information Base (MRIB) clears the IPv6 multicast route entries in its IPv6 multicast route table, and removes the entries from the multicast forwarder. The MRIB sends a “clear” message to the multicast protocols. Each multicast protocol has its own “clear” multicast route command.

Example

```
awplus# clear ipv6 mroute 2001::2 ff08::1
```

clear ipv6 mroute statistics

Use this command to delete multicast route statistics entries from the IPv6 multicast routing table.

Syntax

```
clear ipv6 mroute statistics {*|<ipv6-group-address>
   [<ipv6-source-address>]}
```

Parameter	Description
*	All multicast route entries.
<ipv6-group-addr>	Group IPv6 address, in hexadecimal notation in the format `X.X::X.X`.

Parameter	Description
<ipv6-source-addr>	Source IPv6 address, in hexadecimal notation in the format X.X::X.X.

Mode Privileged Exec

Examples

```
awplus# clear ipv6 mroute statistics 2001::2 ff08::1

awplus# clear ipv6 mroute statistics *
```

clear ipv6 multicast route

Use this command to delete all IPv6 static multicast routes from the IPv6 multicast routing table configured with the **ipv6 multicast route** command.

Syntax `clear ipv6 multicast route *`

Parameter	Description
*	Deletes all IPv6 static multicast routes.

Mode Privileged Exec

Usage When this command is used, the MRIB clears the IPv6 multicast route entries in its IPv6 multicast route table, and removes the entries from the multicast forwarder. The MRIB sends a “clear” message to the multicast protocols. Each multicast protocol has its own “clear” multicast route command.

Example

```
awplus# clear ipv6 multicast route *
```

ipv6 multicast route

Use this command to add an IPv6 static multicast route for a specific multicast source and group IPv6 address to the MRIB. This IPv6 multicast route is used to forward IPv6 multicast traffic from a specific source and group ingressing on an upstream VLAN to a single or range of downstream VLANs.

Use the **no** variant of this command to either remove an IPv6 static multicast route set with this command or to remove a specific downstream VLAN interface from an IPv6 static multicast route for a specific IPv6 multicast source and group address.

Syntax

```
ipv6 multicast route <ipv6-source-addr> <ipv6-group-addr>
   <upstream-vlan-id> [<downstream-vlan-id>]

no ipv6 multicast route <ipv6-source-addr> <ipv6-group-
   addr> [<upstream-vlan-id> <downstream-vlan-id>]
```

Parameter	Description
<ipv6-source-addr>	Source IPv6 address, in dotted decimal notation in the format X.X::X.X.

Parameter	Description
<ipv6-group-addr>	Group IP address, in dotted decimal notation in the format `X.X::X.X`.
<upstream-vlan-id>	Upstream VLAN interface on which the multicast packets ingress.
<downstream-vlan-id>	Downstream VLAN interface or range of VLAN interfaces to which the multicast packets are sent.

Default By default, this feature is disabled.

Mode Global Configuration

Usage Only one multicast route entry per IPv6 address and multicast group can be specified. Therefore, if one entry for an IPv6 static multicast route is configured, PIM will not be able to update this multicast route in any way.

If a dynamic multicast route exists, you cannot create a static multicast route with the same source IPv6 address, group IPv6 address, upstream VLAN, and downstream VLANs. An error message is displayed and logged. To add a new static multicast route, either wait for the dynamic multicast route to timeout or clear the dynamic multicast route with the **clear ipv6 multicast route** command.

To update an existing IPv6 static multicast route entry with more or a new set of downstream VLANs, you must firstly remove the existing static multicast route and then add the new static multicast route with all downstream VLANs specified. If you attempt to update an existing static multicast route entry with an additional VLAN or VLANs, an error message is displayed and logged.

To create a blackhole or null route where packets from a specified source and group address coming from an upstream VLAN are dropped rather than forwarded, do not specify the optional `<downstream-vlan-id>` parameter when entering this command.

To remove a specific downstream VLAN from an existing static multicast route entry, specify the VLAN you want to remove with the ***<downstream-vlan-id>*** parameter when entering the **no** variant of this command.

Examples To create an IPv6 static multicast route for the multicast source IPv6 address **2001::1** and group IPv6 address **ff08::1**, specifying the upstream VLAN interface as **vlan10** and the downstream VLAN interface as **vlan20**, use the following commands:

```
awplus# configure terminal
awplus(config)# ipv6 multicast route 2001::1 ff08::1
vlan10 vlan20
```

To create a blackhole route for the IPv6 multicast source IP address `2001::1` and group IP address `ff08::1`, specifying the upstream VLAN interface as **vlan10**, use the following commands:

```
awplus# configure terminal
awplus(config)# ipv6 multicast route 2001::1 ff08::1
vlan10
```

To create an IPv6 static multicast route for the multicast source IPv6 address `2001::1` and group IPv6 address `ff08::1`, specifying the upstream VLAN

interface as **vlan10** and the downstream VLAN range as **vlan20-25**, use the following commands:

```
awplus# configure terminal
awplus(config)# ipv6 multicast route 2001::1 ff08::1
vlan10 vlan20-25
```

To remove the downstream VLAN **23** from the IPv6 static multicast route created with the above command, use the following commands:

```
awplus# configure terminal
awplus(config)# no ipv6 multicast route 2001::1 ff08::1
vlan10 vlan23
```

To delete an IPv6 static multicast route for the multicast source IPv6 address `2001::1` and group IPv6 address `ff08::1`, use the following commands:

```
awplus# configure terminal
awplus(config)# no ipv6 multicast route 2001::1 ff08::1
```

ipv6 multicast route-limit

Use this command to limit the number of multicast routes that can be added to an IPv6 multicast routing table.

Use the **no** variant of this command to return the IPv6 route limit to the default.

Syntax `ipv6 multicast route-limit <limit> [<threshold>]`

`no ipv6 multicast route-limit`

Parameter	Description
<limit>	**<***1-2147483647***>** Number of routes.
<threshold>	**<***1-2147483647***>** Threshold above which to generate a warning message. The mroute warning threshold must not exceed the mroute limit.

Default The default limit and threshold value is 2147483647.

Mode Global Configuration

Usage This command limits the number of multicast IPv6 routes (mroutes) that can be added to a router, and generates an error message when the limit is exceeded. If the threshold parameter is set, a threshold warning message is generated when this threshold is exceeded, and the message continues to occur until the number of mroutes reaches the limit set by the limit argument.

Examples

```
awplus# configure terminal
awplus(config)# ipv6 multicast route-limit 34 24
```

```
awplus# configure terminal
awplus(config)# no ipv6 multicast route-limit
```

ipv6 multicast-routing

Use this command to enable IPv6 multicast routing on the switch.

Use the **no** variant of this command to disable IPv6 multicast routing after enabling it. When it is disabled, the switch does not perform multicast functions.

Syntax

```
ipv6 multicast-routing

no ipv6 multicast-routing
```

Default By default, IPv6 multicast routing is disabled.

Mode Global Configuration

Usage When the **no** variant of this command is used, the Multicast Routing Information Base (MRIB) cleans up Multicast Routing Tables (MRT, and stops relaying multicast forwarder events to multicast protocols.

When multicast routing is enabled, the MRIB starts processing any MRT addition/deletion requests, and any multicast forwarding events.

You must enable multicast routing before issuing other multicast commands.

Examples

```
awplus# configure terminal
awplus(config)# ipv6 multicast-routing

awplus# configure terminal
awplus(config)# no ipv6 multicast-routing
```

Validation Commands `show running-config`

show ipv6 mroute

Use this command to display the contents of the IPv6 multicast routing (mroute) table.

Syntax

```
show ipv6 mroute [<ipv6-group-addr>] [<ipv6-source-addr>]
   [{count|summary}]
```

Parameter	Description
<ipv6-group-addr>	Group IPv6 address, in hexadecimal notation in the format `X.X::X.X`.
<ipv6-source-addr>	Source IPv6 address, in hexadecimal notation in the format `X.X::X.X`.
count	Display the route and packet count from the IPv6 multicast routing (mroute) table.
summary	Display the contents of the IPv6 multicast routing (mroute) table in an abbreviated form.

Mode User Exec and Privileged Exec

Examples

```
awplus# show ipv6 mroute

awplus# show ipv6 mroute count

awplus# show ipv6 mroute summary

awplus# show ipv6 mroute 2001::2 ff08::1 count

awplus# show ipv6 mroute 2001::2 ff08::1

awplus# show ipv6 mroute 2001::2 summary
```

Output The following is a sample output of this command displaying the IPv6 multicast routing table for a single static IPv6 multicast route:

Figure 4: Example output from the **show ipv6 mroute** command

```
awplus#show ipv6 mroute
IPv6 Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder installed
Timers: Uptime/Stat Expiry
Interface State: Interface
(2001::2, ff08::1), uptime 03:18:38
Owner IMI, Flags: F
  Incoming interface: vlan2
  Outgoing interface list:
    vlan3
```

The following is a sample output of this command displaying the IPv6 multicast routing count table for a single static IPv6 Multicast route:

Figure 5: Example output from the **show ipv6 mroute count** command

```
awplus#show ipv6 mroute count

IPv6 Multicast Statistics
Total 1 routes using 152 bytes memory
Route limit/Route threshold: 1024/1024
Total NOCACHE/WRONGmif/WHOLEPKT recv from fwd: 6/0/0
Total NOCACHE/WRONGmif/WHOLEPKT sent to clients: 6/0/0
Immediate/Timed stat updates sent to clients: 0/0
Reg ACK recv/Reg NACK recv/Reg pkt sent: 0/0/0
Next stats poll: 00:01:14

Forwarding Counts: Pkt count/Byte count, Other Counts: Wrong If pkts
Fwd msg counts: WRONGmif/WHOLEPKT recv
Client msg counts: WRONGmif/WHOLEPKT/Imm Stat/Timed Stat sent
Reg pkt counts: Reg ACK recv/Reg NACK recv/Reg pkt sent

(2001::2, ff08::1), Forwarding: 0/0, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/0,  Reg: 0/0/0
```

The following is a sample output of this command displaying the IPv6 multicast routing summary table for a single static IPv6 multicast route:

Figure 6: Example output from the **show ipv6 mroute summary** command

```
awplus#show ipv6 mroute summary
IPv6 Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder
installed
Timers: Uptime/Stat Expiry
Interface State: Interface

(2001::2, ff08::1), 03:20:28/-, IMI, Flags: F
```

show ipv6 mif

Use this command to display the contents of the IPv6 Multicast Routing Information Base (MRIB) MIF table.

Syntax `show ipv6 mif [<interface>]`

Parameter	Description
<interface>	The interface to display information about.

Mode User Exec and Privileged Exec

Example

```
awplus# show ipv6 mif

awplus# show ipv6 mif vlan2
```

Output Figure 7: Example output from the **show ipv6 mif** command

```
awplus#show ipv6 mif
Interface   Mif   Owner            Uptime
            Idx   Module
vlan3       0     MLD/MLD Proxy-Service 03:28:48
vlan2       1     MLD/MLD Proxy-Service 03:28:48
vlan1       2     MLD/MLD Proxy-Service 03:28:48
```

Issues Resolved in 5.4.2-3.16

AlliedWare Plus maintenance version 5.4.2-3.16 includes the following resolved issue:

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR-53560	**EPSR**	Previously, if dynamic VLAN and guest VLAN were configured on the same VLAN, port-authentication might not work correctly. This issue has been resolved.	Y	Y	Y	Y

Issues Resolved in 5.4.2-3.14

AlliedWare Plus maintenance version 5.4.2-3.14 includes the resolved issues in the following tables. The issues addressed in this document include a level number. This number reflects the importance of the issue that has been resolved.

The levels are:

Level 1 This issue will cause significant interruption to network services, and there is no work-around.

Level 2 This issue will cause interruption to network service, however there is a work-around.

Level 3 This issue will seldom appear, and will cause minor inconvenience.

Level 4 This issue represents a cosmetic change and does not affect network operation.

Level 2

CR	Module	Level	Description	x900/x908	x600	x610	SBx8100
CR00040143	EPSR	2	Previously, on an x900 or x908 VCStack, a joining member could incorrectly drop packets. This issue has been resolved.	Y	–	–	–
CR00039854	IGMP Snooping	2	Previously, a system reboot could occur when configuring an IGMP proxy. This issue has been resolved.	Y	Y	Y	Y

Level 3

CR	Module	Level	Description	x900/x908	x600	x610	SBx8100
CR00040106	Policy-based Routing	3	Previously, the removal and re-addition of a policy map to a port could cause the policy rules to be applied to hardware in the wrong order. The correct order of the rules would not be restored until the unit was rebooted. This issue has been resolved.	–	Y	Y	–
CR00040112	Port Authentication	3	Previously, if a supplicant re-authentication happened before the RADIUS server timeout on the deleted supplicant, the authd process would end unexpectedly. This issue has been resolved.	Y	Y	Y	Y
CR00040176	ACL	3	Previously, with an access-list configured and attached to a port, rebooting a single SBx908/x900 stack member could lead to issues when attempting to alter the access-list while it remained attached to the port. This issue has been resolved.	Y	–	–	Y
CR00040190	QoS	3	Previously, on a VCStack, the running configuration could report incorrect values for the **`wrr-queue egress-rate-limit`** command. This issue has been resolved.	Y	Y	Y	Y

Level 3 (cont.)

CR	Module	Level	Description	x900/x908	x600	x610	SBx8100
CR00039006	LPD	3	Previously, packets arriving at the CPU were discarded after loop-protection link-down action and timeout were repeatedly executed. This issue has been resolved.	Y	Y	Y	Y
CR00039349	RADIUS	3	Previously, a small memory leak would occur upon successful web-authentication. This issue has been resolved.	Y	Y	Y	Y
CR00039386	VCStack	3	Previously, the egress-rate-limit settings would not be retained on all stack members after reboot or failover. This issue has been resolved.	–	–	Y	–
CR00039436	Port Authentication	3	Previously, an internal software race condition could lead to ports configured for MAC-AUTH and dynamic VLAN assignment not being assigned to the dynamic VLAN, and remaining in the native VLAN. This issue has been resolved.	–	Y	Y	–
CR00039565	Switching	3	Previously, an SBx8100 could experience an system restart when a large number of MAC movements occurred. This issue has been resolved.	–	–	–	Y
CR00039674	RADIUS	3	Previously, accounting packets sometimes failed to be sent to the RADIUS server after a reload. This issue has been resolved.	Y	Y	Y	Y
CR00039702	Port Authentication	3	Previously, port-authentication sent only a RADIUS accounting START packet, and no STOP packet, when re-authentication was successful. This issue has been resolved, now both START and STOP packets are sent upon successful re-authentication.	Y	Y	Y	Y
CR00039753	Port Authentication	3	Previously, when a web-authentication supplicant's status moved from REAUTHENTICATING to AUTHENTICATED, the authenticated page was not displayed to the user. This issue has been resolved.	Y	Y	Y	Y
CR00040062	IPv6	3	Previously, the interface state of a tunnel interface was not changed when its underlying VLAN interface went up and down. This issue has been resolved.	Y	Y	Y	Y

Level 4

CR	Module	Level	Description	x900/x908	x600	x610	SBx8100
CR00038957	Command shell	4	Previously, the command **`show platform memory`** would generate excessive log messages on a switch if the switch contained any of the following XEMs: XEM-2XS, XEM-2XT, XEM-2XP, XEM-12Tv2, XEM-12Sv2, XEM-24T. This issue has been resolved.	Y	–	–	–

Level 4 (cont.)

CR	Module	Level	Description	x900/x908	x600	x610	SBx8100
CR00039348	**DHCPv4**	**4**	Three issues have been resolved in the output of the command: `show ip dhcp binding` ? With some time zone offset configurations, the output would fail to show dynamic leases that have infinite expiry. ? The output would show dynamic infinite leases as having an expiry of "19 Jan 2038 03:14:06" Such leases are now shown with an expiry of "Infinite". ? Expired leases with an expiry date with a single digit day of month were being shown in the output. Such leases are now not shown.	Y	Y	Y	Y
CR00039407	**802.1x**	**4**	Previously, when a switch was operating as an 802.1x authenticator, and had sent a maximum number of unanswered EAP-Request/Identity packets, and then received an EAPOL-Start, it would not send another EAP-Request/Identity immediately. Instead, it would wait for tx-period (default 30sec), before sending the EAP-Request/Identity. This issue has been resolved.	Y	Y	Y	Y

Issues Resolved in 5.4.2-3.13

AlliedWare Plus maintenance version 5.4.2-3.13 includes the resolved issues in the following tables.

No Level 1 Issues

Level 2 Issues

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00033963	**NTP**	Previously, configuring NTP on a stacked device could cause a warning message to be printed in error. This issue has been resolved.	Y	Y	Y	Y
CR00036978	**IGMP Snooping**	Previously, a static router port would be removed upon receiving an IGMP multicast query message on that port. This issue has been resolved.	Y	Y	Y	Y
CR00038712	**Pluggable Transceivers**	Previously, some SFP+ modules were incorrectly being reported as having locked up. This issue has been resolved	–	–	–	Y
CR00038950	**SMTP**	Previously, a memory leak could occur when a high rate of log email was generated. This issue has been resolved.	Y	Y	Y	Y
CR00039132	**IGMP**	Previously, in certain limited circumstances, the **ip multicast forward-first-packet** command would not work as expected. This issue has been resolved.	Y	–	–	Y
CR00039163	**RADIUS**	Previously, during 802.1X authentication, if the authenticator port received another EAP Response/ Identity, while waiting for a RADIUS response, it would abort the current authentication request, and start again. This cycle could continue indefinitely. This issue has been resolved.	Y	Y	Y	Y
CR00039332	**IGMP**	Previously, a system reboot would occur when static SSM mapped groups were joining and leaving. This issue has been resolved.	Y	Y	Y	Y
CR00039341	**IGMP**	Previously, static IGMP group entries would not propagated correctly to PIM when a switch started up. This issue has been resolved.	Y	Y	Y	Y
CR00039403	**802.1x**	Previously, when a combination of 802.1X EAPOL-Version 2 and AuthFailVLAN were used, the authenticator would not correctly respond to an authentication request from a supplicant that had just recently had a login failure. This issue has been resolved.	Y	Y	Y	Y

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00038816	**Port Security**	Previously: ? Removing a port configured with port security from a VLAN would not delete the port security MAC entries. ? Disabling port-security on a port would not delete port-security MAC entries ? Shutting down the VLAN that a port was a member of sometimes would not remove all the learned MAC entries for the port, if port security was enabled on that port. ? If the maximum number of port-security entries was reduced to a smaller value, the FDB entries for learnt dynamic entries would not be deleted. ? If a port configured with port-security was deleted from multiple VLANS, the MAC entries on that port would not be deleted. These issues have been resolved.	Y	Y	Y	Y
CR00038905	**IGMP**	Previously, IGMP would not send membership information to upper layer protocols when IGMP snooping was turned off on a switch that was operating as a querier. This issue has been resolved.	Y	Y	Y	Y
CR00038951	**802.1x**	Previously, in dot1x authentication, the combination of the **max-auth-req** and **auth-fail** VLAN feature would not work as expected. This issue. has been resolved.	Y	Y	Y	Y
CR00039213	**Web Auth**	Previously, the Web Authentication process could occasionally stop accepting new connections. This issue has been resolved.	Y	Y	Y	Y
CR00039564	**SNMP**	Previously, when a switch was configured to use VLAN 4094, performing an SNMP GET for on object in the VLAN4094 dot1q MIB table would cause an expected reload. This issue has been resolved.	Y	Y	Y	Y

Level 3

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00039030	**Chassis Management**	Previously, on a SBx8100, when a high rate of traffic was received at the CPU via both network ports and management ports, it was possible for a CFC failover to occur. This issue has been resolved.	–	–	–	Y
CR00039052	**ARP**	Previously, SBx8100 switches could not resolve the ARP for a device that had performed a silent station movement. This issue has been resolved.	–	–	–	Y
CR00034770	**DOS Detection**	Previously, the time displayed in "Last LDF Rx" in the output of **show loop-protection counters** was incorrect. This issue has been resolved.	Y	Y	Y	Y
CR00038917	**IGMP**	Previously, IGMP would not start properly if IGMP snooping was disabled This issue has been resolved.	Y	Y	Y	Y
CR00036831	**MLD Snooping**	Previously, when a switch received a DHCPv6 solicit message, MLD incorrectly registered the interface that received the message as a router port. This issue has been resolved.	Y	Y	Y	Y
CR00038837	**DHCP Snooping**	Previously, if DHCP snooping trapped a Jumbo DHCP packet (of 1500 or more bytes in length) the packet would fail a message length check, and be incorrectly dropped. This issue has been resolved.	Y	Y	Y	Y
CR00039568	**IPv6**	Previously, on rare occasions, where a switch was used in an IPv6 multicast network without configuring a VLAN with **IPv6 enable** and there were other switches in the VLAN, it was possible for the switch to re-forward an MLD protocol packet back to an originating switch using the original MAC address. This issue has been resolved.	Y	Y	Y	Y
CR00039641	**IGMP**	Previously, if the downstream side of a IGMP proxy was configured before the upstream side, an unexpected system reboot could occur. This issue has been resolved.	Y	Y	Y	Y
CR00039717	**SNMP**	Previously, an SNMP GET for an object in the dor1Q MIB could result in a memory leak if VLAN 4094 was configured. This issue has been resolved.	Y	Y	Y	Y

Level 4 Issues

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00038469	Log	Previously, the following message was logged if the ARP cache was cleared whilst traffic was passing through a XEM-2XT, XEM-2XS, XEM-12Tv2 or XEM12Sv2: `X908-1 kernel: Unexpected Regular DSA tag received on pkt from PP 40501064 68004000` This issue has been resolved	Y	–	–	–

Issues Resolved in 5.4.2-3.12

AlliedWare Plus maintenance version 5.4.2-3.12 includes the resolved issues in the following tables.

No Level 1 Issues

Level 2

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036453	Pluggable Transceivers	Previously, when inserting / removing SFP+ modules from the x610 switch, it was possible that the port would lock up. This issue has been resolved.	Y	–	–	–
CR00037432	PIM-SM v4	Previously, when PIM was restarted following a healthcheck, it did not restart correctly. This issue has been resolved.	Y	Y	Y	Y
CR00038040	DHCP Snooping	When using DHCP Snooping on ports that belong to a v2 XEM (2XP, 12Tv2, 12Sv2, 24T), the hardware filter used to provide secure access to the DHCP client was not correctly updated with the clients IP address, so all traffic from the client would incorrectly be dropped by the switch. This issue has been resolved. DHCP Snooped clients on v2 XEM ports will now apply correct security.	–	–	–	Y
CR00038390	VCStack	On x908 and x900 stacks, if a late-joining stack member contains either 12Tv2, 12Sv2, 24T, 2XP, 2XS or XEM, there is a small chance that during the joining of the new stack member, broadcast packets received by other stack members on vlan 1 may form a loop within the switching silicon of the XEM on the joining unit, which can result in a flood of duplicate packets being sent out other ports on the stack. This has been resolved - these XEM modules will no longer create broadcast packet storms due to a stack late-join event.	Y	Y	Y	Y
CR00038941	Loop Protection	Previously, EPSR could interfere with loop protection, even though EPSR was not configured. This issue has been resolved.	Y	Y	Y	Y

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00039028	VCStack	Previously, adding a static route that conflicted with the stack management network could cause associated repeating error messages such as Route x.x.x.x/x not added as it interferes with the Stack Management Network flooding the log if the VLAN was restarted. This issue has been resolved.	Y	Y	Y	Y

Level 3

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00037681	OSPFv2	In OSPF v2, the recycling of LSAs could cause memory usage inefficiency as small LSAs were being allocated the space previously used by a large LSA. The effect has been reduced to much smaller proportions, allowing more efficient memory usage by this process. This issue has been resolved.	Y	Y	Y	Y
CR00038465	ARP	When an ARP entry on a switch transitioned from REACHABLE to PROBE and at the same time a gratuitous ARP request for the same address was received on a different port, the ARP entry would get stuck in the STALE state. This has been resolved by probing (sending ARP requests) to the device that send the gratuitous ARP in this scenario.	Y	Y	Y	Y
CR00038466	VCStack	A loose fibre cable connection into XEMV2 could result in unnecessarily high CPU Utilization. This issue has been resolved.	Y	–	–	–

CR00038468	VCStack	Under exceptional circumstances, an unexpected internal software process restart on a VCStack member could occur. The AW+ device would recover and continue operation without any noticeable impact on the network. However, two problems would result:
		1. The VCStack could consume very high amounts of CPU usage. When this occurred, SHOW CPU would report high CPU usage for the 'ospfd' and 'ospf6d' software processes. This problem would only occur if OSPF or OSPFv3 was configured on the stack, or a stack member was subsequently rebooted.
		2. Sub-optimal recovery of OSPF and OSPFv3 routing would occur if the VCStack master subsequently rebooted any time afterward.
		If a VCStack was affected by these problems, you would see the following messages in 'SHOW LOG':
		`OSPF[1644]: Couldn't join CPG group ospfd-ffo, result:14, waited:0ms`
		`OSPF[1644]: Initializing CPG failed, -1`
		`OSPF[1644]: Could not synchronize FFO event 0 with other nodes`
		These issues have been resolved.
			Y	Y	Y	–

CR00038928	VCStack	Previously, after restoring the device configuration and software release by autoboot from an SD card, the VCStack virtual-mac and stack member priority configuration was not fully applied upon forming a stack. This issue has been resolved.	Y	–	–	–

CR00039115	**LACP**	Removal of LACP from an interface via the command **no channel-group** was not accepted if an IGMP static group was also configured to use the same interface. This issue has been resolved.	Y	Y	Y	Y

Level 4

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00038153 **CR00038839**	**VCStack**	Previously, if an 8-unit VCStack with a large configuration file underwent a staggered power-up, an “auditing inconsistency” message was occasionally logged to the console. This issue has been resolved.	–	–	Y	–
CR00038548	**MLD Logging**	Previously, MLD snooping on a large number of VLANs (>100) could produce unnecessary log messages. This issue has been resolved.	Y	–	–	–

Issues Resolved in 5.4.2-3.11

AlliedWare Plus maintenance version 5.4.2-3.11 includes the resolved issues in the following tables.

No Level 1 Issues

Level 2

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00037372	**VCStack**	Previously, if an SBx8112 had a single CFC line card installed and its line card rebooted with some unsaved running configuration, the line card interfaces sometimes remained in the Configuring state (displayed in **`show card`** command output) after reboot. This issue has been resolved.	–	–	–	Y
CR00034305	**L3 Switching**	Previously, when the CPU in an SBx8112 was subjected to very high rates of traffic, a system reboot could occasionally occur. This issue has been resolved.	–	–	–	Y
CR00036273	**Hot Swap**	Previously, repeatedly hotswapping an SFP very quickly could result in a system reboot. This issue has been resolved.	Y	Y	Y	Y

Level 2 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036405	IPv6	Previously, the maximum number of IPv6 routes that the CPU could successfully route was limited to a lower number than specified in the datasheets for: ? SBx908 was using the **extended** silicon profile ? x610 ? SBx8112 This meant that IPv6 routes were not able to be learnt in hardware while the number of entries were still below advertised limits. This issue has been resolved.	Y	Y	Y	Y
CR00036458	VCStack	Previously, on a fully populated or almost fully populated SBx8100 chassis where a card was either inserted or rebooted there was a small chance that a master failover could occur. This issue has been resolved.	–	–	–	Y
CR00037357	L3 Switching	Previously, after a topology change (e.g. spanning tree or EPSR port up/down), IP traffic flows which failed to resolve ARP entries could take a long time to recover. This is now improved.	–	Y	Y	–
CR00037447	IGMP	Previously, on rare occasions when deleting a **static ip igmp ssm-map**, a system reboot could occur. This issue has been resolved.	Y	Y	Y	Y
CR00037480	IGMP	Previously, if an aggregated interface was deleted prior to deleting any IGMP static-groups depending on it, the device would reboot. This issue has been resolved.	Y	Y	Y	Y
CR00037643	IGMP	Previously, IGMP IPv4 Source Specific Multicast Mapped groups were not released in a timely manner on IGMP leave. Groups were left to eventually time out, resulting in multicast traffic unnecessarily being sent during this period. This issue has been resolved.	Y	Y	Y	Y
CR00037835	Provisioning	Previously, in rare cases, a stack member could fail to fully join an existing stack and would remain stuck in the 'init' state. This issue has been resolved.	Y	Y	Y	Y
CR00037836	Pluggable Transceivers	Previously, when a fibre SFP was removed from an SBx81XS6 or SBx81XS16 line card in an SBx8112, the line card continued to report that the link was up. This issue has been resolved.	–	–	–	Y
CR00037901	System	Previously, configuring VLAN stacking on the same port could overwrite the configured MRU command setting. This issue has been resolved.	–	Y	Y	Y
CR00038057	System	Previously, occasionally on CFC failover, the SBx8112 could incorrectly detect that a valid line card or backup CFC card was unsupported and disable the card. When this problem occurred, the affected card(s) would appear as 'Disabled' in the output from the **show card** command. If this situation occurred, to recover from the problem would require either a physical hotswap of the affected card, or a reboot of the Active CFC. This issue has been resolved.	–	–	–	Y
CR00038201	ETH	Previously, when network traffic was being synchronised between Active and Standby CFCs, under some ethernet port traffic conditions a system reboot could occur. This issue has been resolved.	–	–	–	Y

Level 2 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036405	IPv6	Previously, the maximum number of IPv6 routes that the CPU could successfully route was limited to a lower number than specified in the datasheets for: ? SBx908 was using the **extended** silicon profile ? x610 ? SBx8112 This meant that IPv6 routes were not able to be learnt in hardware while the number of entries were still below advertised limits. This issue has been resolved.	Y	Y	Y	Y
CR00036458	VCStack	Previously, on a fully populated or almost fully populated SBx8100 chassis where a card was either inserted or rebooted there was a small chance that a master failover could occur. This issue has been resolved.	–	–	–	Y
CR00037357	L3 Switching	Previously, after a topology change (e.g. spanning tree or EPSR port up/down), IP traffic flows which failed to resolve ARP entries could take a long time to recover. This is now improved.	–	Y	Y	–
CR00037447	IGMP	Previously, on rare occasions when deleting a **static ip igmp ssm-map**, a system reboot could occur. This issue has been resolved.	Y	Y	Y	Y
CR00037480	IGMP	Previously, if an aggregated interface was deleted prior to deleting any IGMP static-groups depending on it, the device would reboot. This issue has been resolved.	Y	Y	Y	Y
CR00037643	IGMP	Previously, IGMP IPv4 Source Specific Multicast Mapped groups were not released in a timely manner on IGMP leave. Groups were left to eventually time out, resulting in multicast traffic unnecessarily being sent during this period. This issue has been resolved.	Y	Y	Y	Y
CR00037835	Provisioning	Previously, in rare cases, a stack member could fail to fully join an existing stack and would remain stuck in the 'init' state. This issue has been resolved.	Y	Y	Y	Y
CR00037836	Pluggable Transceivers	Previously, when a fibre SFP was removed from an SBx81XS6 or SBx81XS16 line card in an SBx8112, the line card continued to report that the link was up. This issue has been resolved.	–	–	–	Y
CR00037901	System	Previously, configuring VLAN stacking on the same port could overwrite the configured MRU command setting. This issue has been resolved.	–	Y	Y	Y
CR00038057	System	Previously, occasionally on CFC failover, the SBx8112 could incorrectly detect that a valid line card or backup CFC card was unsupported and disable the card. When this problem occurred, the affected card(s) would appear as 'Disabled' in the output from the **show card** command. If this situation occurred, to recover from the problem would require either a physical hotswap of the affected card, or a reboot of the Active CFC. This issue has been resolved.	–	–	–	Y
CR00038201	ETH	Previously, when network traffic was being synchronised between Active and Standby CFCs, under some ethernet port traffic conditions a system reboot could occur. This issue has been resolved.	–	–	–	Y

Level 2 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00038348	MLD Snooping	Previously, when receiving ICMPv6 packets with type 0, CPU usage could become high. This issue has been resolved.	Y	Y	Y	Y
CR00038426	Environment Monitoring	Previously, when a stack member joined the stack, there was a small chance non-critical internal processes could be unnecessarily restarted. This issue has been resolved.	Y	Y	Y	Y
CR00035780	SNMP MIB Support	Previously, using SNMPWALK on the AT-LOG MIB in a VCStack could result in a system reboot. This issue has been resolved.	Y	Y	Y	Y

Level 3

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00037187	VCStack	Previously, it was possible to configure a VLAN with an IP subnet which overlapped with the internally reserved VCStacking management VLAN IP subnet. This had the potential to cause communication failure between line cards and the controller card. This issue has been resolved—it is no longer possible to configure IP addresses in the reserved management VLAN subnet.	–	–	–	Y
CR00037190	VCStack	Previously, when the silicon profile on x900 or x908 was set to "none", and resiliency links were configured on any type of version 2 XEM, this could result in higher than expected CPU usage. If the silicon profile was set to **`extended`**, and resiliency links configured on any type of version 2 XEM, the resiliency links would fail to work. Both of these issues have been resolved.	Y	–	–	–
CR00037318	IGMP Snooping	Previously, IGMP snooping settings were not always restored correctly after disabling/enabling IGMP snooping. This issue has been resolved.	Y	Y	Y	Y
CR00037663	Pluggable Transceivers	Previously, hot-swapping out SFP+ devices from an SBx81XS6 line card in an SBx8112 may not have resulted in a port down event. This issue has been resolved.	–	–	–	Y
CR00037706	Port Authentication	Previously, when the switch was configured for tri-authentication and auth-web-server blocking-mode, a greyed-out logout button was shown on the Web-authentication success page. This issue has been resolved—the logout button is no longer displayed on this page.	Y	Y	Y	Y
CR00037717	802.1x	Previously, if the **`max-auth-req`** option (**`auth supplicant-mac`** command) was set too low, the supplicant could fail to be authenticated. This issue has been resolved.	Y	Y	Y	Y
CR00038021	DHCPv6	Previously, routes dynamically added to the DHCPv6 Relay for Delegated Prefixes associated with long (24.8 day+) DHCPv6 lease times would expire immediately. This issue has been resolved.	Y	Y	Y	Y

Level 3 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00038178	System	Previously, the interface command **egress-rate-limit** would not appear in the running config of a backup stack member after a master failover. This issue has been resolved.	Y	–	–	Y

Level 4

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00038164	System	Previously, VLAN interface configuration on some stack members could occasionally be displayed in an inconsistent and non-numerical order in the output from show commands, such as **show running-config**. This issue has been resolved.	Y	Y	Y	Y

Issues Resolved in 5.4.2-3.10

AlliedWare Plus maintenance version 5.4.2-3.10 includes the resolved issues in the following table.

No Level 1 Issues

Level 2

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00037349	DHCPv4	Previously, if a loopback interface was configured, the DHCP server sometimes terminated unexpectedly and generated an error message such as the following: `Unsupported device type 772 for "lo"` This issue has been resolved.	Y	Y	Y	Y

No Level 3 Issues

No Level 4 Issues

Issues Resolved in 5.4.2-3.9

AlliedWare Plus maintenance version 5.4.2-3.9 includes the resolved issues in the following table.

No Level 1 Issues

Level 2

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00035148	MLD Snooping	Previously, receiving invalid short IPv6 packets with no UDP/TCP header could result in MLD snooping freezing. This issue has been resolved.	Y	Y	Y	–
CR00036917	IPv6	Previously, if a VCStack master left the stack, sometimes IPv6 traffic was lost and not recovered. This issue has been resolved.	Y	Y	Y	–
CR00036973	Hot swap	Previously, a XEM-12T sometimes lost its configuration after many repeated reboots. This issue has been resolved.	Y	Y	Y	–
CR00037372	VCStack	Previously, if an SBx8112 had a single CFC line card installed and its line card rebooted with some unsaved running configuration, the line card interfaces sometimes remained in the Configuring state (displayed in **`show card`** command output) after reboot. This issue has been resolved.	–	–	–	Y
CR00037597	Provisioning	Previously, if a VCStack or SBx8112 is started up in a staggered manner, in some rare cases it was possible for some stack members or a CFC to remain in the 'Init' state and never fully join the stack or switch. This issue has been resolved.	Y	Y	Y	Y
CR00035707	Provisioning	Previously, if a VCStack member joined a stack and then immediately left the stack again (within a couple of seconds of joining), then the next time it attempted to join the stack it sometimes failed and never successfully joined. This issue has been resolved.	–	Y	Y	–
CR00036756	ARP	Previously, if a nexthop remained valid for longer than 25 days, then moved port or became unreachable for a short period and then returned, there was a chance that the ARP entry for that nexthop would not be successfully relearnt. As a workaround, an entry in this state could be successfully relearnt by momentarily shutting down the VLAN and then renabling the VLAN. This issue has been resolved; the work around is no longer necessary.	Y	Y	Y	Y
CR00037393	Provisioning	Previously, if VCStack members or CFCs booted up in a staggered manner, the port states were occasionally inconsistent between stack members or CFCs. This issue has been resolved.	Y	Y	Y	Y
CR00037812	802.1x	Previously, when the DHCP server used by the authentication process received an irregular DHCP Request packet whose option length field was incorrect, the authentication process sometimes failed. This issue has been resolved.	Y	Y	Y	Y
CR00037018	File System	Previously, creating, editing and copying files in Flash memory sometimes failed. This issue has been resolved.	Y	Y	Y	Y

Level 2 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00037554	XEM1XP	Previously, when a XEM1XP port connected to an external link partner was brought down (**shutdown**) and up again (**no shutdown**), the link sometimes failed to come up again. This issue has been resolved.	Y	–	–	–

No Level 3 Issues

No Level 4 Issues

Issues Resolved in 5.4.2-3.8

AlliedWare Plus maintenance version 5.4.2-3.8 includes the resolved issues in the following table.

No Level 1 Issues

Level 2

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036446	System	Previously, setting a large number (256 or more) of static IPv6 multicast commands sometimes resulted in a system restart. This issue has been resolved.	Y	Y	Y	Y
CR00036448	MLD Snooping	Previously, MLD reports could in some cases be sent with no source address, causing them to be dropped by compliant MLDv2 routers. This issue has been resolved.	Y	Y	Y	Y
CR00036609	Port Security	Previously, when port security was used on an x900, SBx908, or SBx8112, the system could sometimes reboot unexpectedly when either: ? entering **switchport port-security** configuration commands during operation for switchport interfaces that were already linked up, or ? **switchport port-security** was configured and a VCStack member or SBx8112 card was rebooted. This issue has been resolved.	Y	–	–	Y
CR00036777	GARP	Previously, when GVRP was creating a large number of dynamic VLANs in response to received GVRP packets, and a GVRP port was linkdown, a system reboot could sometimes occur. This issue has been resolved.	Y	Y	Y	Y

Level 2 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036834	RIP	Previously, when unicast RIP neighbors were configured with the **neighbor** command, the RIP version configured on the outgoing interface was overridden by the globally configured RIP version. This issue has been resolved— if unicast neighbors are configured for RIP, the RIP version configured on the interface (if any) will be used instead of the globally configured RIP version.	Y	Y	Y	Y
CR00036836	System	Previously, if an internal process (IMI) restarted while the switch was running, some interface commands in the previous running configuration were occasionally not fully restored. This issue has been resolved.	Y	Y	Y	Y
CR00036856	PIM-DM v4	Previously, if an upstream router was sending a multicast stream, but was not sending state refresh messages, PIM dense-mode sometimes temporarily lost multicast routes and therefore dropped some multicast traffic. This issue has been resolved.	Y	Y	Y	–
CR00037089	DHCPv4	Previously: ? the DHCP server sometimes responded to requests that arrived on VLANs that had no IP address, and ? when both the DHCP server and DHCP relay were configured on the same switch, the DHCP server sometimes generated malformed packets. These issues have been resolved.	Y	Y	Y	Y
CR00037090	ACL	Previously, if a large number of ACLs and/or QoS class-maps were configured, it was possible for a policy-based route entry to operate incorrectly. This issue has been resolved.	Y	–	–	–
CR00037320	Port Configuration	Previously, when the Tx strand in the fibre connected to a XEM-2XP port went down, this was not always detected. This issue has been resolved.	Y	–	–	–
CR00035636	DHCPv6	Previously, IPv6 DHCP-relay operated correctly on one VLAN, but did not always relay correctly on multiple VLANs. This issue has been resolved. The switch can now relay multiple DHCP requests from multiple DHCP IPv6 clients on different VLANs.	Y	Y	Y	–

Level 3

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00035391	DHCPv6	Previously, DHCPv6 relay did not correctly add a static route for delegated prefixes. This issue has been resolved. DHCP6 relay now adds static routes to downstream routers that request and are delegated IPv6 prefixes from DHCP servers. These routes are removed when the client releases the prefix or when the route expires.	Y	Y	Y	Y

Level 3 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036102	File System	Previously, a VCStack or SBx8112 could sometimes fail to synchronise for a new release because of a lack of available flash memory. This issue has been resolved. There is now increased memory available for synchronisation.	Y	Y	Y	Y
CR00036516	SNMP	Previously, the following interface counters returned 0 when queried with SNMP: ifInMulticastPkts 1.3.6.1.2.1.31.1.1.1.2 ifInBroadcastPkts 1.3.6.1.2.1.31.1.1.1.3 ifOutMulticastPkts 1.3.6.1.2.1.31.1.1.1.4 ifOutBroadcastPkts 1.3.6.1.2.1.31.1.1.1.5 ifInOctets 1.3.6.1.2.1.2.2.1.10 ifInUnicastPkts 1.3.6.1.2.1.2.2.1.11 ifInNonUnicastPkts 1.3.6.1.2.1.2.2.1.12 ifInErrors 1.3.6.1.2.1.2.2.1.14 ifOutOctets 1.3.6.1.2.1.2.2.1.16 ifOutUnicastPkts 1.3.6.1.2.1.2.2.1.17 ifOutNonUnicastPkts 1.3.6.1.2.1.2.2.1.18 ifOutErrors 1.3.6.1.2.1.2.2.1.2 This issue has been resolved.	–	–	–	Y
CR00036852	802.1x	Previously, when a web-authentication client sent an irregular DHCP request (one that has the length of the Requested IP Address option set to zero), this could cause the authentication process to restart. This issue has been resolved.	Y	Y	Y	Y

Level 4

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00035833	802.1x	Previously, if a port was configured with MAC-AUTH, 802.1x and the auth-fail VLAN, then if a supplicant failed authentication, a static FDB entry, containing that supplicant's MAC address, became permanently lodged in the MAC table. The only way to clear out this entry was to reboot the switch. This issue has been resolved.	Y	Y	Y	Y

Issues Resolved in 5.4.2-3.7

AlliedWare Plus maintenance version 5.4.2-3.7 includes the resolved issues in the following table.

No Level 1 Issues

Level 2

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00035996	IPv4	Previously, transferring jumbo frames using the IP helper to a receive buffer that was too small to receive them resulted in a system reboot. This issue has been resolved—frames that are too large for the receive buffer are discarded, and a log message is generated to record this.	Y	Y	Y	Y
CR00036219	L2 Switching	Previously, when Loop Protection was enabled, after a VCStack master failover, Loop Detection Frames (LDF) were not correctly transmitted. This issue has been resolved.	Y	Y	Y	Y
CR00036359	OSPFv2	A previous change to the range for the **ip ospf retransmit-interval** interface command allowed a range of 1-65535, but this change was accidently removed and the original range of 5-65535 was reinstated. This issue has been resolved—the range has now been restored to 1-65535.	Y	Y	Y	Y
CR00035466	DOS Detection	Previously, DoS detection did not correctly detect or take action (e.g., shutdown) for ports attacked by ping-of-death or tear-drop attacks. This issue has been resolved.	–	Y	Y	–
CR00035636	DHCPv6	Previously, IPv6 DHCP-relay operated correctly on one VLAN, but did not always relay correctly on multiple VLANs. This issue has been resolved. The switch can now relay multiple DHCP requests from multiple DHCP IPv6 clients on different VLANs.	Y	Y	Y	–
CR00035823	CPU	Previously, a high traffic load on a VCStack using port security could occasionally result in a system reboot. This issue has been resolved.	–	Y	Y	–
CR00036096	GARP	Previously, if the GVRP CLI debugging was enabled (**debug gvrp cli** command), using the command **no gvrp timer** could result in a system reboot. This issue has been resolved.	Y	Y	Y	Y
CR00036414	L2 Switching	Previously, if loop protection put a port into the link down state, the **shutdown** command could result in the port being put into the link up state. This issue has been resolved.	–	Y	Y	–
CR00036442	MLD	Previously, the switch recorded solicited-node multicast addresses—such as IPv6 Neighbor Solicitation messages for tracking MLD. It is not necessary to track these solicited-node addresses. This issue has been resolved—solicited-node multicast addresses are now not tracked for MLD.	Y	Y	Y	Y
CR00035840	L2 Switching	Previously, when using silicon profile extended mode on an SBx908 VCStack with a very heavy traffic load, a stack separation could occur. This issue has been resolved	Y	–	–	–
CR00036327	VCStack, Triggers	Previously, if a trigger was setup to run a script on a master failover, then this sometimes caused a system reboot. This issue has been resolved.	Y	Y	Y	Y

Level 2 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036360	PIM-DMv4	Previously, a newly added PIM-DM interface was always added as a downstream interface of existing multicast route entries regardless of the existence of local joins or neighbouring router on this interface. As a result, data may have flowed unnecessarily to the switch from the upstream router and never be pruned. This issue has been resolved.	Y	Y	Y	Y
CR00036432	System	Previously, static routing entries that overlapped directly connected routes could sometimes be lost from hardware and were not re-added automatically when a port unrelated to those static routes was shut down. This issue has been resolved.	Y	Y	Y	Y
CR00035354	DHCPv4	Previously, for very short defined DHCP lease times (only seconds), simultaneous requests of several leases could trigger a DHCP relay exception. This issue has been resolved.	Y	Y	Y	Y
CR00035818	System	Previously, if a VCStack was under a high load for a prolonged period from traffic with random source MAC addresses, then the following could occur: ? the **show platform table fdb** command could become non-responsive; ? the **clear mac address-table dynamic** command could result in a system reboot. These issues have been resolved.	Y	Y	Y	–
CR00035858	Port Security	Previously on a VCStack, when a switch port entered security shutdown after detection of excessive end-user source MAC addresses (that is, after the specified **port-security maximum** limit was exceeded), these symptoms could occur: ? the **show int st** command response could be sluggish; ? a system reboot sometimes occurred. This issue has been resolved.	–	Y	Y	–
CR00036241	DHCPv4	Changes were made to the DHCP Server to secure against DOS security vulnerabilities identified by ISC (Internet Systems Consortium) ISC references CVE-2012-3571 and CVE-2012-3954.	Y	Y	Y	Y
CR00036262	VCStack	Previously, when Web-authentication was applied to a static or LACP link aggregator on a VCStack, a DHCP broadcast packet on a port in the link aggregator was sometimes duplicated, leading to a failure of client IP allocation. This issue has been resolved.	–	Y	Y	–
CR00036452	RSTP	Previously, if a switch-port's up/down state was toggled repeatedly and rapidly, a small memory loss could occur. If this was repeated a very large number of times, a system reboot could eventually occur. This issue has been resolved.	Y	Y	Y	Y
CR00035747	Port Configuration	Previously, if a VCStack member contained a 10Gbps XEM module, then on reboot the port configuration settings for that XEM were sometimes lost. This issue has been resolved.	Y	–	–	–

Level 2 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036688	QoS	Previously on XEM-12Sv2 and XEM-2XP modules, QoS configurations to classify CPU control traffic and change priority queues (e.g., the **set queue** command and premarking DSCP map) would sometimes fail to change the queues. During times of congestion, this QoS failure could adversely affect control protocols such as OSPF. This issue has been resolved.	Y	–	–	–
CR00035770	Port Authentication	Previously, when port authentication was enabled on an aggregated port, communication between supplicants and the guest VLAN sometimes failed. This issue has been resolved.	Y	Y	Y	Y

Level 3

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036171	Log	Previously, jumbo frames could caused a large number of unnecessary messages for what was the normal operation of Jumbo frames when handled by CPU. An example message was: `'Puma but i (1) != 0'` This issue has been resolved—these log messages are no longer generated.	Y	–	–	–
CR00034959	L2 Switching	Previously, when Forwarding Database entries were deleted, for example when VLANS were deleted, unnecessary EXFX debug log messages were generated. This issue has been resolved.	Y	–	–	Y
CR00036346	Triggers	Previously, an interface trigger on a VCStack sometimes failed after master failover. This issue has been resolved.	Y	Y	Y	–
CR00036046	L2 Switching	Previously, using the **no mru** command reset the MRU to 1500. This issue has been resolved. It now resets it correctly to the default value 1518.	–	Y	Y	–
CR00036085	L2 Switching	Previously, on an SBx908 in silicon profile extended mode, incorrect error messages were generated periodically from an unused switch silicon facility. The messages incorrectly appeared to show Forwarding Database entry deletion problems. This issue has been resolved.	Y	–	–	–
CR00036276	System	Previously, using the unsupported command **show platform phy** on a port on a VCStack backup member could result in a system reboot of the backup member. This issue has been resolved.	–	Y	Y	–
CR00036490	L3 Switching	Previously, on some XEM models (XEM-2XT, XEM-2XS, XEM-2XP, XEM-12Tv2, XEM-12Sv2, XEM-24T), small throughput loss (0.001% of line rate) could occur on high traffic loads. This was caused through CRC errors. This issue has been resolved.	Y	–	–	–

Level 3 (cont.)

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00036547	SNMP MIB Support	Previously, when a link aggregator interface was created, it was not set in the SNMP MIB dot1dBasePortTable and the dot1dTpPortTable until after the config was saved and the device was rebooted. This issue has been resolved. The new interface is now stored in the MIBs when it is created.	Y	Y	Y	Y

Level 4

CR	Module	Description	x900/x908	x600	x610	SBx8100
CR00035892	EDAC	Previously, log messages occasionally reported spurious and incorrect EDAC PCI/X errors. These messages no longer appear.	Y	Y	Y	–

Issues Resolved in 5.4.2-3.6

AlliedWare Plus maintenance version 5.4.2-3.6 includes the resolved issues in the following table.

No Level 1 Issues

Level 2 Issues

CR	Module	Description
CR00035296	OSPFv2	Previously, OSPF routers on a common Ethernet network could occasionally, under very rare circumstances, get into a state where the **show ip ospf route** command would show no routes and new routes were not being added correctly. The circumstances were: ? two neighbours lose 2-way connectivity ? other neighbours (especially the Designated Router) do not lose connectivity ? no other changes in the routing state of the network occur (route recalculation would clear the error state) This issue has been resolved.
CR00035844	SSH	Previously, it was not possible to disallow SSH authentication by Challenge Response (also known as keyboard-interactive). This has been resolved by expanding the scope of the **password** authentication setting to also cover Challenge Response, so that if **password** is switched off, so is Challenge Response. Similarly, enabling password authentication also enables Challenge Response authentication.

Level 2 Issues (cont.)

CR	Module	Description
CR00035853	ACL	x600 and x610 only. Previously, when DHCP snooping was enabled on an aggregated link in a VCStack, then a master failover in the stack caused the new master to reboot. This issue has been resolved.
CR00036149	VLAN	Previously, a small amount of memory that was allocated when a VLAN was created was not freed when the VLAN was deleted. This issue has been resolved.

Issues Resolved in 5.4.2-2.5

AlliedWare Plus maintenance version 5.4.2-2.5 includes the resolved issues in the following tables.

Level 1 Issues

CR	Module	Description
CR00035515	Boot Setup	Previously, in rare cases the switch could get into an unresponsive state when a syslog restart occurred. This issue has been resolved.

Level 2 Issues

CR	Module	Description
CR00035444	IPv4 Unicast Routing	Previously, static routes whose nexthop fell within the address ranges covered by both an interface route and a blackhole route were rendered inactive as a result of the blackhole route. This behaviour was incorrect and in contradiction to previous AlliedWare Plus behaviour. This issue has been resolved. The correct behaviour has been restored and as long as there is a selected route that can resolve the nexthop, blackhole routes will not stop the static route being made active.
CR00035527	VRRP	Previously, if the switch received a VRRPv3 advertisement packets, it would cause the switch to become unresponsive due to high CPU load. This issue has been resolved.
CR00035543	ARP	Previously, when the silicon profile on an SBx908 was configured to **`extended`** mode, ARP broadcasts were not sent to the CPU to be learnt/processed. This issue has been resolved.
CR00035697	PoE	Previously, at higher temperatures, the x610-POE switches could incorrectly read that the PSU fans were running too slowly, and turn off PoE power. This issue has been resolved.

Level 2 Issues (cont.)

CR	Module	Description
CR00035701	L3 Switching	SBx908, x908 only. Previously, a VCStack healthcheck failure could occasionally result in a system restart. This issue has been resolved.
CR00035725	SNMP MIB Support	Previously, if SNMP GETNEXT was used with a large index for the 'logIndex' MIB variable in the AT-LOG MIB, a system reboot could occur. This issue has been resolved.
CR00035780	SNMP MIB Support	Previously, using SNMPWALK on the AT-LOG MIB in a VCStack could result in a system reboot. This issue has been resolved.
CR00036022	EPSR	x610 only. Previously, when the switch CPU was under heavy load, loop protection could fail to detect a loop. This issue has been resolved.
CR00035724	L3 Switching	Previously, when the SBx908 or x900 was operating in **`platform silicon-profile extended`** mode, it ignored DHCP requests made to it. This issue has been resolved.
CR00035770	Port Authenticati on	Previously, when port authentication was enabled on an aggregated port, communication between supplicants and the guest VLAN sometimes failed. This issue has been resolved.
CR00035851	IGMP	Previously, if IGMP was enabled on several IP-enabled VLAN interfaces, IGMP could get into an enabled but inactive state on some VLANs, either on boot-up, or after the VLAN state changed to “up”. This issue has been resolved.
CR00035859	L3 Switching	Previously, in extremely rare situations, a XEM-2XS would not correctly detect when a link partner interface was disabled. This issue has been resolved.
CR00035965	ARP	Previously, when running VRRP on an x900 or SBx908, heavy system load could sometimes lead to internal congestion resulting in some dropped packets. In this state it was also possible for an alignment trap to be generated while processing ARPs. These issues have been resolved.
CR00035715	OSPFv2	x600, x610 only. Previously it was possible for a stack member that became the Disabled master, and then shortly afterwards the Active master to loose OSPF configuration entries. This issue has been resolved.

Level 3 Issues

CR	Module	Description
CR00035396	802.1x	Previously, when the **supplicant-mac** feature was enabled, **max-reauth-req** was always set to zero, resulting in the switch never restarting authentication. This issue has been resolved—the switch will now use the value of **max-reauth-req** to determine the number of times to restart authentication.
CR00035470	VCStack	Previously, in rare cases, a unit failed to fully join a VCStack on the first attempt, then automatically rebooted and joined the stack successfully. This issue has been resolved—the unit will join successfully the first time.
CR00035525	SD card reader	In previous 5.4.2 versions, the SD card reader failed to read Sandisk 2Gb cards. This issue has been resolved.
CR00035679	SNMP MIB Support	Previously, the BRIDGE_MIB MIB objects dot1dStpPortPriority, dot1dStpPortPathCost, and dot1dStpPortEnable could not be set. This issue has been resolved.
CR00035782	L2 Switching	Previously on a XEM-2XP, XEM-2XT or XEM-2XS installed in an x900 or SBx908, if DHCP snooping was enabled globally but disabled locally on a particular VLAN, and a DHCP packet was received on an STP blocked port in that VLAN, the packet was processed and re-flooded when it should have been dropped. A similar issue occurred with the loop-detection feature where a loop-detect frame would be incorrectly re-flooded when received on an STP blocked port. These issues has been resolved.
CR00035795	ACL	x900, SBx908 only. Previously, when the platform routing ratio was **ipv4only**, standard ACLs with rule indexes greater than 1535 were not correctly removed from ports by the **no access-group** command.
CR00035798	Port Authentication	Previously, auth roaming would fail when used with either the dynamic VLAN creation feature or guest VLAN feature. This issue has been resolved.
CR00035930	SNMP MIB Support	Previously, it was not possible to read objects in the dot1qVlanStaticTable. This issue has been resolved.
CR00035481	VLAN	Previously, a private VLAN could not be deleted through the command line interface. This issue has been resolved.
CR00035726	VCStack	Previously, in extremely rare cases when a VCStack backup member was power cycled, it would fail to properly rejoin the stack, and remain in the init state. This issue has been resolved.
CR00035914	ACL	Previously, when an SBx908 was configured with silicon profile extended mode, global ACLs added to a XEM-2XP, XEM-2XS, and XEM-2XT used more hardware entries than required. This issue has been resolved.

Level 3 Issues (cont.)

CR	Module	Description
CR00035929	**Port Configuration**	Previously, on XEM-2XP, XEM-2XS and XEM-2XT, it was possible for the link state in software to get out of sync with the link state in hardware, following multiple link state changes in a short time period. This issue has been resolved.
CR00036033	**PIM-SM v4**	Previously, in mixed mode PIM-DM and PIM-SM networks, when a switch acting as RP for a PIM-SM domain received a (*,*,RP) join from a directly connected border PIM router which was forwarding a multicast group from a PIM-DM domain, the interface receiving the group could be incorrectly added to the outgoing interface list, resulting in a multicast loop between the RP and PIM border router. This issue has been resolved.
CR00036045	**Ping Polling**	Previously, ping polling pings were often sent up to one second later than the configured frequency. This issue has been resolved.

CR00036189	**RADIUS**	Previously, when a user cancelled the removal of a CA certificate:
		`no crypto pki trustpoint local  > y/n?`
		`> n`
		the certificate remained in Flash (correct behaviour), however the command line indicated the removal was successful, and the CA configuration was removed from the running config despite the user selecting to cancel the operation.
		This issue has been resolved.

CR00036243	**Port Configuration**	Previously, on an SBx908 with a XEM-STK installed in bay 1, if the **`clear port counter`** command was entered, unnecessary log messages could be generated, and in rarer cases a system reboot could occur. This issue has been resolved.
CR00034082	**VCStack**	When multiple stack members started up at the same time it was possible on a 4 x x600 stack for one of the members to join the stack late, even though it started up at the same time as the other members. This issue has been resolved and all members now join at the same time.

CR00035379	**VCStack**	At startup on an SBx908 VCStack, under exceptional circumstances, an error condition was encountered which meant switchport interfaces would not function correctly across the VCStack. This error condition was always characterised by the following log message:
		`user.err awplus-2 HSL[1431]: HSL: ERROR: Failed to connect to stack member 2 HIP (Connection timed out)`
		This problem may have also occurred on AW+ releases 5.3.4 to 5.4.2 inclusive. One symptom of this problem was that all ports on one VCStack member would fail to be added to any static-channels, resulting in error messages loading the startup configuration, such as:
		`static-channel-group 1 -- % Failed to aggregate port1.1.1`
		This issue has been resolved.

Level 4 Issues

CR	Module	Description
CR00035228	IGMP	Previously, when an x600 and x610 VCStack tried to send IGMP packets from a VCStack backup-member, it generated packet transmission error messages, e.g.: `NSM[1237]: [IGMP-ENCODE] : sendto() failed on vlan1: Operation not permitted(1)` This issue has been resolved.
CR00035475	Pluggable Transceivers	Previously, an SFP+ module still showed as inserted after it had been removed from a XEM-2XS, and the SFP information was not be updated until a new SFP+ was inserted. This issue has been resolved.

CR00035493	SSH	SFTP download is not supported on the switch. Previously, when the command:
		`copy <filename> sftp://<server-ip-address>`
		was entered, a misleading error message was generated. The error message has been corrected to:
		`This feature is not supported.`
		(SFTP upload is supported.)

CR00035613	SNMP MIB Support	Previously, when performing SNMP GETNEXT for objects with large instance-IDs in the AT-LOG-MIB, lexicographic ordering of the returned values failed. This issue has been resolved.
CR00035624	Scripting, VCStack	Previously, remote-login to a back-up stack member failed if the user was remotely authenticated via either TACACS+ or RADIUS. This issue has been resolved.
CR00035665	Console	SBx908 and x908 only. Previously, the **`show platform`** command displayed the silicon profile defined in the running config. It now displays the actual hardware silicon profile.
CR00035666	SNMP MIB Support	Previously, when performing SNMP GETNEXT for certain objects in the dot1qBridge MIB, lexicographic ordering of the returned values failed. This issue has been resolved.
CR00035686	SNMP MIB Support	Previously, when performing SNMP GETNEXT for certain objects in the dot1dBridge MIB, lexicographic ordering of the returned values failed. This issue has been resolved.
CR00035689	ACL	SBx908 and x908 only. Previously, when the switch was configured for **`platform routingratio ipv4andipv6`** (the default), the **`show platform classifier statistics utilization brief`** displayed the classifier utilisation and limits incorrectly by a factor of 2. This issue has been resolved.

Level 4 Issues (cont.)

CR	Module	Description
CR00035942	SNMP MIB Support	Previously, an SNMP GET of the logMessage MIB object in the AT-LOG MIB could return a string longer than 255 characters, which conflicted with the MIB definition. This issue has been resolved.

Issues Resolved in 5.4.2-1.4

AlliedWare Plus maintenance version 5.4.2-1.4 includes the resolved issues in the following tables.

No Level 1 Issues

Level 2

CR	Module	Description
CR00035524	VCStack	Previously, when a SBx908 VCStack member was rebooted and rejoined the stack, messages could be displayed, such as the following: *NSM[1438]: 3 audit inconsistencies detected - stack member 1 should reboot* These messages could be displayed for a few minutes at which point the system would recover and continue operating as normal. This issue only affected SBx908 VCStacks running AW+ 5.4.1. This issue has been resolved.
CR00035557	VCStack	In rare circumstances, a VCStack backup member could sometimes end up with an inconsistent interface state and configuration. If this problem occurred, the following log messages would be seen: local6.crit awplus-2 NSM[1470]: 3 audit inconsistencies detected - stack member 2 should reboot This issue only occurred if: ? A command failed when the backup member loaded the startup configuration, but the command succeeded on the VCStack master. ? Another VCStack member rejoined This issue has been resolved.
CR00035638	ACL	Previously, when a Stack of x610 switches started in a standalone environment, the ACL configuration was not set on the provisioned static aggregator. This issue has been resolved.
CR00035715	OSPFv2	Previously it was possible for a stack member that became the Disabled master, and then shortly afterwards the Active master to loose OSPF configuration entries. This issue has been resolved.

No Level 3 Issues

No Level 4 Issues

Issues Resolved in 5.4.2-1.3

AlliedWare Plus maintenance version 5.4.2-1.3 includes the resolved issues in the following tables.

No Level 1 Issues

Level 2

CR	Module	Description
CR00034946	Hotswap	Previously, hotswapping a XEM-2XP could cause a system re-boot. This issue has been resolved.
CR00035006	Command Shell	Previously, use of the small selection of commands that support the ability to accept an interface range could have resulted in a memory leak, or in rare cases a process termination. This issue has been resolved.

Level 3

CR	Module	Description
CR00034082	VCStack	When multiple stack members started up at the same time it was possible on a 4 x x600 stack for one of the members to join the stack late, even though it started up at the same time as the other members. This issue has been resolved and all members now join at the same time.
CR00034941	VCStack	Previously, the forwarding of data could be slow on one backup member following the formation of an 8-unit VCStack. This issue has been resolved.
CR00034979	Pluggable Transceivers	Previously, flow control could be on momentarily when the configuration for flow control was off. This issue has been resolved.
CR00035007	OSPFv2	If OSPF was gracefully restarted many times in quick succession, the OSPF routes could be prematurely removed from the hardware, causing a brief period of traffic loss. This issue has been resolved.
CR00035054	VLAN	x600 switches only. Previously, hosts in an isolated VLAN could ping the IP address of the primary VLAN of the switch. This issue has been resolved.
CR00035139	SNMP	Previously, setting the **ipForwarding** MIB variable to **notForwarding** via a MIB Browser would cause all IP based operations to fail, including responses to PING and SNMP requests. This issue has been resolved. Now support for configuring the **ipForwarding** MIB variable via SNMP has been disabled.

CR00035147	**Mirroring**	x600 switch only. Previously, configuring a mirror port to **`direction transmit`** then changing the configuration to **`direction receive`** would cause the **`no mirror`** command to fail to remove the mirror port entry from hardware. This would prevent the next configured mirror port from working. This issue has been resolved.
CR00035165	**ACL**	Previously, deleting an interface with an access group attached and attaching the group to a new interface could produce misleading error messages. This issue has been resolved.
CR00035246	**Pluggable Transceivers**	The transmit and receive power values and thresholds displayed by DDM (Digital Diagnostics Monitoring) in **`show system pluggable`** diagnostics were incorrectly interpreted. This issue has been resolved.
CR00035249	**Pluggable Transceivers**	Previously, when an 10G SFP+ installed in an x610 series switch with 7m passive cables connected was under full load, a very small amount of packet loss may have occurred. This issue has been resolved.
CR00035271	**PoE**	If **`no service power-inline`** was configured in the startup-config, the switch could hang and timeout during startup. This issue has been resolved.
CR00035288	**Logging**	When the **`log email`** command was configured on a stack, both the Master and Backup members tried to send emails to the SMTP Server. Since only the Master is actually able to send emails, the Backup members were using up Flash memory to store the unsent emails and unterminated processes. This caused unnecessary log messages and memory leaks on the backup members. This issue has been resolved.
CR00035319	**DHCPv4**	Previously, the DHCP server would not NAK (Negatively Acknowledge) some lease renewal requests from a client that had been moved to a different subnet. This issue has been resolved Most DHCP clients, upon physically being moved to a new subnet, go to the INIT-REBOOT state and immediately request a lease reconfirmation, in which case this issue does not occur. However, some clients, e.g. Windows XP DHCP client, do not change state when moved to a new subnet, and it is for these clients that this change has been made.
CR00035379	**VCStack**	SBx908 only At startup on an SBx908 VCStack, under exceptional circumstances, an error condition was encountered which meant switchport interfaces would not function correctly across the VCStack. This error condition was always characterised by the following log message: `user.err awplus-2 HSL[1431]: HSL: ERROR: Failed to connect to stack member 2 HIP (Connection timed out)` This problem may have also occurred on AW+ releases 5.3.4 to 5.4.2 inclusive. One symptom of this problem was that all ports on one VCStack member would fail to be added to any static-channels, resulting in error messages loading the startup configuration, such as: static-channel-group 1 -- % Failed to aggregate port1.1.1 This issue has been resolved.

CR00035389	**OSPFv3**	Previously, static IPv6 blackhole routes (with a Null destination interface) failed to be redistributed to routing protocols. This issue has been resolved.
CR00035390	**OSPFv3**	Previously, static IPv6 blackhole routes (with a Null destination interface) would fail to be added to the hardware table and FIB. This issue has been resolved.

Level 4

CR	Module	Description
CR00035222	**Logging**	Previously, any unsent email messages were stored in memory until memory ran out. This could produce a memory leak. This issue has been resolved. Now the maximum allowable memory storage size is 5% of total RAM for emails. A memory check ensures that new messages are not added to the queue when the limit is reached. This will prevent emails from consuming too much of the system resource.
CR00035302	**SSH**	An error message "sh: a: unknown operand" is printed on the console when a SCP or SFTP is used on a PC to copy a file from a switch. The cause was found to be missing quotes around a string in the startshell script. This issue has been resolved.
CR00035332	**Pluggable Transceivers**	Previously, errors messages logged by the switch when initialising pluggable transceivers did not include which port was affected. This information is now included in the log message. This issue has been resolved.
CR00035410	**BGPv4**	Previously, the switch would not accept 0 as a valid parameter in the **`ip community-list`** command: (in the form of X:0 nor 0:X). This issue has been resolved.

Issues Resolved in 5.4.2-0.2

AlliedWare Plus maintenance version 5.4.2-0.2 includes the resolved issues in the following tables.

No Level 1 Issues

Level 2

CR	Module	Description
CR00035164	**VRF**	If BGP is used to distribute routes between VRF instances, then previously a stack master failover would result in some traffic interruption. This has been resolved.
CR00035227	**OSPF**	Previously, if OSPF was used on a stack of 5 or more switches, and a double device failure occurred in that stack, and one of the failing units was the stack master, then the redistribution of connected routes into OSPF could sometimes fail to operate correctly after the fail over. This issue has been resolved.

No Level 3 Issues

No Level 4 Issues

Errata to the Software Reference

The following update is a correction to the Software Reference for AlliedWare Plus 5.4.2

ping ipv6

This command sends a query to another IPv6 host (send Echo Request messages).

Note `Use of the interface parameter keyword, plus an interface or an interface range, with this command is only valid when pinging an IPv6 link local address.`

Syntax `ping ipv6 [<host>|<ipv6-address>] [repeat {<1-2147483647>|continuous}] [size <10-1452>] [interface <interface-list>] [timeout <1-65535>]`

Parameters in the output of the **`ping ipv6`** command

Parameter	Description
<ipv6-addr>	The destination IPv6 address. The IPv6 address uses the format X:X::X:X.
<hostname>	The destination hostname.
repeat	Specify the number of ping packets to send.
<1-2147483647>	Specify repeat count. The default is 5.

Parameters in the output of the **ping ipv6** command

Parameter	Description
size <*10-1452*>	The number of data bytes to send, excluding the 8 byte ICMP header. The default is 56 (64 ICMP data bytes).
interface <*interface-list*>	The interface or range of configured IP interfaces to use as the source in the IP header of the ping packet.
timeout <*1-65535*>	The time in seconds to wait for echo replies if the ARP entry is present, before reporting that no reply was received. If no ARP entry is present, it does not wait.
repeat	Specify the number of ping packets to send.
<*1-2147483647*>	Specify repeat count. The default is 5.
continuous	Continuous ping.
size <*10-1452*>	The number of data bytes to send, excluding the 8 byte ICMP header. The default is 56 (64 ICMP data bytes).
timeout <*1-65535*>	The time in seconds to wait for echo replies if the ARP entry is present, before reporting that no reply was received. If no ARP entry is present, it does not wait.

Mode User Exec and Privileged

Example

```
awplus# ping ipv6 2001:2001:0db8::a2
```

Related Commands traceroute ipv6

show ip sockets

Use this command to display information about the IP or TCP sockets that are present on the switch. It includes TCP, UDP listen sockets, displaying associated IP address and port. The information displayed for established TCP sessions includes the remote IP address, port, and session state. Raw IP protocol listen socket information is also displayed for protocols such as VRRP and ICMP6, which are configured to receive IP packets with the associated protocol number.

Syntax `show ip sockets`

Mode User Exec and Privileged Exec

Usage Use this command to verify that the socket being used is opening correctly. If there is a local and remote endpoint, a connection is established with the ports indicated.

Note that this command does not display sockets that are used internally for exchanging data between the various processes that exist on the device and are involved in its operation and management. It only displays sockets that are present for the purposes of communicating with other external devices.

Example To display ip sockets currently present on the device, use the command:

```
awplus# show ip sockets
```

Output Example output from the **show ip sockets** command

```
Socket information

Not showing 40 local connections
Not showing 7 local listening ports

Typ Local Address            Remote Address           State
tcp 0.0.0.0:111              0.0.0.0:*                LISTEN
tcp 0.0.0.0:80               0.0.0.0:*                LISTEN
tcp 0.0.0.0:23               0.0.0.0:*                LISTEN
tcp 0.0.0.0:443              0.0.0.0:*                LISTEN
tcp 0.0.0.0:4743             0.0.0.0:*                LISTEN
tcp 0.0.0.0:873              0.0.0.0:*                LISTEN
tcp :::23                    :::*                     LISTEN
udp 0.0.0.0:111              0.0.0.0:*
udp 226.94.1.1:5405          0.0.0.0:*
udp 0.0.0.0:161              0.0.0.0:*
udp :::161                   :::*
raw 0.0.0.0:112              0.0.0.0:*                112
raw :::58                    :::*                     58
raw :::112                   :::*                     112
```

Parameters in the output of the **show ip sockets** command

Parameter	Description
Not showing <number> local connections	This field refers to established sessions between processes internal to the device, that are used in its operation and management. These sessions are not displayed as they are not useful to the user. <number> is some positive integer.

Parameters in the output of the **show ip sockets** command

Parameter	Description
Not showing <number> local listening ports	This field refers to listening sockets belonging to processes internal to the device, that are used in its operation and management. They are not available to receive data from other devices. These sessions are not displayed as they are not useful to the user. <number> is some positive integer.
Typ	This column displays the type of the socket. Possible values for this column are: **tcp** : IP Protocol 6. **udp** : IP Protocol 17. **raw** : Indicates that socket is for a non port-orientated protocol (i.e. a protocol other than TCP or UDP) where all packets of a specified IP protocol type are accepted. For raw socket entries the protocol type is indicated in subsequent columns.
Local Address	For TCP and UDP listening sockets this shows the destination IP address (either IPv4 or IPv6) and destination TCP or UDP port number for which the socket will receive packets. The address and port are separated by ':'. If the socket will accept packets addressed to any of the switch's IP addresses, the IP address will be 0.0.0.0 for IPv4 or :: for IPv6. For active TCP sessions the IP address will display which of the switches addresses the session was established with. For raw sockets this displays the IP address and IP protocol for which the socket will accept IP packets. The address and protocol are separated by ':'. If the socket will accept packets addressed to any of the switch's IP addresses, the IP address will be 0.0.0.0 for IPv4 and ::. IP Protocol assignments are described at: http://www.iana.org/assignments/protocol-numbers
Remote Address	For TCP and UDP listening sockets this shows the source IP address (either IPv4 or IPv6) and source TCP or UDP port number for which the socket will accept packets. The address and port are separated by ':'. If the socket will accept packets addressed from any IP address, the IP address will be 0.0.0.0 for IPv4 or :: for IPv6. This is the usual case for a listening socket. Normally for a listen socket any source port will be accepted. This is indicated by ". For active TCP sessions the IP address will display the remote address and port the session was established with. For raw sockets the entry in this column will be 0.0.0.0: or ::: for IPv4 and IPv6, respectively.
State	This column shows the state of the socket. For TCP sockets this shows the state of the TCP state machine. For UDP sockets this column is blank. For raw sockets it contains the IP protocol number. The possible TCP states are: **LISTEN, SYN-SENT,SYN-RECEIVED, ESTABLISHED, FIN-WAIT-1, FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK, TIME-WAIT, CLOSED** RFC793 contains the TCP state machine diagram with Section 3.2 describing each of the states.
Typ	This column displays the type of the socket. Possible values for this column are: **tcp** : IP Protocol 6. **udp** : IP Protocol 17. **raw** : Indicates that socket is for a non port-orientated protocol (i.e. a protocol other than TCP or UDP) where all packets of a specified IP protocol type are accepted. For raw socket entries the protocol type is indicated in subsequent columns.

show ip traffic

Use this command to display statistics regarding IP traffic sent and received by all interfaces on the switch, showing totals for IP and IPv6 and then broken down into sub-categories such as TCP, UDP, ICMP and their IPv6 equivalents when appropriate.

Syntax `show ip traffic`

Mode User Exec and Privileged Exec

Example To display IP traffic statistics, use the command:

```
awplus# show ip traffic
```

Output Example output from the **show ip traffic** command

```
IP:
     261998 packets received
     261998 delivered
     261998 sent
     69721 multicast packets received
     69721 multicast packets sent
     23202841 bytes received
     23202841 bytes sent
     7669296 multicast bytes received
     7669296 multicast bytes sent
IPv6:
     28 packets discarded on transmit due to no route
ICMP6:
UDP6:
UDPLite6:
TCP:
     0 remote connections established
     40 local connections established
     7 remote listening ports
     7 local listening ports
     261 active connection openings
     247 passive connection openings
     14 connection attempts failed
     122535 segments received
     122535 segments transmitted
     14 resets transmitted
     227 TCP sockets finished time wait in fast timer
     155 delayed acks sent
     21187 headers predicted
     736 pure ACKs
     80497 pure ACKs predicted
UDP:
     139468 datagrams received
     139468 datagrams sent
UDPLite:
```

Parameters in the output of the **show ip traffic** command

Parameter	Description
IPv4	IPv4 counters
IPv6	IPv6 counters
received packets with no route	Received packets with no route
truncated packets received	Truncated packets received
multicast packets received	Multicast packets received
multicast packets sent	Multicast packets sent
broadcast packets received	Broadcast packets received
broadcast packets sent	Broadcast packets sent
bytes received	Bytes received
bytes sent	Bytes sent
multicast bytes received	Multicast bytes received
multicast bytes sent	Multicast bytes sent
broadcast bytes received	Broadcast bytes received
broadcast bytes sent	Broadcast bytes sent
packets received	Packets received
packets received with invalid headers	Packets received with invalid headers
oversize packets received	Oversize packets received
packets received with no route	Packets received with no route
packets received with invalid address	Packets received with invalid address
packets received with unknown protocol	Packets received with unknown protocol
truncated packets received	Truncated packets received
received packets discarded	Received packets discarded
received packets delivered	Received packets delivered
forwarded packets transmitted	Forwarded packets transmitted
packets transmitted	Packets transmitted
packets discarded on transmit	Packets discarded on transmit
packets discarded on transmit due to no route	Packets discarded on transmit due to no route
fragment reassembly timeouts	Fragment reassembly timeouts
fragment reassembly required	Fragment reassembly required
fragment reassembly OK	Fragment reassembly OK
fragment reassembly failures	Fragment reassembly failures
fragmentations succeeded	Fragmentations succeeded
fragmentations failed	Fragmentations failed
fragments created	Fragments created
ICMP6	ICMPv6 counters
messages received	Messages received
errors received	Errors received
messages sent	Messages sent
TCP	TCP counters
remote connections established	Remote connections established
local connections established	Local connections established
remote listening ports	Remote listening ports

Parameters in the output of the **show ip traffic** command (cont.)

Parameter	Description
local listening ports	Local listening ports
active connection openings	Active connection openings
passive connection openings	Passive connection openings
connection attempts failed	Connection attempts failed
connection resets received	Connection resets received
segments received	Segments received
segments transmitted	Segments transmitted
retransmits	Retransmits
bad segments received	Bad segments received
resets transmitted	Resets transmitted
datagrams received	Datagrams received
received for unknown port	Received for unknown port
datagrams sent	Datagrams sent
syncookies sent	Syncookies sent
syncookies received	Syncookies received
syncookies failed	Syncookies failed
embryonic resets	Embryonic resets
sockets pruned	Sockets pruned
ICMPs out of window	ICMPs out of window
ICMPs dropped due to lock	ICMPs dropped due to lock
ARPs filtered	ARPs filtered
TCP sockets finished time wait in fast timer	TCP sockets finished time wait in fast timer
time wait sockets recycled by time stamp	Time wait sockets recycled by time stamp
time wait sockets killed	Time wait sockets killed
delayed acks sent	Delayed acks sent delayed acks further delayed because of locked socket
delayed acks lost	Delayed acks lost
listening socket overflows	Listening socket overflows
listening socket drops	Listening socket drops
headers predicted	Headers predicted
pure ACKs	Pure ACKs
pure ACKs predicted	Pure ACKs predicted
losses recovered by TCP Reno	Losses recovered by TCP Reno
losses recovered by SACK	Losses recovered by SACK
SACKs renegged	SACKs renegged
detected reordering by FACK	Detected reordering by FACK
detected reordering by SACK	Detected reordering by SACK
detected reordering by TCP Reno	Detected reordering by TCP Reno
detected reordering by sequence	Detected reordering by sequence
full undos	Full undos
partial undos	Partial undos
SACK undos	SACK undos
loss undos	Loss undos

Parameters in the output of the **show ip traffic** command (cont.)

Parameter	Description
segments lost	Segments lost
lost retransmits	Lost retransmits
TCP Reno failures	TCP Reno failures
SACK failures	SACK failures
loss failures	Loss failures
fast retransmits	Fast retransmits
forward retransmits	Forward retransmits
retransmits in slow start	Retransmits in slow start
timeouts	Timeouts
TCP Reno recovery failures	TCP Reno recovery failures
SACK recovery failures	SACK recovery failures
collapsed segments received	Collapsed segments received
DSACKs sent for old packets	DSACKs sent for old packets
DSACKs sent for out of order segments	DSACKs sent for out of order segments
DSACKs received	DSACKs received
DSACKs received for out of order segments	DSACKs received for out of order segments
connections reset due to unexpected SYN	Connections reset due to unexpected SYN
connections reset due to unexpected data	Connections reset due to unexpected data
connections reset due to early user close	Connections reset due to early user close
connections aborted due to lack of memory	Connections aborted due to lack of memory
connections aborted due to timeout	Connections aborted due to timeout
connections aborted due to lingering	Connections aborted due to lingering
connection aborts due to connection failure	Connection aborts due to connection failure
TCP memory pressure events	TCP memory pressure events
SACKs discarded	SACKs discarded
Old DSACKs ignored	Old DSACKs ignored
DSACKs ignored without undo	DSACKs ignored without undo
Spurious RTOs	Spurious RTOs
TCP MD5 Not Found	TCP MD5 Not Found
TCP MD5 Unexpected	TCP MD5 Unexpected
TCP SACKs shifted	TCP SACKs shifted
TCP SACKs merged	TCP SACKs merged
TCP SACK shift fallback	TCP SACK shift fallback
UDP	UDP Counters
UDPLite	UDPLite Counters
UDP6	UDPv6 Counters
UDPLite6	UDPLitev6 Counters
datagrams received	Datagrams received
datagrams received for unknown port	Datagrams received for unknown port
datagram receive errors	Datagram receive errors
datagrams transmitted	Datagrams transmitted

Parameters in the output of the **show ip traffic** command (cont.)

Parameter	Description
datagrams received	Datagrams received
datagrams received for unknown port	Datagrams received for unknown port
datagram receive errors	Datagram receive errors
datagrams transmitted	Datagrams transmitted