Management Software

AT-S101

User's Guide

For use with the AT-GS950/8POE
Gigabit Ethernet WebSmart Switch

Version 1.0.0

613-000985 Rev. A

Allied Telesis™

Contents

Figures

Preface

The AT-S101 Management Software is the operating system for the AT-GS950/8POE Gigabit Ethernet WebSmart Switch. This guide explains how to use the management software to control and monitor the operating parameters of the AT-GS950/8POE switch.

This Preface contains the following sections:

Document Conventions

This document uses the following conventions:

Note
Notes provide additional information.

Caution
Cautions inform you that performing or omitting a specific action may result in equipment damage or loss of data.

Warning
Warnings inform you that performing or omitting a specific action may result in bodily injury.

Where to Find Web-based Guides

The installation and user guides for all Allied Telesis products are available in portable document format (PDF) on our web site at **www.alliedtelesis.com**. You can view the documents online or download them onto a local workstation or server.

For details about the features and functions of the AT-GS950/8POE switch, see the following installation guides on our web site:

- *AT-GS950/8POE Gigabit Ethernet WebSmart Installation Guide* (part number 613-000989)

Contacting Allied Telesis

This section provides Allied Telesis contact information for technical support as well as sales and corporate information.

Online Support

You can request technical support online by accessing the Allied Telesis Knowledge Base: **www.alliedtelesis.com/support/kb.aspx**. You can use the Knowledge Base to submit questions to our technical support staff and review answers to previously asked questions.

Email and Telephone Support

For Technical Support via email or telephone, refer to the Support section of the Allied Telesis web site: **www.alliedtelesis.com**.

Warranty

The AT-GS950/8POE Gigabit Ethernet WebSmart Switch is covered under a Lifetime Warranty (Two Years Fan & Power Supply). For warranty information, go to the Allied Telesis web site at **www.alliedtelesis.com**.

Returning Products

Products for return or repair must first be assigned a return materials authorization (RMA) number. A product sent to Allied Telesis without an RMA number will be returned to the sender at the sender's expense. For instructions on how to obtain an RMA number, go to the Support section on our web site at **www.alliedtelesis.com/support.rma.aspx**.

Sales or Corporate Information

You can contact Allied Telesis for sales or corporate information through our web site at **www.alliedtelesis.com**.

Management Software Updates

New releases of the management software for our managed products are available from the following Internet sites:

- Allied Telesis web site: **www.alliedtelesis.com**
- Allied Telesis FTP server: **ftp://ftp.alliedtelesis.com**

If the FTP server prompts you to log on, enter "anonymous" as the user name and your email address as the password.

Chapter 1

Starting a Web Browser Management Session

This chapter contains the procedures for starting, using, and quitting a web browser management session on the AT-GS950/8POE switch. This chapter includes the following sections:

Establishing a Remote Connection to the Web Browser Interface

The AT-GS950/8POE switch is shipped with a pre-assigned IP address of 192.168.1.1. You must set your local PC on the same subnet as the pre-assigned IP address for your initial logon.

After your initial login, you may want to assign a new IP address to your switch. To manually assign an IP address to the switch, refer to “Configuring an IP Address, Subnet Mask and Gateway Address” on page 22. To configure the switch to obtain its IP configuration from a DHCP server, refer to “Enabling and Disabling the DHCP Client” on page 27.

You must set your local PC to the same subnet as the preassigned IP address.

Note
Enhanced stacking, a feature of other Allied Telesis Layer 2 and Layer 2+ managed switches, is not supported by the AT-GS950/8POE switch.

To start a web browser management session, perform the following procedure:

1. Start your web browser.
2. In the URL field of the browser, enter 192.168.1.1 which is the default IP address of the switch. See Figure 1.

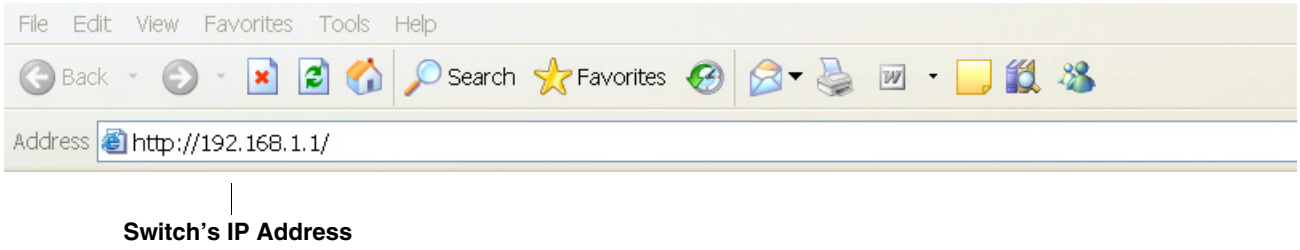

Figure 1. Entering a Switch’s IP Address in the URL Field

The AT-S101 Management Software displays the login dialog box, shown in Figure 2.

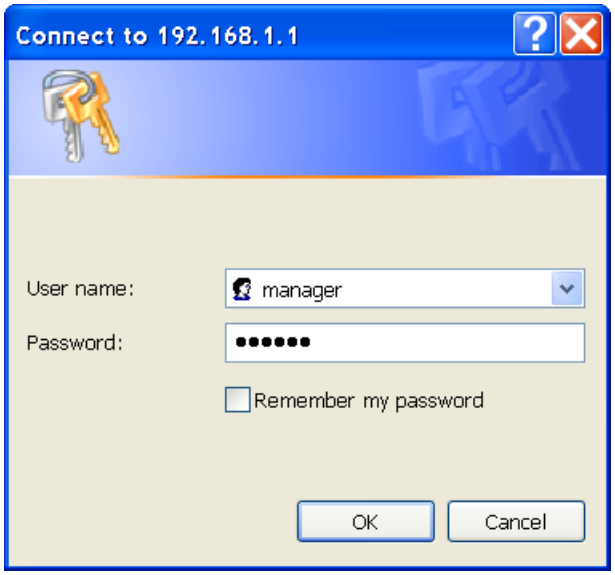

Figure 2. AT-S101 Login Dialog Box

3. Enter the AT-S101 management login user name and password. The default user name is "manager" and the default password is "friend." Then press OK. The login name and password are case-sensitive.

   The Switch Information Page is displayed. See Figure 3 on page 16.

To change the user name and password, refer to “Configuring System Management Information” on page 29.

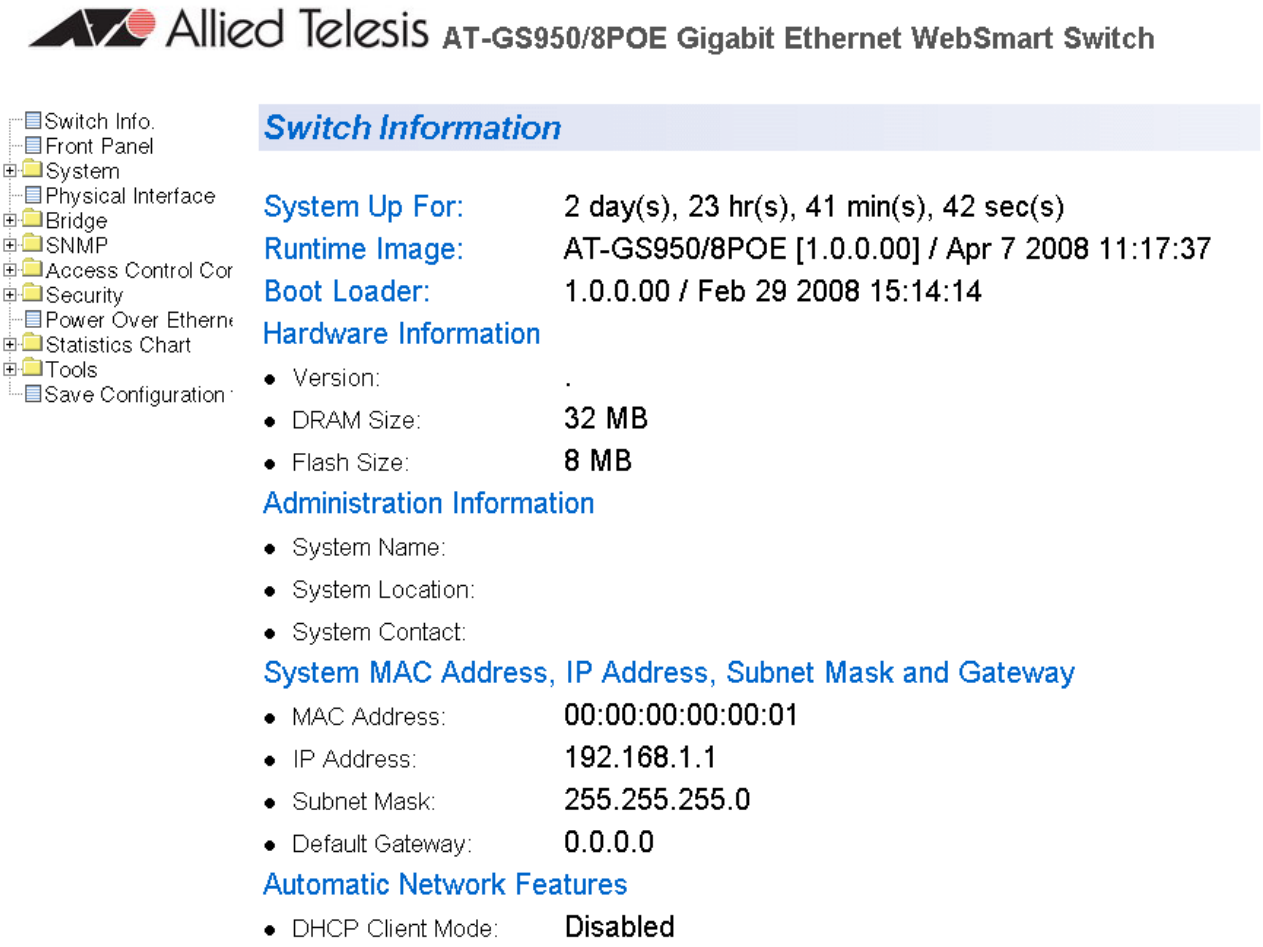

Figure 3. Switch Information Page

The main menu is on the left side of the home page. It consists of the following folders and web pages:

- Switch Info.
- Front Panel
- System
- Physical Interface
- Bridge
- SNMP
- Access Control Config.
- Security
- Power Over Ethernet
- Statistics Chart

- Tools
- Save Configuration

4. To see the front panel of the switch, select **Front Panel** from the menu on the left side of the page.

   The AT-S101 Management Software displays the front of the switch. The window contains an image of the front of the switch. Ports that have a link to an end node are green. Ports without a link are grey. An example of a front panel is shown in Figure 4.

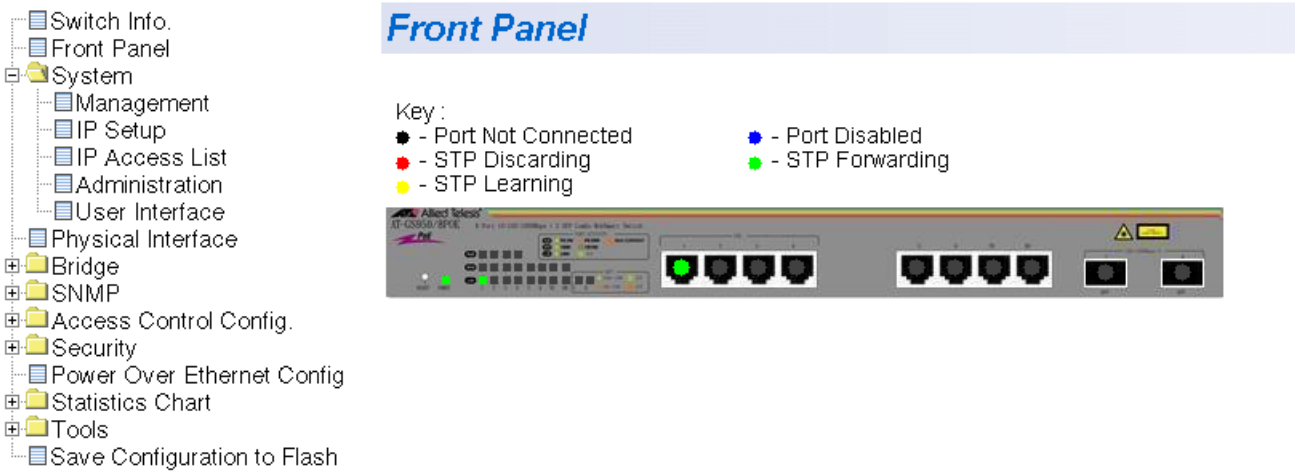

Figure 4. Front Panel Page

A web browser management session remains active even if you link to other sites. You can return to the management web pages anytime as long as you do not quit the browser.

Web Browser Tools

You can use the web browser tools to move around the management pages. Selecting **Back** on your browser's toolbar returns you to the previous display. You can also use the browser's **bookmark** feature to save the link to the switch.

Quitting a Web Browser Management Session

To exit a web browser management session, close the web browser.

Chapter 2

Basic Switch Parameters

This chapter provides procedures to perform basic switch activities such as reassigning the IP address, enabling the DHCP Client, configuring new usernames and passwords, and rebooting the system.

This chapter contains the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Configuring an IP Address, Subnet Mask and Gateway Address

This procedure explains how to change the IP address, subnet mask, and gateway address to the switch. Before performing the procedure, note the following:

- A gateway address is only required if you want to remotely manage the device from a management station that is separated from the switch by a router.
- To configure the switch to automatically obtain its IP configuration from a DHCP server on your network, go to “Enabling and Disabling the DHCP Client” on page 27.

To change the switch’s IP configuration, perform the following procedure:

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. From the **System** folder, select **IP Setup**.

   The IP Setup Page is shown in Figure 5.

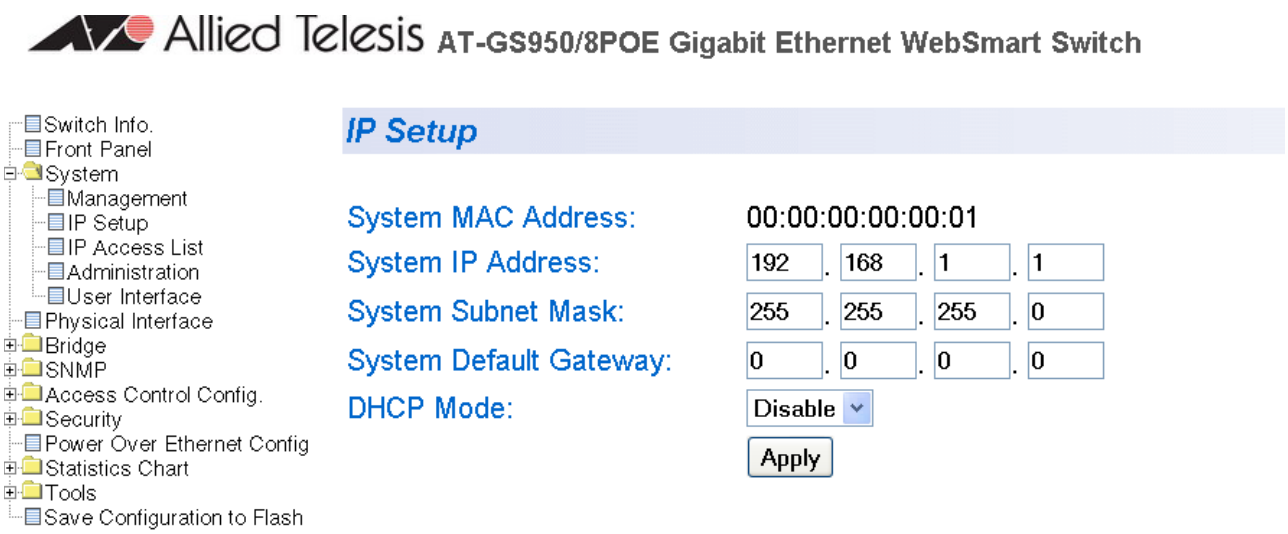

Figure 5. IP Setup Page

3. Change the IP configuration parameters by entering new information in the following fields:

   System MAC Address
   This parameter displays the MAC address of the switch. You cannot change this parameter.

   System IP Address
   Displays the current IP address of the switch. To change the IP address, enter a new IP address.

   System Subnet Mask
   Displays the current subnet mask of the switch. To change the subnet mask, enter a new subnet mask.

   System Default Gateway
   Displays the default gateway of the switch. To change the default gateway, enter a new gateway.

   DHCP Mode
   For information about setting this parameter, refer to "Enabling and Disabling the DHCP Client" on page 27.

4. Click **Apply**.

   Note
   Changing the IP address ends your management session. To resume managing the device, enter the new IP address of the switch in the web browser's URL field, as shown in Figure 1 on page 14.

5. After you log on to the switch with the new IP address, select **Save Configuration to Flash** to save the new IP address to memory.

   Caution
   If you do not select **Save Configuration to Flash**, the IP address may revert to its default setting when you power cycle the switch.

Setting Up the IP Access List

The IP Access List feature, when enabled, restricts remote access to management software by means of a user-configured list of IP addresses. It does not restrict the management ping response activity, only web access to the management software.

Note
By default, the IP Access List feature is disabled.

The procedures in this section describe how to enable or disable the IP Access List feature and how to add or remove IP addresses from the list. See the following sections:

- “Creating an IP Access List” on page 24
- “Deleting an IP Address” on page 26

Note
You cannot modify an existing IP address.

Creating an IP Access List

To create a list of restricted IP addresses, perform the following procedure:

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. From the **System** folder, select **IP Access List**.

   The IP Access List Page is shown in Figure 6.

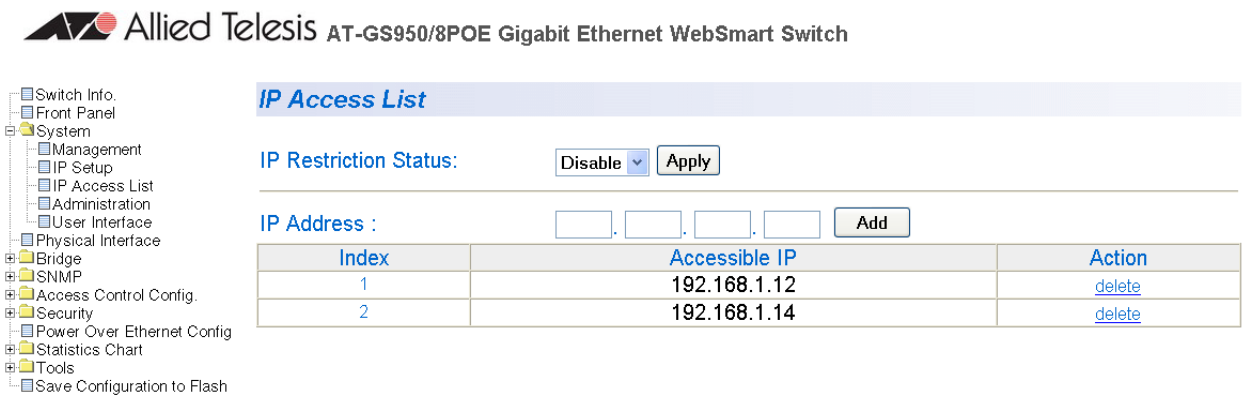

Figure 6. IP Access List Page

3. To set the IP restriction status, select Disable or Enable in the pull-down menu next to the **IP Restriction Status** field. Then click **Apply**.

   By default, the IP Restriction Status field is set to Disable.

4. Enter an IP address that you want to prevent from accessing the switch in the xxx.xxx.xxx.xxxx format next to the **IP Address** field. Then click **Add**.

   The IP address is added to the IP Access List Table.

5. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting an IP Address

To delete an IP address from the IP Access List, perform the following procedure:

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. From the **System** folder, select **IP Access List**.

   The IP Access List Page is displayed. See Figure 6 on page 24.

3. Select **delete** next to the IP address that you want to remove.

   The IP address is removed from the IP Access List Table.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Enabling and Disabling the DHCP Client

Since the AT-GS950/8POE switch is a web-only switch and does not have a local console connection, you must be careful when you change the IP address of the switch by enabling the DHCP client. To look up the IP address on a DHCP server, you must have the MAC address of the AT-GS950/8POE switch. Once the switch obtains a new IP address from the DHCP server, the switch becomes inaccessible and the MAC address can no longer be viewed in the AT-S101 software.

Before you enable the DHCP client, record the switch's MAC address. You can view the MAC address on the System Information Page when you first log onto the switch. See "Viewing System Information" on page 36. Or, you can see the MAC address on the label affixed to the switch.

If the switch power cycles before you save the new configuration, the software reverts to the default IP address value. Or, if you press the Reset button before you save the DHCP client on the switch, the software reverts the default IP address value. In either case, the IP address value is 192.168.1.1.

This procedure explains how to activate and deactivate the DHCP client on the switch. When the client is activated, the switch obtains its IP configuration, its IP address and subnet mask, from a DHCP server on your network. Before performing the procedure, note the following:

- By default, the DHCP client is disabled on the switch.
- The DHCP client does not support BOOTP.
- After you enable DHCP, you will end the current management session. Log on with the new IP address (provided by your system administrator) using the procedure described in "Establishing a Remote Connection to the Web Browser Interface" on page 14.

Caution
Record the MAC address of your switch before you begin this procedure.

To activate or deactivate the DHCP client on the switch, perform the following procedure:

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. From the **System** folder, select **IP Setup**.

   The IP Setup Page is shown in Figure 5 on page 22.

3. From the pull-down menu next to the **DHCP Mode** field, select **Enable** or **Disable**.

   By default, this field is set to **Disable**.

4. Click **Apply**.

   If you enable the DHCP client, the web server connection to the switch is lost.

   If you disable the DHCP client, note the new **System IP Address** value that you assigned to the switch. Record this value for future use.

Caution
Enabling or disabling DHCP ends your management session.

Caution
If you do not select **Save Configuration to Flash**, the DHCP mode reverts to its default setting of 192.168.1.1 when you power cycle the switch.

5. Log on to the switch with the new IP address and immediately save your configuration by selecting **Save Configuration to Flash** from the menu on the left side of the page.

   If you enable DHCP and then save your configuration, you save the IP address on the DHCP server.

   If you disable DHCP, enter a new IP address, and then save your configuration, you have saved the DHCP setting and the new IP address on the switch.

Configuring System Management Information

This section explains how to assign a name to the switch, as well as the location of the switch and the name of the switch's administrator. Entering this information is optional.

To set a switch's administration information, perform the following procedure:

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. From the **System** folder, select **Management**.

   The Management Page is shown in Figure 7.

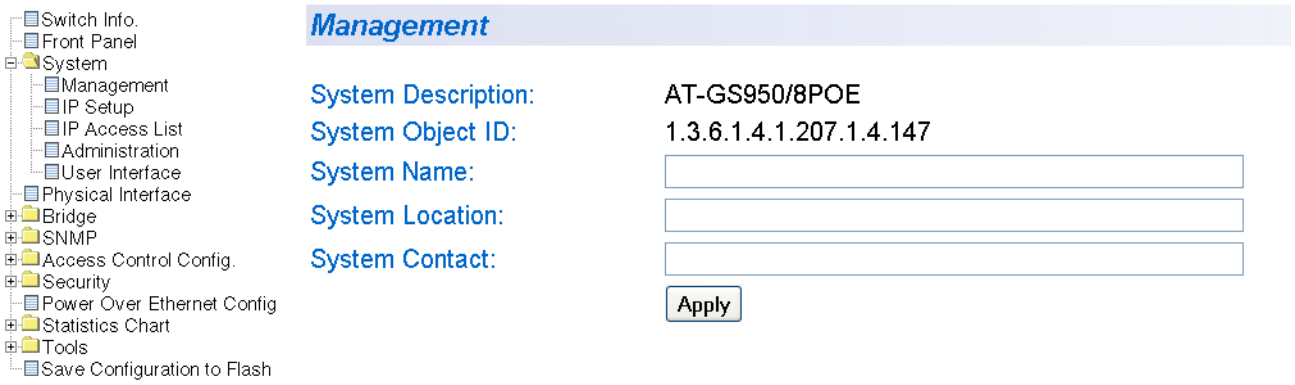

Figure 7. Management Page

3. Configure the following parameters as necessary:

   System Description
   Specifies the model number of the switch. You cannot change this parameter.

   System Object ID
   Indicates the unique SNMP MIB object that identifies the AT-GS950/8POE switch model. You cannot change this parameter.

   System Name
   Specifies a name for the switch, for example, Sales. The name is optional and may contain up to 50 characters.

Note
Allied Telesis recommends that you assign a name to the switch. A name can help you identify the switch when you manage it and can also help you avoid performing a configuration procedure on the wrong switch.

System Location
Specifies the location of the switch. The location is optional and may contain up to 50 characters.

System Contact
Specifies the name of the network administrator responsible for managing the switch. This contact name is optional and may contain up to 50 characters.

4. Click **Apply**.

5. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Configuring System Administration Information

This section explains how to enable password protection and create users in the web interface. See the following sections:

Adding System Administration Information

To set a switch's administration information, perform the following procedure:

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. From the **System** folder, select **Administration**.

The Administration Page is shown in Figure 8.

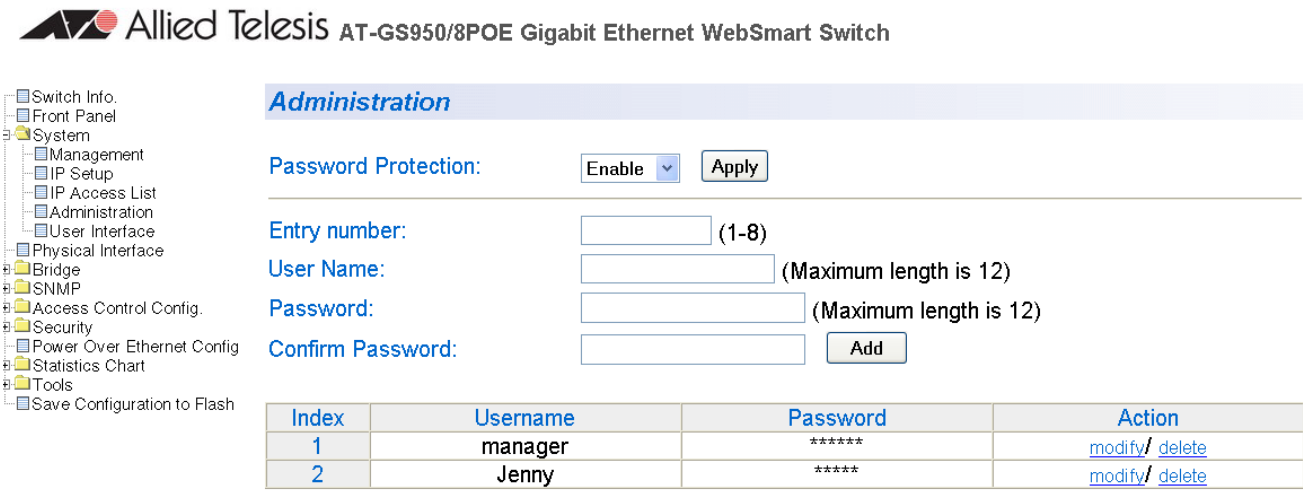

Figure 8. Administration Page

3. To enable or disable password protection, select Enable or Disable from the pull-down menu next to the **Password Protection** field. Then click **Apply**.

   You can control login authentication by enabling password protection which requires a user to supply a password when logging onto the switch. If you disable password protection, a user can login without inputting a password. By default, this field is set to Enable.

4. To create an entry number, type 1 through 8 in the box next to the Entry number field.

   This value appears as the Index value in the Administration table at the bottom of the page.

5. To create a user name, enter a user name in the box next to the **User Name** field.

   You can enter a value of up to 12 alphanumeric characters.

6. To add a password for the above user name, enter a password of up to 12 alphanumeric characters in the box next to the **Password** field.

7. To confirm the above password, retype the password in the box next to the **Confirm Password** field.

8. Click **Add** to activate your changes on the switch.

9. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Modifying Administration Information

To modify the a user name password, perform the following procedure.

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. From the **System** folder, select **Administration**.

   The Administration Page is shown in Figure 8 on page 31.

3. Select the user name that you want to change and click **modify**.

The Modify Administration Page is displayed. See Figure 9.

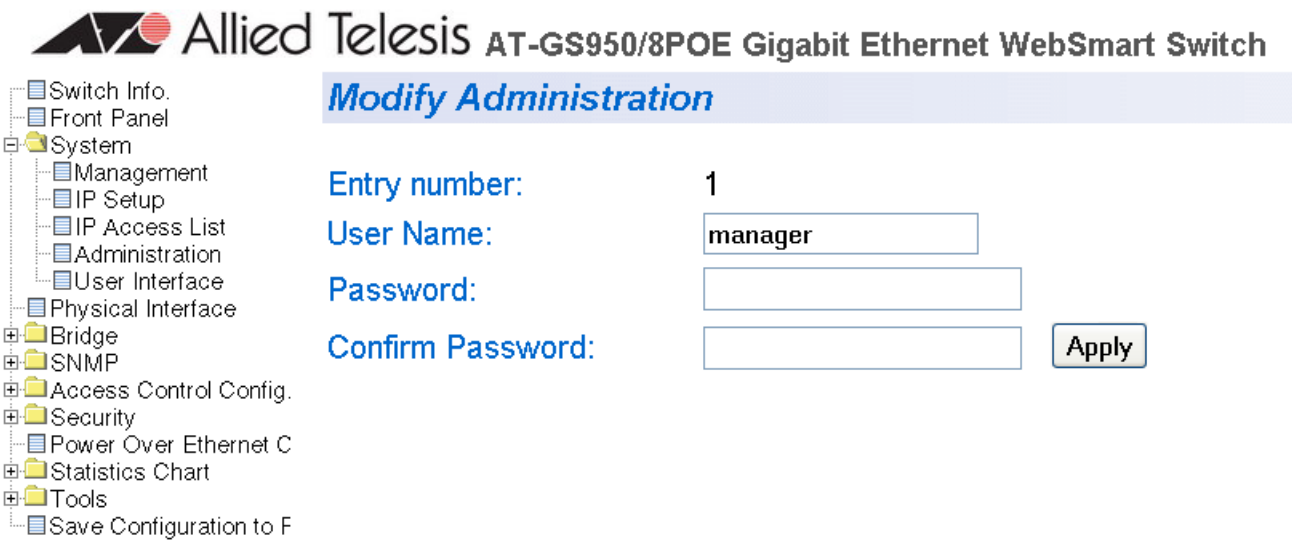

Figure 9. Modify Administration Page

4. To change a password, enter a password of up to 12 alphanumeric characters in the box next to the **Password** field.
5. To confirm the above password, retype the password in the box next to the **Confirm Password** field.
6. Click **Apply** to activate your changes on the switch.
7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting Administration Information

To delete a user name, perform the following procedure.

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.
2. From the **System** folder, select **Administration**.

   The Administration Page is shown in Figure 8 on page 31.
3. Select the user name that you want to delete and click **delete**.

   The user name is removed from the Administration Table.
4. Click **Add** to activate your changes on the switch.

Setting the User Interface Configuration

This procedure explains how to adjust the user interface and security features on the switch. With this procedure you can:

- Enable an SNMP Agent. To configure the SNMP feature, see Chapter 8, “Simple Network Management Protocol (SNMP)” on page 99.
- Enable and disable the web server.

To set the switch’s user interface configuration, perform the following procedure:

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. From the **System** folder, select **User Interface**.

   The User Interface Page is shown in Figure 10.

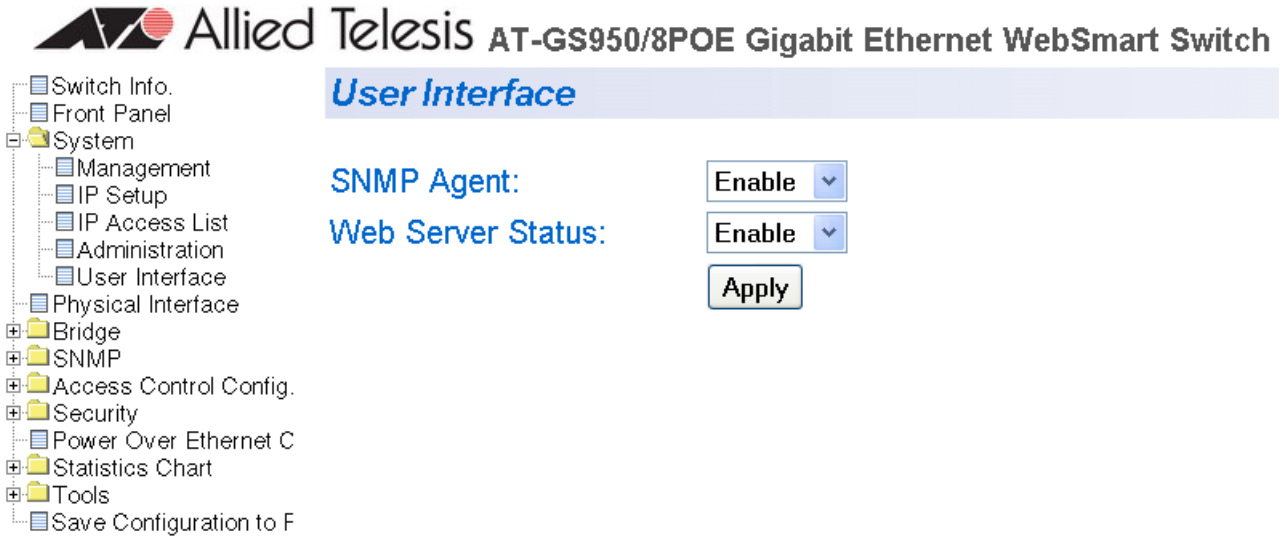

Figure 10. User Interface Page

3. To enable or disable an SNMP agent, do the following:

   a. Click the **SNMP Agent** parameter and choose **Enable** or **Disable** from the list. The default is Enable. When you enable this parameter, the SNMP agent is enabled.

   b. Click **Apply**.

4. To enable or disable the web server, do the following:

   a. Click the **Web Server** parameter and choose **Enable** or **Disable** from the pull-down menu. The default is **Enable**. When you enable this parameter, you can use a web browser to manage the switch remotely.

   Note
   Disabling the web browser automatically ends your remote management session. If this occurs, press the Reset button to recycle the power to the switch. Then logon to the switch.

   b. Click **Apply**.

5. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Viewing System Information

To view general information about the switch, perform the following procedure:

1. From the menu on the left side of the page, click the **System** folder.

   The **System** folder expands.

2. Select **Switch Info**.

   The Switch Information Page is shown in Figure 11.

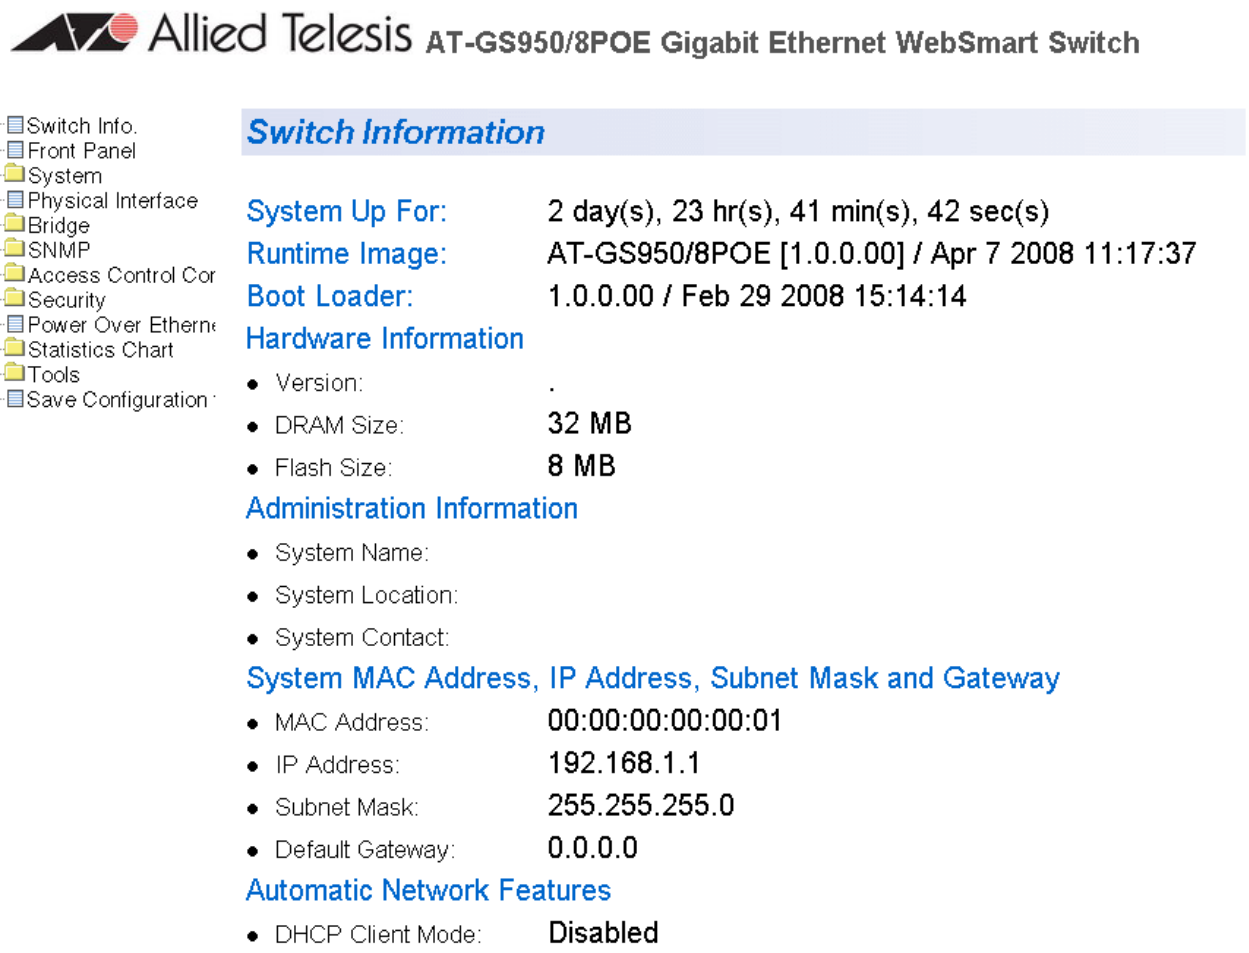

Figure 11. Switch Information Page

The Switch Information Page displays the following information:

System Up For
The number of days, hours, and minutes that the switch has been running since it was last rebooted.

Runtime Image
The version number and build date of the runtime firmware.

Boot Loader
The version number and build date of the bootloader firmware.

Hardware Information Section:

Version
The hardware version number.

DRAM Size
The size of the DRAM, in megabytes.

Flash Size
The size of the flash memory, in megabytes.

Administration Information Section:

Switch Name
The name assigned to the switch. To give the switch a name, refer to "Configuring System Management Information" on page 29.

Switch Location
The location of the switch. To specify the location, refer to "Configuring System Management Information" on page 29.

Switch Contact
The contact person responsible for managing the switch. To specify the name of a contact, refer to "Configuring System Management Information" on page 29.

System MAC Address, IP Address, Subnet Mask, and Gateway Section:

MAC Address
The MAC address of the switch. You cannot change this value.

IP Address
The IP address of the switch. Refer to "Configuring an IP Address, Subnet Mask and Gateway Address" on page 22 to manually assign an IP address or "Enabling and Disabling the DHCP Client" on page 27 to activate the DHCP client.

Subnet Mask
The subnet mask for the switch. Refer to "Configuring an IP Address, Subnet Mask and Gateway Address" on page 22 to manually assign a

subnet mask or “Enabling and Disabling the DHCP Client” on page 27 to activate the DHCP client.

Default Gateway
Default gateway’s IP address. Refer to “Configuring an IP Address, Subnet Mask and Gateway Address” on page 22 to manually assign a gateway address or “Enabling and Disabling the DHCP Client” on page 27 to activate the DHCP client.

Automatic Network Features Section:

DHCP Mode
The status of the DHCP client on the switch. For information about setting this parameter, refer to “Enabling and Disabling the DHCP Client” on page 27.

Rebooting a Switch

This procedure reboots the switch and reloads the AT-S101 Management Software from flash memory. You may want to reboot the device if you believe it is experiencing a problem.

Caution
The switch does not forward network traffic during the reboot process. Some network traffic may be lost.

To reboot a switch, perform the following procedure:

1. From the menu on the left side of the page, select the **Tools** folder.

   The **Tools** folder expands.

2. From the **Tools** folder, select **Reboot**.

   The Reboot Page is shown in Figure 12.

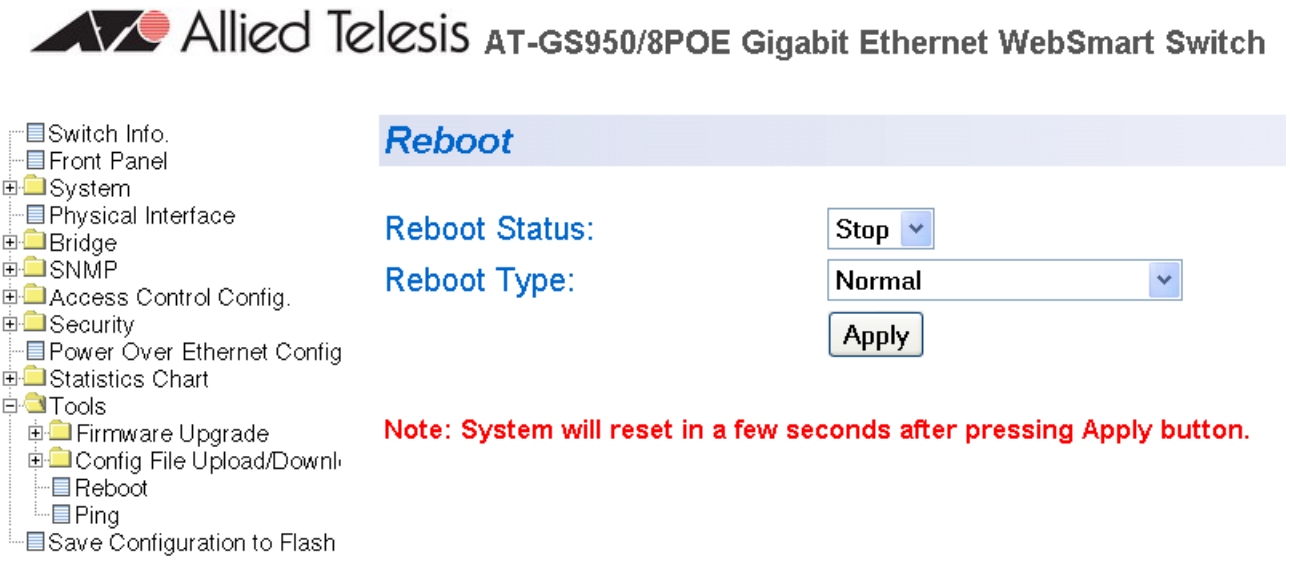

Figure 12. Reboot Page

3. For the Reboot Type, select **Normal** from the pull-down menu. This is the default setting.

Note
Two additional Reboot Type options, **Factory Default** and **Reset to Factory Default Except IP Address,** are described in “Returning the AT-S101 Management Software to the Factory Default Values” on page 44.

4. For the **Reboot Status**, use the pull-down menu to select **Start** to begin the reboot.

5. Click **Apply**.

   The switch immediately begins to reload the AT-S101 Management Software. This process takes approximately one minute to complete. You can not manage the device during the reboot. After the reboot is finished, you can log in again if you want to continue to manage the device.

Pinging a Remote System

This procedure instructs the switch to ping a node on your network. This procedure is useful in determining whether an active link exists between the switch and another network device. Note the following before performing the procedure:

- The device you are pinging must be a member of the Default VLAN. In other words, the port on the switch through which the node is communicating with the switch must be an untagged or tagged member of the Default VLAN.

To ping a network device, perform the following procedure:

1. From the menu on the left side of the page, select the **Tools** folder.

   The **Tools** folder expands.

2. From the **Tools** folder, select **Ping**.

   The Ping Test Configuration Page is displayed. See Figure 13 on page 41.

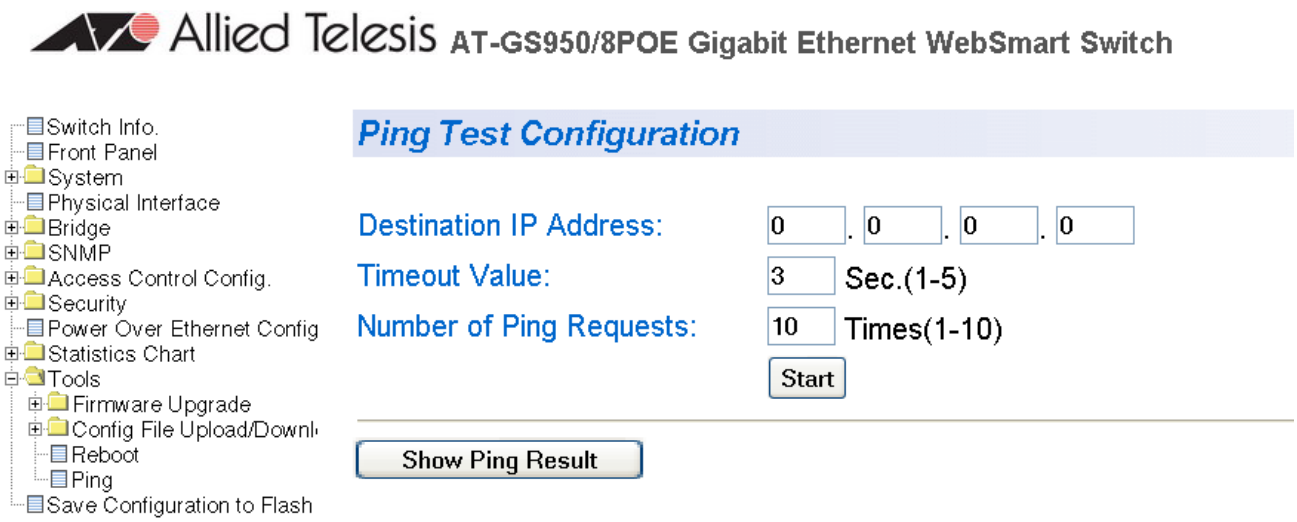

Figure 13. Ping Test Configuration Page

3. Configure the following parameters:

   Destination IP Address
   The IP address of the node you want to ping.

Timeout Value
Specifies the length of time, in seconds, the switch waits for a response before assuming that a ping has failed. The default is 3 seconds.

Number of Ping Requests
Specifies the number of ping requests you want the switch to perform. The default is 10.

4. Click **Start**.

5. To view the ping results, click **Show Ping Results**.

   A sample Ping Test Results Page is displayed. See Figure 14.

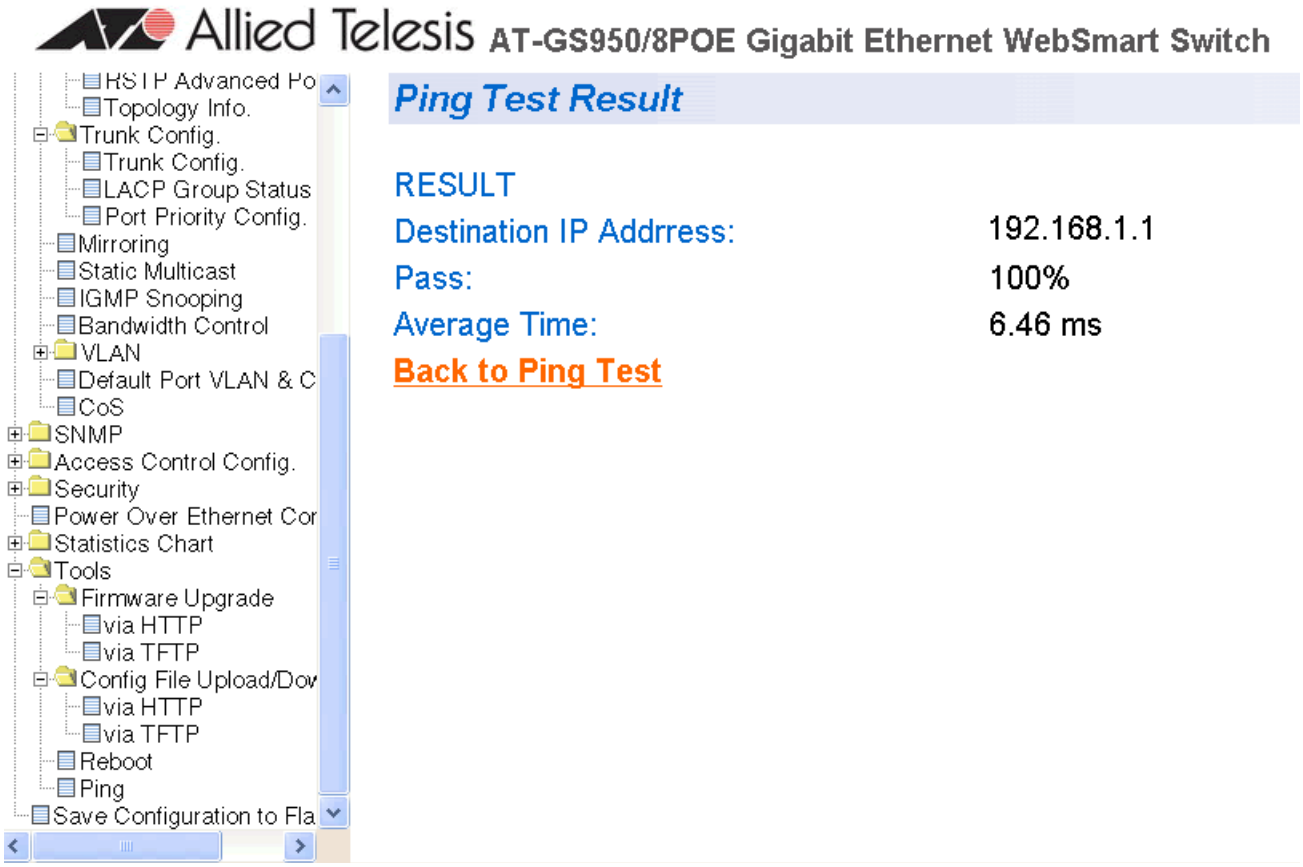

Figure 14. Ping Test Results Page

The following information is provided:

Destination IP Address
Indicates the IP address of the unit that receives the ping.

Pass
Indicates the percentage of times the ping passed.

Average Time
Indicates the time, in milliseconds, the ping was received.

6. Click **Back to Ping Test** to return to the Ping Test Configuration Page.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Returning the AT-S101 Management Software to the Factory Default Values

This procedure returns all AT-S101 Management Software parameters to their default values and deletes all tagged and port-based VLANs on the switch. The AT-S101 Management Software default values are listed in Appendix A, “AT-S101 Management Software Default Settings” on page 231.

Caution
This procedure causes the switch to reboot. The switch does not forward network traffic during the reboot process. Some network traffic may be lost.

To return the AT-S101 software to the default settings, perform the following procedure:

1. From the Tools folder, select **Reboot**.

   The Reboot Page is shown in Figure 12 on page 39.

2. For the **Reboot Type** field, use the pull-down menu to select one of the following:

   Factory Default
   Resets all switch parameters to the factory default settings, including the IP address, subnet mask, and gateway address.

   Factory Default Except IP Address
   Resets all switch parameters to the factory default settings, but retains the IP address, subnet mask, and gateway settings. If the DHCP client is enabled, it remains enabled after this reset.

3. For the **Reboot Status** field, use the pull-down menu to select **Start** to begin the reboot.

4. Click **Apply**.

   The switch is rebooted. You must wait for the switch to complete the reboot process before reestablishing your management session.

Chapter 3

Virtual LANs

This chapter contains a description of Virtual Local Area Networks (VLANs) and procedures for creating, modifying, and deleting port-based and tagged VLANs from a web browser management session. This chapter contains the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

VLAN Overview

A VLAN is a group of ports on an Ethernet switch that form a logical Ethernet segment. The ports of a VLAN form an independent traffic domain where the traffic generated by the nodes of a VLAN remains within the VLAN.

With VLANs, you can segment your network through the switch's AT-S101 Management Software and group nodes with related functions into their own separate, logical LAN segments. These VLAN groupings can be based on similar data needs or security requirements. For example, you can create separate VLANs for each department in your company, such as one for Sales and another for Accounting.

VLANs offer several important benefits:

- Improved network performance

  Network performance often suffers as networks grow in size and as data traffic increases. The more nodes on each LAN segment vying for bandwidth, the greater the likelihood overall network performance decreases.

  VLANs improve network perform because traffic stays within the separate, logical LAN segment of the VLAN. The nodes of a VLAN receive traffic only from nodes of the same VLAN. This reduces the need for nodes to handle traffic that is not destined for them. It also frees up bandwidth within all the logical workgroups.

  In addition, because each VLAN constitutes a separate broadcast domain, broadcast traffic remains within the VLAN. This too can improve overall network performance.

- Increased security

  Because data traffic generated by a node in a VLAN is restricted only to the other nodes of the same VLAN, you can use VLANs to control the flow of packets in your network and prevent packets from flowing to unauthorized end nodes.

- Simplified network management

  In addition, VLANs can simplify network management. Before the advent of VLANs, physical changes to the network often had to been made at the switches in the wiring closets. For example, if an employee changed departments, changing the employee's LAN segment assignment might require a change to the cabling of the switches.

  With VLANS, you can change the LAN segment assignment of an end node connected to the switch through the AT-S101 software. Also, you

can change the VLAN memberships through the management software without moving the workstations physically or change group memberships without moving cables from one port to another.

In addition, a virtual LAN can span more than one switch. This means that the end nodes of a VLAN do not need to be connected to the same switch and so are not restricted to being in the same physical location.

The AT-GS950/8POE switch supports the following types of VLANs:

- Port-based VLANs
- Tagged VLANs

Both types of VLANs are described in the following sections.

Port-based VLAN Overview

As explained in the "VLAN Overview" on page 46, a VLAN consists of a group of ports on an Ethernet switch that form an independent traffic domain. Traffic generated by the end nodes of a VLAN remains within the VLAN and does not cross over to the end nodes of other VLANs unless there is an interconnection device, such as a router or Layer 3 switch.

A port-based VLAN is a group of ports on an AT-GS950/8POE switch that form a logical Ethernet segment. A port-based VLAN can have as many or as few ports as needed. The VLAN can consist of all the ports on an Ethernet switch, or just a few ports.

The parts of a port-based VLAN in the AT-S101 Management Software are:

- VLAN name
- Group ID

VLAN Name

To create a port-based VLAN, you must give it a name that reflects the function of the network devices that are VLAN members, such as Sales, Production, and Engineering.

Group ID

Each VLAN in a network must have a unique number assigned to it. This number is called the Group ID. This number uniquely identifies a VLAN in the switch.

Each port of a port-based VLAN can belong to as many VLANs as needed. Therefore, traffic can be forwarded to the members of the groups which the port is assigned to. For example, port 1 and port 2 are members of group 1 and ports 1 and 3 are members of group 2. In this case, traffic from port 1 is forwarded to ports 2 and 3, traffic from port 2 is forwarded only to port 1, and traffic from port 3 is forwarded only to port 1.

General Rules for Creating a Port-based VLAN

Below is a summary of the general rules to observe when creating a port-based VLAN.

- Assign a name to each port-based VLAN.
- Assign each port-based VLAN a Group ID.
- The AT-GS950/8POE switch can support up to 52 port-based VLANs.

Tagged VLAN Overview

The second type of VLAN supported by the AT-S101 Management Software is the *tagged VLAN*. In this type of VLAN, membership is determined by information within the frames that are received on a port and the VLAN configuration of each port.

The VLAN information within an Ethernet frame is referred to as a *tag* or *tagged header*. A tag, which follows the source and destination addresses in a frame, contains the Group ID of the VLAN to which the frame belongs (IEEE 802.3ac standard). This number uniquely identifies each VLAN in a network.

When a switch receives a frame with a VLAN tag, referred to as a *tagged frame*, the switch forwards the frame only to those ports whose Group ID equals the VLAN tag.

A port that receives or transmits tagged frames is referred to as a *tagged port*. Any network device connected to a tagged port must be IEEE 802.1Q-compliant. This is the standard that outlines the requirements and standards for tagging. The device must be able to process the tagged information on received frames and add tagged information to transmitted frames.

A tagged VLAN consists of the following:

- VLAN Name
- Group ID
- Tagged and Untagged Ports
- Port VLAN identifier (PVID)

Tagged and Untagged Ports

When you specify that a port is a member of a tagged VLAN, you need to specify that it is tagged or untagged. You can have a combination of tagged and untagged ports in the same VLAN.

Packet transmission from a tagged port differs from packet transmission from an untagged port. When a packet is transmitted from a tagged port, the tagged information within the packet is maintained when it is transmitted to the next network device. If the packet is transmitted from an untagged port, the VLAN tag information is removed from the packet before it is transmitted to the next network device.

The IEEE 802.1Q standard describes how tagging information within a packet is used to forward or discard traffic throughout the switch. If the incoming packet has a VLAN tag that matches one of the Group IDs of which the port is a member, the packet is accepted and forwarded to the appropriate port(s) within that VLAN. If the incoming packet's VLAN tag does not match one of the Group IDs assigned to the port, the packet is discarded.

Port VLAN Identifier

When an untagged packet is received on a port in a tagged VLAN, it is assigned to one of the VLANs of which that port is a member. The deciding factor in this process is the Port VLAN Identifier (PVID). Both tagged and untagged ports in a tagged VLAN must have a PVID assigned to them. The default value of the PVID for each port is 1. The switch associates a received untagged packet to the Group ID that matches the PVID assigned to the port. As a result, the packet is only forwarded to those ports that are members of that VLAN.

General Rules for Creating a Tagged VLAN

Below is a summary of the rules to observe when you create a tagged VLAN.

- Each tagged VLAN must be assigned a unique VID. If a particular VLAN spans multiple switches, each part of the VLAN on the different switches must be assigned the same VID.
- A tagged port can be a member of multiple VLANs.
- The AT-GS950/8POE switch can support up to 200 tagged VLANS.

Displaying Ports and Assigning Ports to a VLAN

By default, all of the ports on the switch are assigned to the Tagged VLAN. The procedure described in this section allows you to display the current VLAN assignment of ports. In addition, it permits you to assign ports to tagged or a port-based VLAN. However, you can assign ports to a port-based VLAN only after you have created a port-based VLAN with the procedure described in “Creating a Port-Based VLAN” on page 56.

To assign ports to a tagged or port-based VLAN, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands.

2. From the **Bridge** folder, select **VLAN**.

   The **VLAN** folder expands.

3. From the **VLAN** folder, select **VLAN Mode**.

   The VLAN Mode Page is shown in Figure 15.

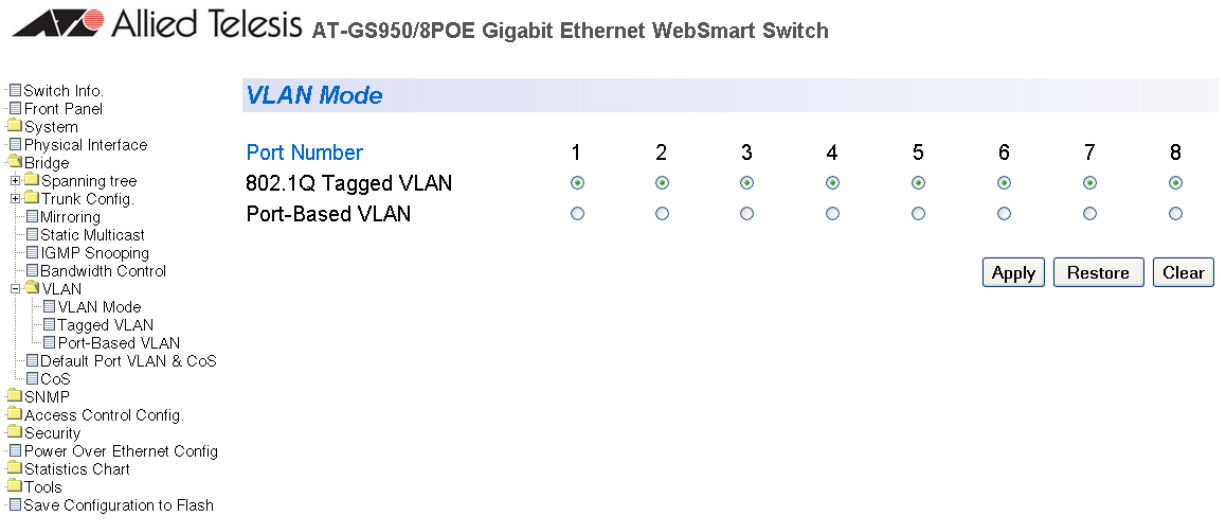

Figure 15. VLAN Mode Page

4. To add ports to a Tagged or Port-Based VLAN, select the ports and then click **Apply**.

Creating a Tagged VLAN

To create a tagged VLAN, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands.

2. From the **Bridge** folder, select **VLAN**.

   The **VLAN** folder expands.

3. From the **VLAN** folder, select **Tagged VLAN**.

   The Tagged VLAN Page is shown in Figure 16

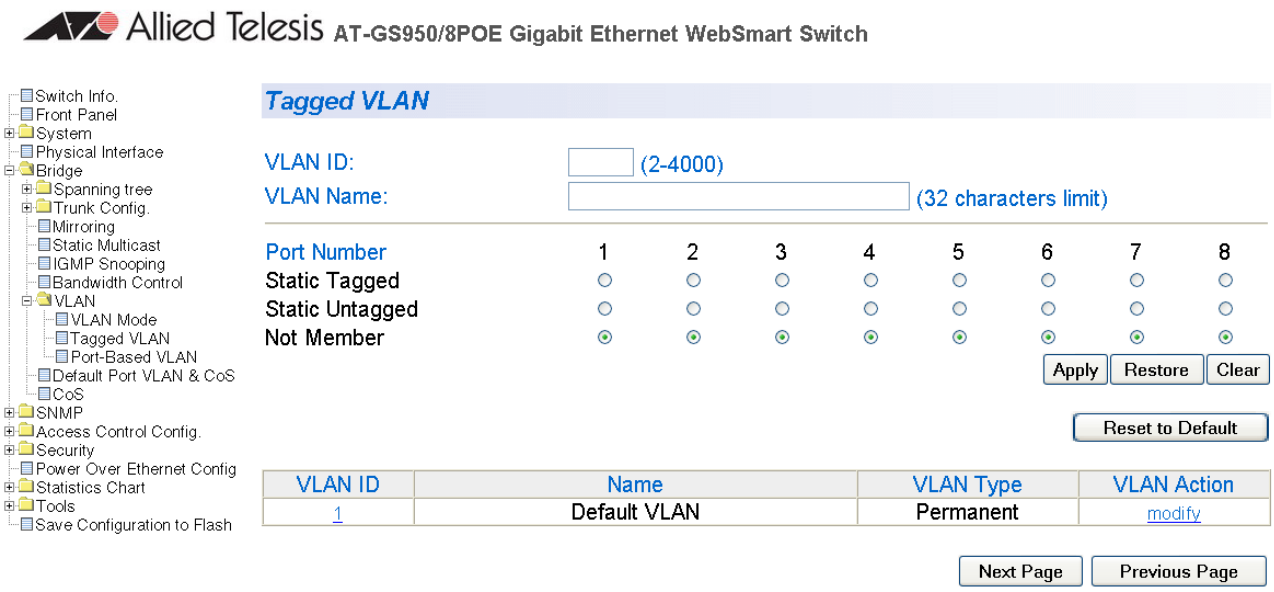

Figure 16. Tagged VLAN Page

4. To assign a VLAN ID, type in a VLAN ID in the **VLAN ID** field.

   Choose a value between 2 and 4,000. You can configure up to 200 tagged VLANs.

5. To assign a name to the VLAN, type in a name in the **VLAN Name** field.

   Enter a value of up to 32 characters.

6. To assign ports to the VLAN, click on the port numbers labeled either Static Tagged or Static Untagged. Then click **Apply**.

   By default, all the ports are assigned to the **Not Member** category.

For an example of Tagged VLANs, see Figure 17.

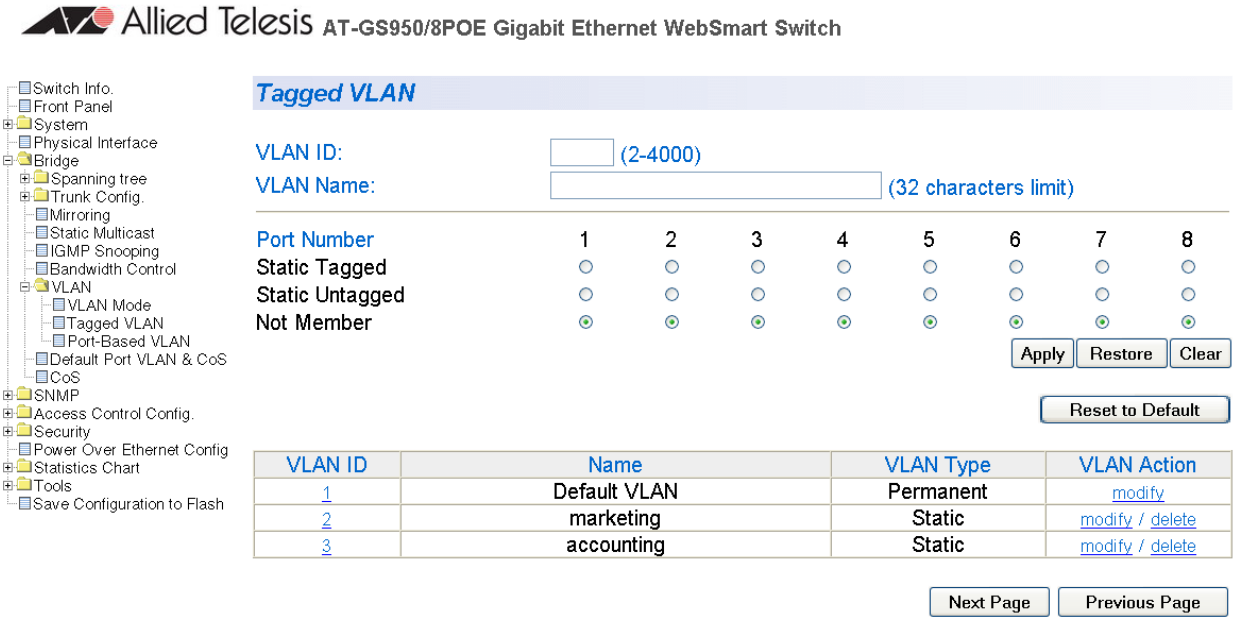

Figure 17. Example of Tagged VLAN Page

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Modifying a Tagged VLAN

To modify the name or port assignments of a tagged VLAN, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands.

2. From the **Bridge** folder, select **VLAN**.

   The **VLAN** folder expands.

3. From the **VLAN** folder, select **Tagged VLAN**.

   An Example of a Tagged VLAN page is shown in Figure 17 on page 52.

4. In the VLAN Action column, click **modify** next to the VLAN that you want to change.

   The Modify VLAN Page is displayed, see Figure 18

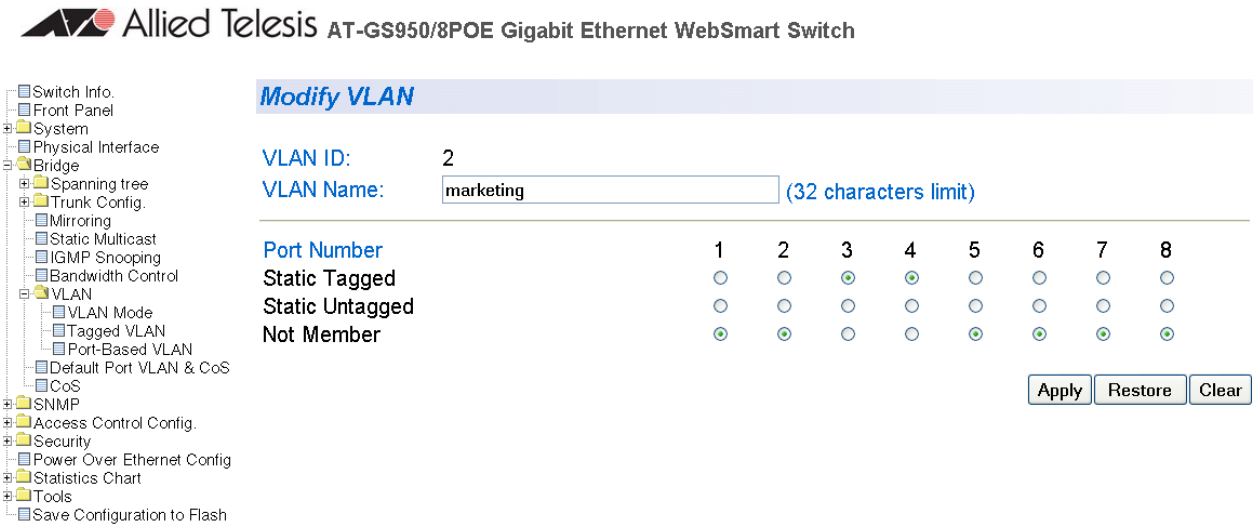

Figure 18. Modify VLAN Page

5. To change the VLAN ID, type in a VLAN ID in the **VLAN ID** field.

   You can choose a value between 2 and 4,000. You can configure up to 52 VLANs.

6. To change the name of the VLAN, type in a name in the **VLAN Name** field.

7. To assign ports to the VLAN, click on the port numbers labeled either Static Tagged or Static Untagged. Then click **Apply**.

8. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting a Tagged VLAN

To delete a tagged VLAN, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands.

2. From the **Bridge** folder, select **VLAN**.

   The **VLAN** folder expands.

3. From the **VLAN** folder, select **Tagged VLAN**.

   An example of the Tagged VLAN Page is shown in Figure 17 on page 52.

4. In the VLAN Action column, click **delete** next to the VLAN that you want to delete.

   A confirmation prompt is displayed.

5. Click **OK** to delete the VLAN or **Cancel** to cancel the deletion.

   Note
   You cannot delete the Default VLAN which has a VID of 1.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Creating a Port-Based VLAN

To create a port-based VLAN, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands.

2. From the **Bridge** folder, select **VLAN**.

   The **VLAN** folder expands.

3. From the **VLAN** folder, select **Port-Based VLAN**.

   The Port-Based VLAN Page is shown in Figure 19.

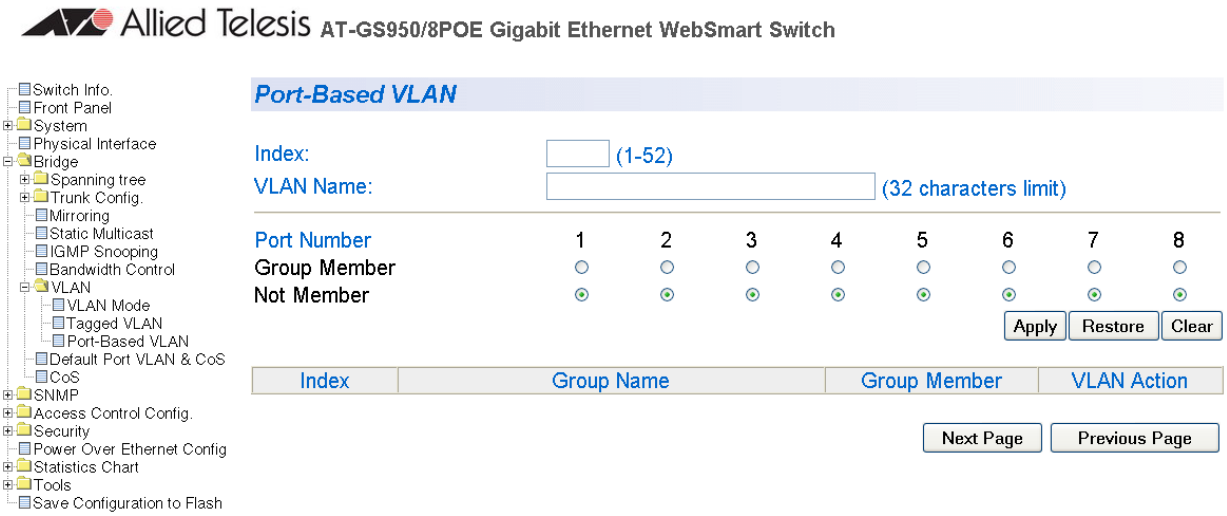

Figure 19. Port-Based VLAN Page

4. To assign a VLAN ID, type a VLAN ID in the **VLAN ID** field.

   Choose a value between 1 and 52.

5. To assign a name to a VLAN, type in a name in the **VLAN Name** field.

   Enter a value of up to 32 characters.

6. To assign ports to the VLAN, click on the port numbers labeled Group Member. Then click **Apply**.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Modifying a Port-Based VLAN

To modify the name or port assignments of a port-based VLAN, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands.

2. From the **Bridge** folder, select **VLAN**.

   The **VLAN** folder expands.

3. From the **VLAN** folder, select **Port-Based VLAN**.

   The Port-Based VLAN Page is shown in Figure 19 on page 56.

4. In the VLAN Action column, click **modify** next to the VLAN that you want to change.

   The Modify Port-based VLAN Page is shown in Figure 20.

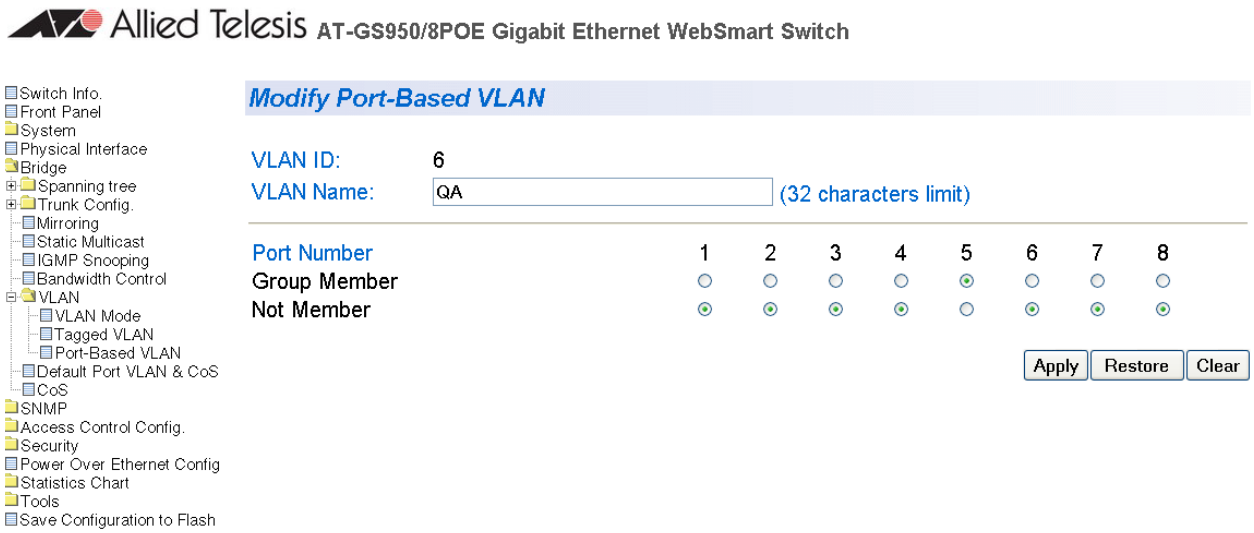

Figure 20. Modify Port-based VLAN

5. To change the name of the VLAN, type in a name in the **VLAN Name** field.

   Enter a value of up to 32 characters.

6. To assign ports to the VLAN, click on the port numbers labeled either Group Member or Not Member. Then click **Apply**.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting a Port-Based VLAN

To delete a port-based VLAN, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands.

2. From the **Bridge** folder, select **VLAN**.

   The **VLAN** folder expands.

3. From the **VLAN** folder, select **Port-Based VLAN**.

   The Port-Based VLAN Page is shown in Figure 19 on page 56.

4. In the VLAN Action column, click **delete** next to the VLAN that you want to delete.

   A confirmation prompt is displayed.

5. Click **OK** to delete the VLAN or **Cancel** to cancel the deletion.

   Note
   You cannot delete the Default VLAN which has a VID of 1.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 4

Quality of Service (QoS)

This chapter contains a description of the QoS feature and the procedures for configuring Quality of Service (QoS). This chapter includes the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

When a port on an Ethernet switch becomes oversubscribed—its egress queues contain more packets than the port can handle in a timely manner—the port may be forced to delay the transmission of some packets, which delays packets from reaching their destinations. A port may be forced to delay transmission of packets while it handles other traffic, and, in some situations, some packets destined to be forwarded to an oversubscribed port from other switch ports may be discarded.

Minor delays are often inconsequential to a switch or its performance. But there are applications, referred to as delay or time sensitive applications, that can be impacted by packet delays. Voice transmission and video conferencing are two examples. If packets carrying data for either of these are delayed from reaching their destination, the audio or video quality may suffer.

This is where the QoS feature can be of value. It allows you to manage the flow of traffic through a switch by having the switch ports give higher priority to some packets, such as delay sensitive traffic, over other packets. This is referred to as prioritizing traffic.

The QoS feature actually consists of several different elements. The element supported by the AT-GS950/8POE switch is called Class of Service (CoS) and which applies primarily to tagged packets and DSCP which priorities IP packets. As explained in “Tagged VLAN Overview” on page 105, a tagged packet contains information that specifies the VLAN to which the packet belongs.

A tagged packet can also contain a priority level. This priority level is used by switches and other networking devices to determine how important (delay sensitive) a packet is in comparison to other packets. Packets of a high priority are typically handled before packets of a low priority.

CoS, as defined in the IEEE 802.1p standard, has eight levels of traffic classes. In the AT-S101 software, the priorities are 0 to 7, with 0 the lowest priority and 7 the highest priority.

When a tagged packet is received by a port, it is examined by the AT-S101 software for its priority. The switch software uses the priority to determine which egress priority queue the packet should be stored in on the egress port.

Each port on the AT-GS950/8POE switch has four priority queues, 0 (low) to 3 (high). When a tagged packet enters a switch port, the switch responds by placing the packet into one of the queues according to the assignments shown in Table 1 on page 63. A packet in a high priority egress queue is typically transmitted from a port sooner than a packet in a low priority queue.

Table 1. Default Mappings of IEEE 802.1p Priority Levels to Egress Port Priority Queues

IEEE 802.1p Traffic Class	AT-GS950/8POE Egress Port Priority Queue
0	0
1	0
2	0
3	1
4	2
5	2
6	3
7	3

For example, a tagged packet with a priority tag of 6 is placed in the egress port's highest priority queue of 3, while a packet with a priority tag of 1 is placed in the lowest priority queue.

Note
QoS is disabled by default on the switch.

You can customize these priority-to-queue assignments using the AT-S101 Management Software. The procedure for changing the default mappings is found in "Mapping CoS Priorities to Egress Queues" on page 65. Note that because all ports must use the same priority-to-egress queue mappings, these mappings are applied at the switch level. They cannot be set on a per-port basis.

You can configure a port to ignore the priority levels in its tagged packets and use a temporary priority level assigned to the port instead. For instance, perhaps you decide that all tagged packets received by port 4 should be assigned a priority level of 5, regardless of the priority level in the packets themselves. The procedure for overriding priority levels is explained in "Configuring CoS" on page 67.

CoS relates primarily to tagged packets rather than untagged packets because untagged packets do not contain a priority level. By default, all untagged packets are placed in a port's Q0 egress queue, the queue with the lowest priority. But you can override this and instruct a port's untagged frames to be stored in a higher priority queue. The procedure for this is also explained in "Configuring CoS" on page 67.

One last thing to note is that the CoS feature does not change the priority level in a tagged packet. The packet leaves the switch with the same priority it had when it entered. This is true even if you change the default priority-to-egress queue mappings.

The default setting for the Quality of Service feature is disabled. When the feature is disabled, all tagged packets are stored in the lowest priority queue of a port.

Mapping CoS Priorities to Egress Queues

This procedure explains how to change the default mappings of CoS priorities to egress priority queues, as shown in Table 1 on page 63. This is set at the switch level. You cannot set these mappings on a per-port level. You can also use this procedure to enable and disable QoS on the switch.

To change the default mappings of CoS priorities to egress priority queues or to enable or disable the QoS feature, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands to show the **VLAN** folder.

2. From the **VLAN** folder, select **CoS**.

   The CoS Page is shown in Figure 21.

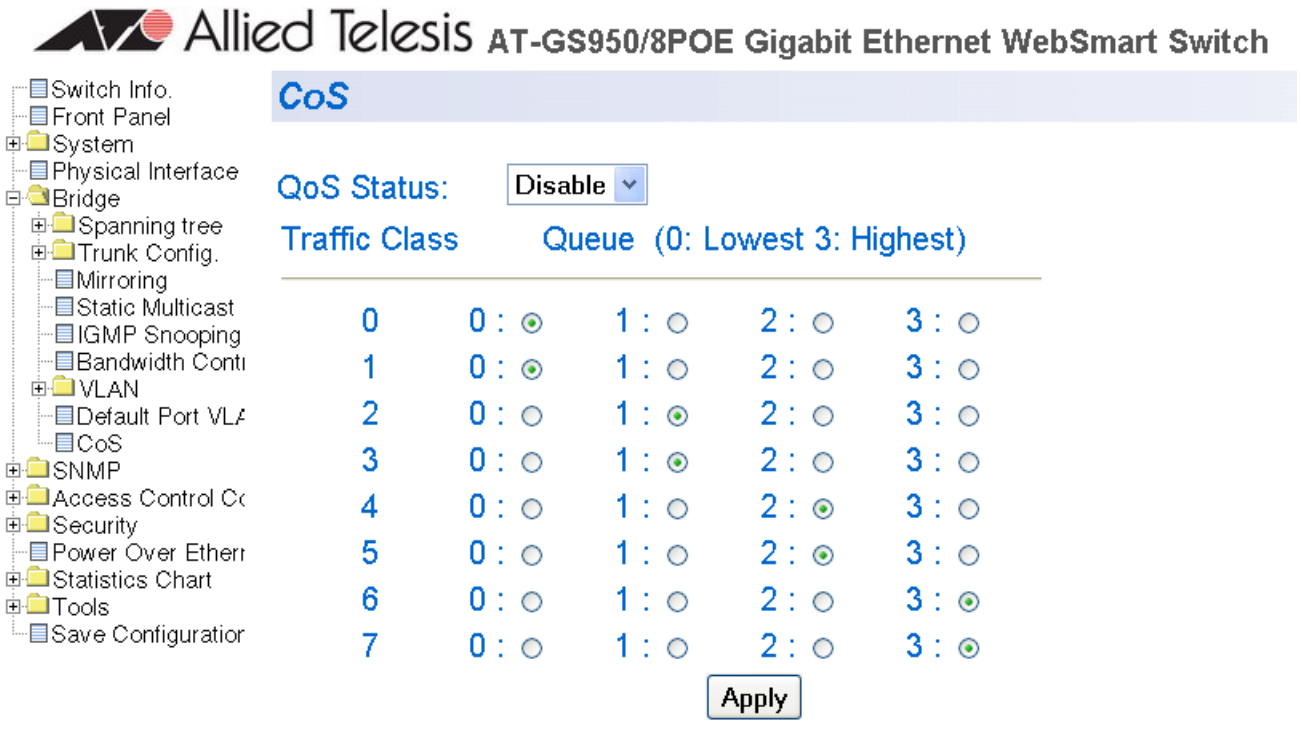

Figure 21. CoS Page

3. To enable or disable QoS, select **Enable** or **Disable** from the QoS Status pull-down menu. The default is **Disable**.

4. To change the egress priority queue assignment of an 802.1p priority class, click the dialog circle of the queue for the corresponding priority.

For example, to direct all tagged traffic with a traffic class of 4 to egress queue 3 on the ports, click the dialog circle for queue 3 in the traffic class 4 row.

5. Click **Apply**.

Note
The switch does not alter the original priority level in tagged frames. Frames leave the switch with the same priority level they had when they entered the switch.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Configuring CoS

As explained in "Overview" on page 62, a packet received by a port is placed it into one of four priority queues on the egress port according to the switch's mapping of 802.1p priority levels to egress priority queues. The default mappings are shown in Table 1 on page 63.

You can override the mappings at the port level by assigning a new default egress queue to a port. Note that this assignment is made on the ingress port before the frame is forwarded to the egress port. Consequently, you need to configure this feature on the ingress port. For example, you can configure a port so that all ingress frames are stored in egress queue 3 of the egress port, regardless of the priority levels that might be in the frames themselves, as found in tagged frames.

Note
The switch does not alter the original priority level of tagged frames. Frames leave the switch with the same priority level they had when they entered the switch.

To configure CoS for a port, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Bridge** folder expands to show the **VLAN** folder.

2. From the **VLAN** folder, select **Default Port VLAN**.

   The Default Port VLAN & CoS Page is shown in Figure 22.

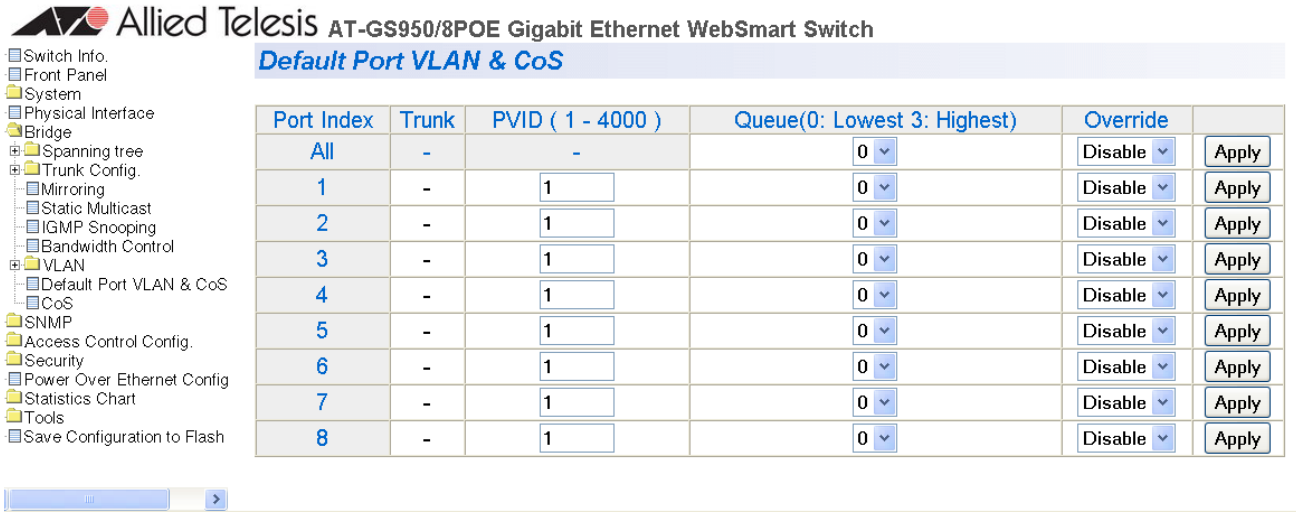

Figure 22. Default Port VLAN & CoS Page

The columns in the menu display the following information:

Port Index
Displays the port number. The All value refers to ports 1 through 8 on the switch.

Trunk
Displays the trunk number if the port is a member of a trunk.

PVID
Displays the Port VLAN identifier (PVID) of the port. For more information about this field, see “Port VLAN Identifier” on page 49.

Queue
Displays the number of the queue where untagged packets received on the port are stored on the egress queue. In this field, 0 is the lowest value and 3 is the highest value.

Override
Displays whether the priority level in ingress tagged frames is being used or not. If the parameter is set to **Disable**, the override is deactivated and the port uses the priority levels contained within the frames to determine the egress queue. If the parameter is set to **Enable**, the override is activated and the tagged packets are stored in the egress queue specified in the Queue column.

3. To change the egress queue where ingress untagged frames received on a port will be stored on the egress port, use the pull-down menu in the **Queue** column and select the desired queue. The range is 0 (lowest) to 3 (highest). The default is 0. For example, if you select 3 for queue 3 for a port, all ingress untagged packets received on the port are stored in egress queue 3 on the egress port. (If you perform Step 3 and override the priority level in ingress tagged packets, this also applies to tagged packets as well.)

   If the selected port is part of a port trunk, all ports in the trunk are automatically assigned the same egress queue.

4. To configure a tagged port so that the switch ignores the priority tag in ingress tagged frames, select **Enable** from the Override column for the corresponding port.

   The default for this parameter is **Disable**, meaning that the priority level of tagged frames is determined by the priority level specified in the frame itself.

5. Click **Apply**.

Note
The tagged information in a frame is not changed as the frame traverses the switch. A tagged frame leaves a switch with the same priority level that it had when it entered.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 5

Port Configuration

This chapter provides a description of the physical characteristics of the ports and a procedure that explains how to view and change the port settings. This chapter includes the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

This chapter describes how to display and modify the physical characteristics of an AT-GS950/8POE switch. You can display and modify the settings of all the ports on one web page. The port characteristics of displayed are:

- Port status
- Port type
- Flow control setting
- Duplex mode setting

Displaying and Configuring Ports Using the Port Configuration Page

This procedure explains how to configure the ports on the switch using the Port Configuration Page. This page allows you to view and configure the parameter settings of all the switch ports at one time.

To configure the ports, perform the following procedure:

1. From the menu on the left side of the page, select **Physical Interface** which is located underneath the **System** folder.

   The Physical Interface Page is shown in Figure 23. The page lists all the ports on the switch and their current settings.

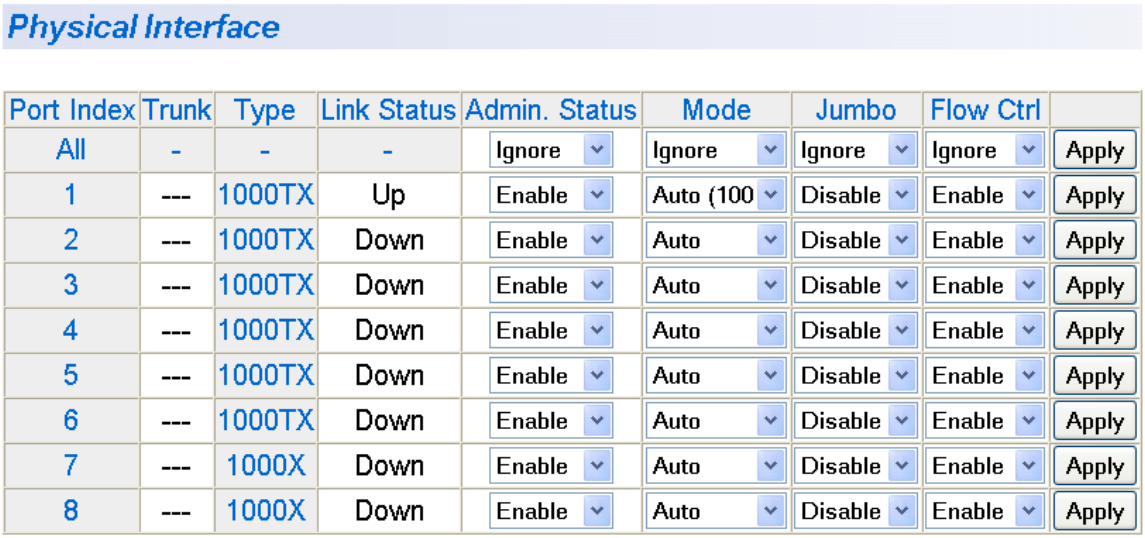

Figure 23. Physical Interface Page

2. Adjust the port settings as needed. Not all parameters are adjustable. The parameters are defined here:

   Port Index
   Specifies the port number. The **All** value indicates ports 1 through 8. You cannot change this parameter.

   Note
   You can use the All value to set the Admin. Status, Mode, Jumbo, and Flow Ctrl fields to the same values on all eight ports.

   Trunk
   Indicates the trunk group number. A number in this column indicates

that the port has been added to a trunk. For information about configuring a trunk, refer to Chapter 6, “Port Trunking” on page 77.

Type
Indicates the port type. The port type is 1000TX for 10/100/1000Base-T twisted-pair ports and 1000X for the optional SFP fiber ports.

Link Status
Indicates the status of the link between the port and the end node connected to the port. The possible values are:

Up - Indicates a valid link exists between the port and the end node.

Down - Indicates the port and the end node have not established a valid link.

Admin. Status
Indicates the operating status of the port.

You can use this parameter to enable or disable a port. You may want to disable a port and prevent packets from being forwarded if a problem occurs with the node or cable connected to the port. You can enable the port to resume normal operation after the problem has been fixed. You can also disable an unused port to secure it from unauthorized connections. The possible values are:

Ignore—Indicates the All setting does not apply to the Admin. Status field. In other words, each port is set individually.

Enabled—The port is able to send and receive Ethernet frames. This is the default setting for a port.

Disabled—The port is disabled.

Mode
Indicates the speed and duplex mode settings for the port.

You can use this parameter to set the speed and duplex mode of a port. Possible settings are:

Ignore—Indicates the All setting does not apply to the Mode field. In other words, each port is set individually.

Auto - The port is using Auto-Negotiation to set the operating speed and duplex mode. This is the default setting for all ports. The actual operating speed and duplex mode of the port are displayed in parentheses (for example, “1000F” for 1000 Mbps full duplex mode) after a port establishes a link with an end node.

Auto (100F) - 1000 Mbps in half-duplex mode

1000/Full - 1000 Mbps in full-duplex mode

100/Full - 100 Mbps in full-duplex mode

10/Full - 10 Mbps in full-duplex mode

100/Half - 100 Mbps in half-duplex mode

10/Half - 10 Mbps in half-duplex mode

When selecting a setting, note the following:

- When a twisted-pair port is set to Auto-Negotiation, the default setting, the end node should also be set to Auto-Negotiation to prevent a duplex mode mismatch. A switch port using Auto-Negotiation defaults to half-duplex if it detects that the end node is not using Auto-Negotiation. This can result in a mismatch if the end node is operating at a fixed duplex mode of full-duplex. To avoid this problem when connecting an end node with a fixed duplex mode of full-duplex to a switch port, disable Auto-Negotiation on the port and set the port's speed and duplex mode manually.
- Allied Telesis does not recommend manually setting a 10/100/1000Base-T twisted-pair port to either 1000 Mbps full duplex or 1000 Mbps half duplex. For 1000 Mbps operation, Allied Telesis recommends setting the port to Auto-Negotiation.
- The only valid setting for an optional SFP port is Auto-Negotiation.

Jumbo
Indicates whether or not jumbo frames can be accepted by the switch. You may want to activate jumbo frames when your switch will transmit video and audio files. The possible values are:

Ignore—Indicates the All setting does not apply to the Jumbo field. In other words, each port is set individually.

Enabled —The port is permitted to accept jumbo frames.

Disabled—The port is not permitted to accept jumbo frames. This is the default setting for all ports on the switch.

Flow Control
The current flow control setting on the port. The switch uses a special pause packet to notify the end node to stop transmitting for a specified period of time. The possible values are:

Ignore—Indicates the All setting does not apply to the Flow Control field. In other words, each port is set individually.

Enabled—The port is permitted to use flow control. This is the default setting for all ports on the switch.

Disabled—The port is not permitted to use flow control.

3. Click **Apply** to save the configuration.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 6

Port Trunking

This chapter contains the following procedures for working with port trunking:

Note
For information about Link Aggregation Control Protocol (LACP) port trunking, see Chapter 7, “LACP Port Trunks” on page 85.

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Port Trunking Overview

A port trunk is an economical way for you to increase the bandwidth between the Ethernet switch and another networking device, such as a network server, router, workstation, or another Ethernet switch. A port trunk is a group of ports that have been grouped together to function as one logical path. A port trunk increases the bandwidth between the switch and another network device and is useful in situations where a single physical link between the devices is insufficient to handle the traffic load.

Static Port Trunk Overview

A static port trunk consists of two to eight ports on the switch that function as a single virtual link between the switch and another device. A static port trunk improves performance by distributing the traffic across multiple ports between the devices and enhances reliability by reducing the reliance on a single physical link.

A static trunk is easy to configure. You designate the ports on the switch that are in the trunk and the management software on the switch automatically groups them together.

The example in Figure 24 illustrates a static port trunk of four links between two AT-GS950/8POE Gigabit Ethernet WebSmart Switches.

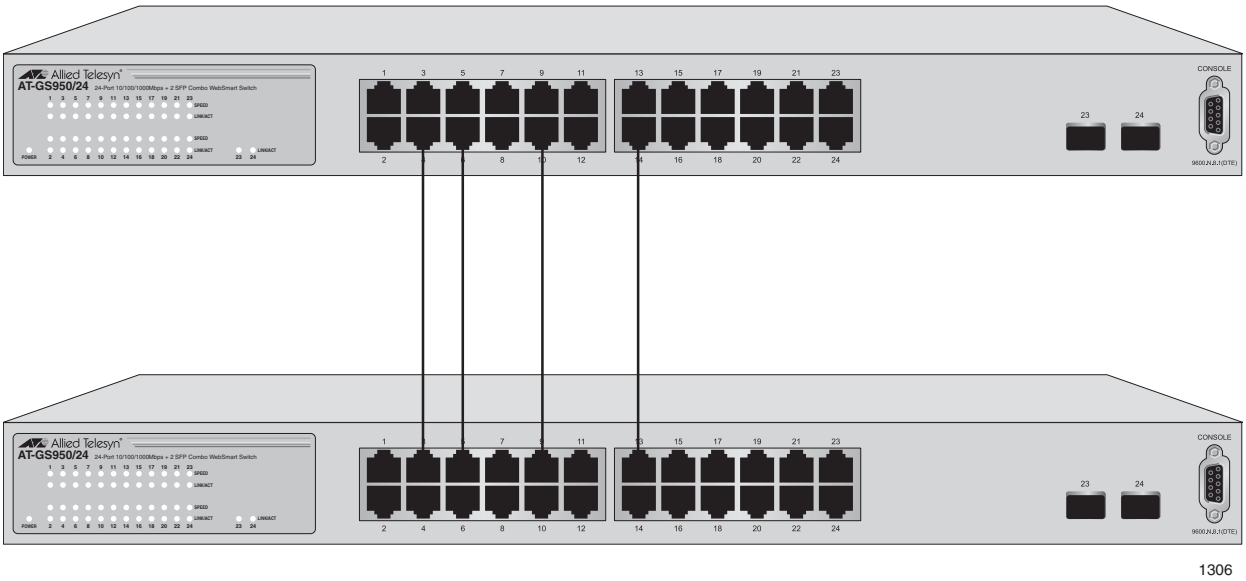

Figure 24. Static Port Trunk Example

Network equipment vendors tend to employ different techniques to implement static trunks. Consequently, a static trunk on one device may be incompatible with the same feature on a device from a different

manufacturer. For this reason static trunks are typically employed only between devices from the same vendor. That is not to say that an Allied Telesis Layer 2 managed switch cannot form a static trunk with a device from another manufacturer; however, there is the possibility that the implementations of static trunking on the two devices may be incompatible.

Also, note that a static trunk does not provide for redundancy or link backup. If a port in a static trunk loses its link, the trunk's total bandwidth is diminished. Although the traffic carried by the lost link is shifted to one of the remaining ports in the trunk, the bandwidth remains reduced until the lost link is reestablished or you reconfigure the trunk by adding another port to it.

Static Port Trunk Guidelines

Following are the guidelines for creating a static trunk:

- Allied Telesis recommends setting static port trunks between Allied Telesis networking devices to ensure compatibility. While an Allied Telesis device may be able to form a static trunk with a device from another equipment vendor, it is possible that the implementation of this feature on the two devices may be incompatible, resulting in undesired switch behavior.
- A static trunk can contain up to eight ports.
- The ports of a static trunk must be of the same medium type. They can be all twisted-pair ports or all fiber optic ports.
- The ports of a trunk can be either consecutive (for example, Ports 2 through 4) or nonconsecutive (for example, ports 3, 5, and 7).
- Before creating a port trunk, examine the speed, duplex mode, flow control, and back pressure settings of all of the ports included in the trunk. Verify that the settings are the same for all ports in the trunk. If these settings are not the same, then the switch does not allow you to create the trunk.
- After you have created a port trunk, a change to the speed, duplex mode, flow control, or back pressure of any port in the trunk automatically implements the same change on all the other member ports.
- A port can belong to only one static trunk at a time.
- The ports of a static trunk can be untagged or untagged members of the same VLAN.

The switch selects a port in the trunk to handle broadcast packets and packets of unknown destination. The switch makes this choice based on a hash algorithm, depending upon the source and destination MAC addresses.

Creating a Port Trunk

This procedure explains how to create a port trunk.

Caution
Do not connect the cables of a port trunk to the ports on the switch until you have configured the ports on both the switch and the end node. Connecting the cables prior to configuring the ports can create loops in your network topology. Loops can result in broadcast storms which can adversely affect the operation of your network.

To create a port trunk, perform the following procedure:

1. Select the **Bridge** folder.

   The **Bridge** folder expands.

2. From the **Bridge** folder**,** select the **Trunk Config.** folder.

   The **Trunk Config.** folder expands.

3. From the **Trunk Config.** folder, select **Trunking**.

   The Trunking Page is displayed. See Figure 25.

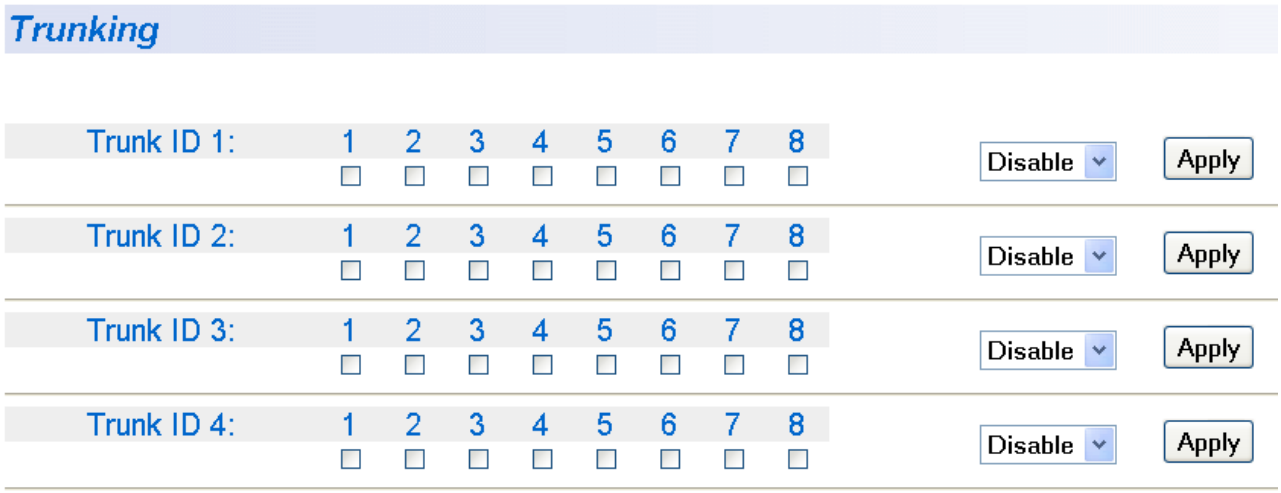

Figure 25. Trunking Page

If the switch does not contain a port trunk, all of the ports on the switch are unchecked. If there is a port trunk, the ports in the trunk are checked.

4. Click the dialog boxes of the ports that will make up the port trunk.

   A check in a box indicates the port is a member of the trunk. No check means the port is not a member. A port trunk can contain up to eight ports.

5. Change the status of the trunk from **Disable** to one of the following:

   Active
   The aggregator will broadcast and respond to LACPDU (LACP Data Unit) packets.

   Passive
   The aggregator will not broadcast LACPDU packets, but it will respond to them.

   Manual
   Enables static port trunking and disables LACP.

6. Click **Apply**.

   The trunk is now operational on the switch.

7. Configure the port trunk on the other switch and connect the cables.

8. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Modifying a Port Trunk

This procedure explains how to change the status of a port trunk and add or remove ports from a port trunk.

Caution
Before you modify a port trunk, disconnect the cables from the ports of the trunk. Adding or removing ports from a trunk without first disconnecting the cables can create loops in your network topology, which can cause broadcast storms and poor network performance.

To add or remove ports from a trunk, perform the following procedure:

1. Select the **Bridge** folder.

   The **Bridge** folder expands.

2. From the **Bridge** folder**,** select the **Trunk Config.** folder.

   The **Trunk Config.** folder expands.

3. From the **Trunk Config.** folder, select **Trunking**.

   The Trunking Page is shown in Figure 25 on page 80.

4. Click the status of the port trunk you want to modify and change the status to one of the following options:

   Disable
   Disable the port trunk.

   Active
   The aggregator will broadcast and respond to LACPDU packets.

   Passive
   The aggregator will not broadcast LACPDU packets, but it will respond to them.

   Manual
   Enables static port trunking and disables LACP.

5. Click **Apply**.

6. To add or remove a port from a trunk, click the dialog box for the port in the corresponding trunk row.

   A check in a box indicates the port is a member of the trunk. No check means the port is not a member. A port trunk can contain up to eight ports.

7. Click **Apply**.
8. Modify the port trunk on the other switch and reconnect the cables.
9. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Disabling a Port Trunk

This procedure explains how to disable a port trunk.

Note
Before you disable a port trunk, disconnect all of the cables from the ports of the trunk. Leaving the cables connected can create loops in your network topology because the ports of a disabled port trunk function as normal network ports, forwarding individual network traffic.

To enable or disable a port trunk, perform the following procedure:

1. Select the **Bridge** folder.

   The **Bridge** folder expands.

2. From the **Bridge** folder**,** select the **Trunk Config.** folder.

   The **Trunk Config.** folder expands.

3. From the **Trunk Config.** folder, select **Trunking**.

   The Trunking Page is shown in Figure 25 on page 80.

4. To disable a port trunk, select **Disable** from the pull-down menu next to the trunk that you want to disable.

5. Then click **Apply**.

6. Modify the port trunk on the other switch and disconnect the cables.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 7

LACP Port Trunks

This chapter contains overview information about LACP port trunks and the procedures for setting this feature. This chapter contains the following sections:

Note
For information about port trunking, see Chapter 6, “Port Trunking” on page 77.

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

LACP Overview

LACP (Link Aggregation Control Protocol) port trunks perform the same function as static trunks. They increase the bandwidth between network devices by distributing the traffic load over multiple physical links. The advantage of an LACP trunk over a static port trunk is its flexibility. While implementations of static trunking tend to be vendor specific, the AT-S101 software implementation of LACP is compliant with the IEEE 802.3ad standard, making it interoperable with equipment from other vendors that also comply with the standard. Therefore, you can create an LACP trunk between an Allied Telesis device and network devices from other manufacturers.

Another advantage is that ports in an LACP trunk can function in a standby mode. This adds redundancy and resiliency to the trunk. If a link in a static trunk goes down, the overall bandwidth of the trunk is reduced until the link is reestablished or another port is added to the trunk. In contrast, an LACP trunk can automatically activate ports in a standby mode when an active link fails so that the maximum possible bandwidth of the trunk is maintained.

For example, assume you create an LACP trunk of ports 11 to 20 on a switch and the switch is using ports 11 to 18 as the active ports and ports 19 and 20 as reserve. If an active port loses its link, the switch automatically activates one of the reserve ports to maintain maximum bandwidth of the trunk.

The main component of an LACP trunk is an *aggregator* which is a group of ports on the switch. The ports in an aggregator are further grouped into one or more trunks, referred to as *aggregate trunks*.

An aggregate trunk can consist of any number of ports on a switch, but only a maximum of eight ports can be active at a time. If an aggregate trunk contains more ports than can be active at once, the extra ports are placed in a standby mode. Ports in the standby mode do not pass network traffic, but they do transmit and accept LACP data unit (LACPDU) packets, which the switch uses to search for LACP-compliant devices.

Only ports on a switch that are part of an aggregator transmit LACPDU packets. If a switch port that is part of an aggregator does not receive LACPDU packets from its corresponding port on the other device, it assumes that the other port is not part of an LACP trunk. Instead, it functions as a normal Ethernet port by forwarding network traffic. However, it does continue to send LACPDU packets. If it begins to receive LACPDU packets, it automatically transitions to an active or standby mode as part of an aggregate trunk.

If there will be more than one aggregate trunk on a switch, each trunk may require a separate aggregator or it may be possible to combine them into a common aggregator. The determining factor is whether the trunks are going to the same device or different devices. If the trunks are going to the same device, each must have its own aggregator. If they are going to different devices, the trunks can be members of a common aggregator. In the latter situation, the switch differentiates the individual aggregate trunks.

Here are two examples. Figure 26 illustrates the AT-GS950/8POE switch with two LACP trunks, each containing three links. Because both aggregate trunks go to the same 802.3ad-compliant device, in this case another Gigabit Ethernet switch, each trunk requires a separate aggregator.

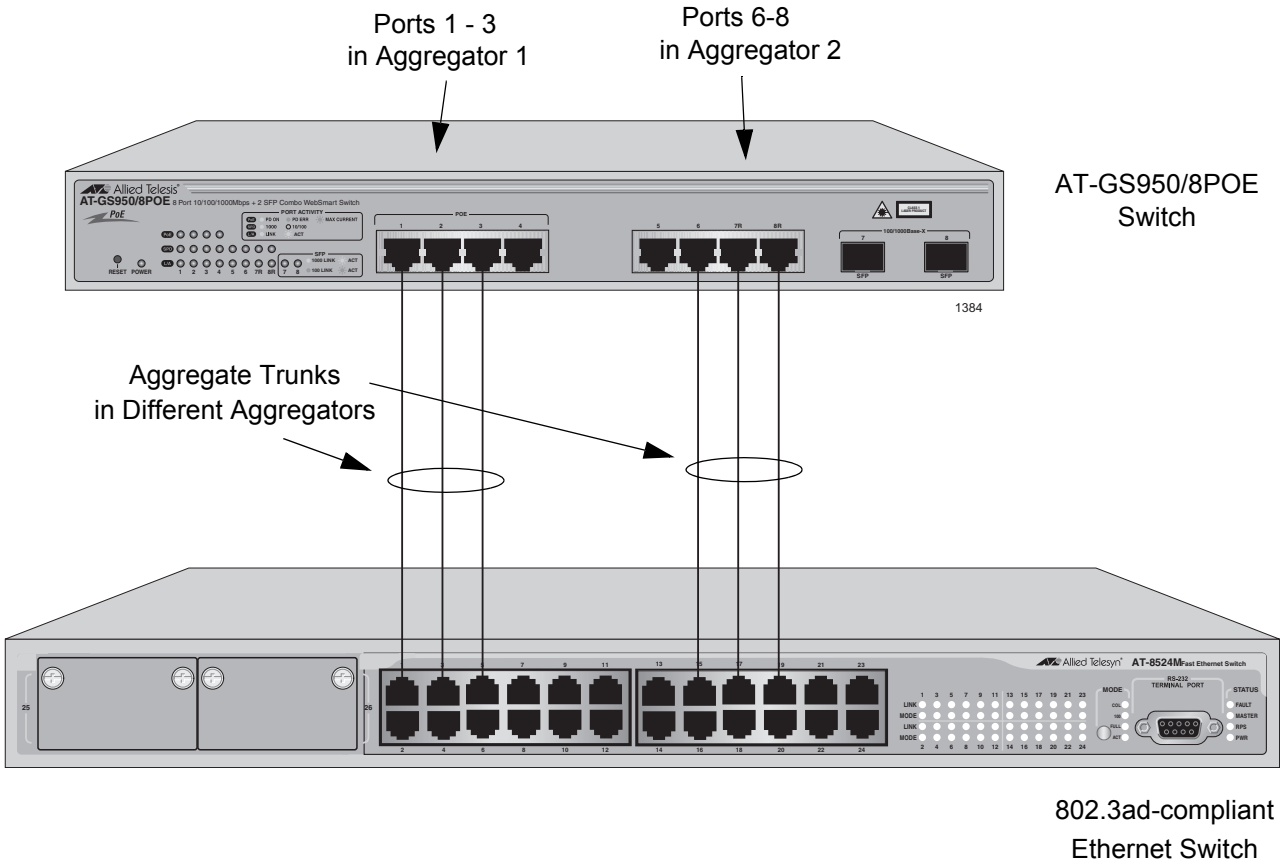

Figure 26. Example of Multiple Aggregators for Multiple Aggregate Trunks

Here is how the example looks in a table format.

Aggregator Description	Aggregator Ports	Aggregate Trunk Ports
Aggregator 1	1-3	1-3
Aggregator 2	6-8	6-8

Caution

The example cited here illustrates a loop in a network. Avoid network loops to prevent broadcast storms.

If the aggregate trunks go to different devices, you can create one aggregator and the AT-GS950/8POE switch forms the trunks automatically. This is illustrated in Figure 27 where the ports of two aggregate trunks on the AT-GS950/8POE switch are members of the same aggregator. It is the switch that determines that there are two separate aggregate trunks.

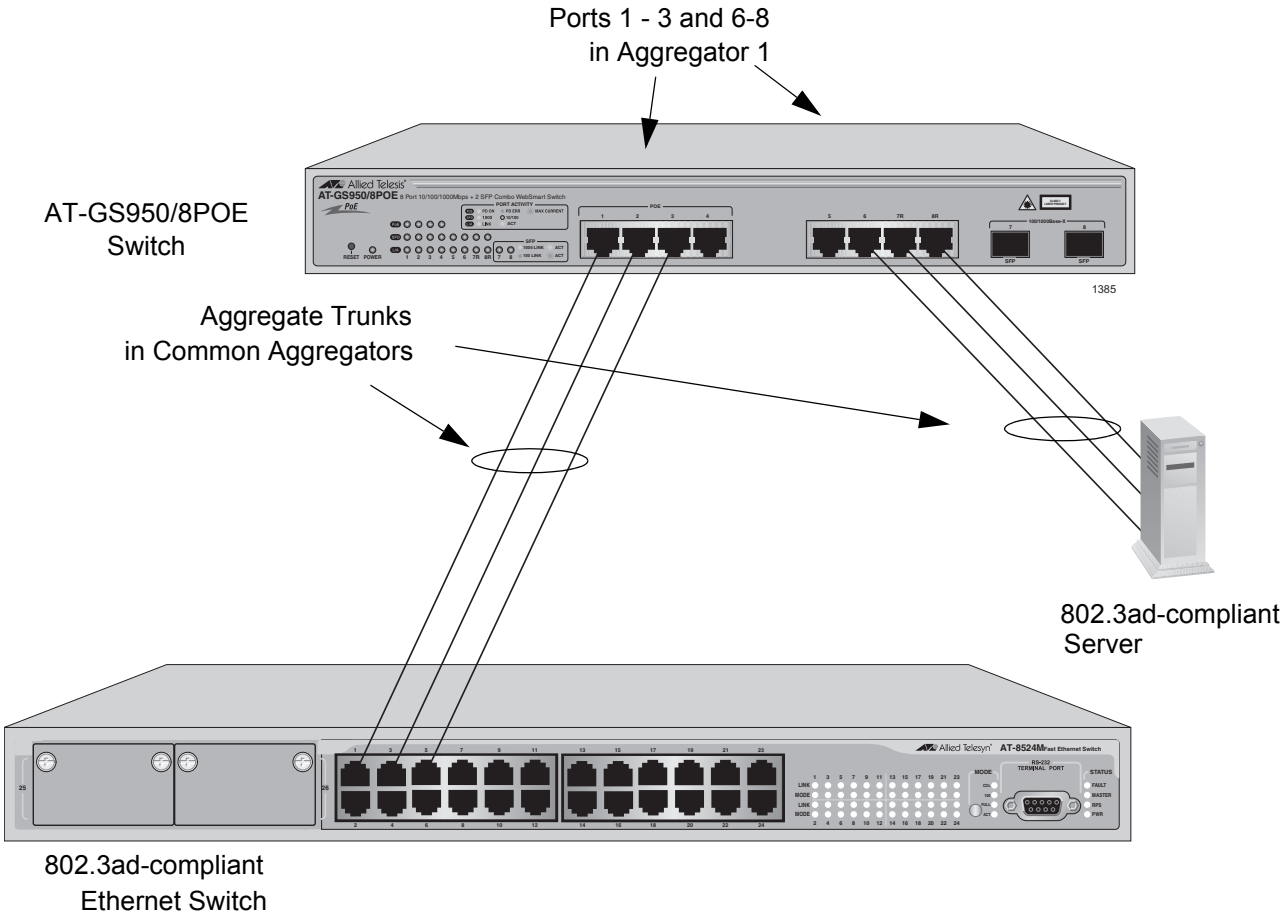

Figure 27. Example of an Aggregator with Multiple Trunks

Here is how this example looks in table format.

Aggregator Description	Aggregator Ports	Aggregate Trunk Ports
Aggregator 1	1-3, 6-8	1-3
		6-8

You could, if you wanted, create separate aggregators for the different aggregate trunks in the example above. But letting the switch make the determination for you whenever possible saves time later if you physically reassign ports to a different trunk connected to another device.

LACP System Priority

It is possible for two devices interconnected by an aggregate trunk to encounter a conflict when they form the trunk. For example, the two devices might not support the same number of active ports in an aggregate trunk or might not agree on which ports are active and which are in standby mode.

If a conflict does occur, the two devices need a mechanism for resolving the problem and deciding whose LACP settings take precedence. This is the function of the system LACP priority value. A hexadecimal value of from 1 to FFFF, this value is used whenever the devices encounter a conflict creating a trunk—the lower the number, the higher the priority. As a result, the settings on the device with the higher priority take precedence over the settings on the other device. If both devices have the same system LACP priority value, the settings on the switch with the lowest MAC address take precedence. In the AT-S101 software, the MAC address is called the System ID.

Key Parameter

The *key* parameter is a hexadecimal value from 1 to FFFF that identifies an aggregator. Each aggregator on a switch must have a unique key parameter value. The key is restricted to a switch. Two aggregators on different switches can have the same key without generating a conflict.

LACP Port Priority Value

The switch uses a port's LACP priority to determine which ports are active and which in the standby mode in situations where the number of ports in the aggregate trunk exceeds the highest allowed number of active ports. This parameter is a hexadecimal value in a range of 1 to FFFF, based on the port number. For instance, the priority values for ports 2 and 11 are 0002 and 000B, respectively. The lower the number, the higher the priority. Ports with the highest priorities are designated as the active ports in an aggregate trunk.

For example, if both 802.3ad-compliant devices support up to eight active ports and there are a total of ten ports in the trunk, the eight ports with the highest priorities (lowest priority values) are designated as the active ports, and the others are placed in the standby mode. If an active link goes down on a active port, the standby port with the next highest priority is automatically activated to take its place.

The selection of the active links in an aggregate trunk is dynamic and changes as links are added, removed, lost, or reestablished. For example, if an active port loses its link and is replaced by another port in the standby mode, the reestablishment of the link on the originally active port causes the port to return to the active state by virtue of having a higher priority value than the replacement port, which returns to the standby mode.

Two conditions must be met for a port in an aggregate trunk to function in the standby mode. First, the number of ports in the trunk must exceed the highest allowed number of active ports and, second, the port must be receiving LACPDU packets from the other device. A port functioning in the standby mode does not forward network traffic. However, it continues to send LACPDU packets. If a port that is part of an aggregator does not receive LACPDU packets, it functions as a normal Ethernet port and forwards network packets along with LACPDU packets.

Note
You can adjust the value of a port's priority.

Guidelines

The following guidelines apply to creating aggregators:

- LACP must be activated on both the switch and the other device.
- The other device must be 802.3ad-compliant.
- An aggregator can consist of any number of ports.
- The AT-S101 Management Software supports up to four active ports in an aggregate trunk at a time.
- The AT-GS950/8POE switch can support up to six static and LACP aggregate trunks at a time (for example, four static trunks and two LACP trunks). An LACP trunk is countered against the maximum number of trunks only when it is active.
- The ports of an aggregate trunk must be the same medium type: all twisted pair ports or all fiber optic ports.
- The ports of a trunk can be consecutive (for example ports 1-5) or nonconsecutive (for example, ports 2, 4, 6, 8).
- A port can belong to only one aggregator at a time.
- A port cannot be a member of an aggregator and a static trunk at the same time.
- The ports of an aggregate trunk must be untagged members of the same VLAN.
- 10/100/1000Base-TX twisted pair ports must be set to Auto-Negotiation or 100 Mbps, full-duplex mode. LACP trunking is not supported in half-duplex mode.
- 100Base-FX fiber optic ports must be set to full-duplex mode.
- You can create an aggregate trunk of transceivers with 1000Base-X fiber optic ports.
- Only those ports that are members of an aggregator transmit LACPDU packets.
- A member port of an aggregator functions as part of an aggregate trunk only if it receives LACPDU packets from the remote device. If it does not receive LACPDU packets, it functions as a regular Ethernet port, forwarding network traffic while also continuing to transmit LACPDU packets.
- The port with the highest priority in an aggregate trunk carries broadcast packets and packets with an unknown destination.
- Prior to creating an aggregate trunk between an Allied Telesis device and another vendor's device, refer to the vendor's documentation to determine the maximum number of active ports the device can support in a trunk. If the number is less than four, the maximum number for the AT-GS950/8POE switch, you should assign the other vendor's device

a higher system LACP priority than the AT-GS950/8POE switch. This can help avoid a conflict between the devices if some ports are placed in the standby mode when the devices create the trunk. For background information, refer to "LACP System Priority" on page 90.

- LACPDU packets are transmitted as untagged packets.

Displaying LACP Group Status

To display the LACP Group Status, perform the following procedure:

1. Select the **Bridge** folder.

   The **Bridge** folder expands.

2. From the **Bridge** folder**,** select the **Trunk Config.** folder.

   The **Trunk Config.** folder expands.

3. From the **Trunk Config.** folder, select **LACP Group Status**.

   The LACP Group Status Page is displayed. See Figure 28 for an example of the default display.

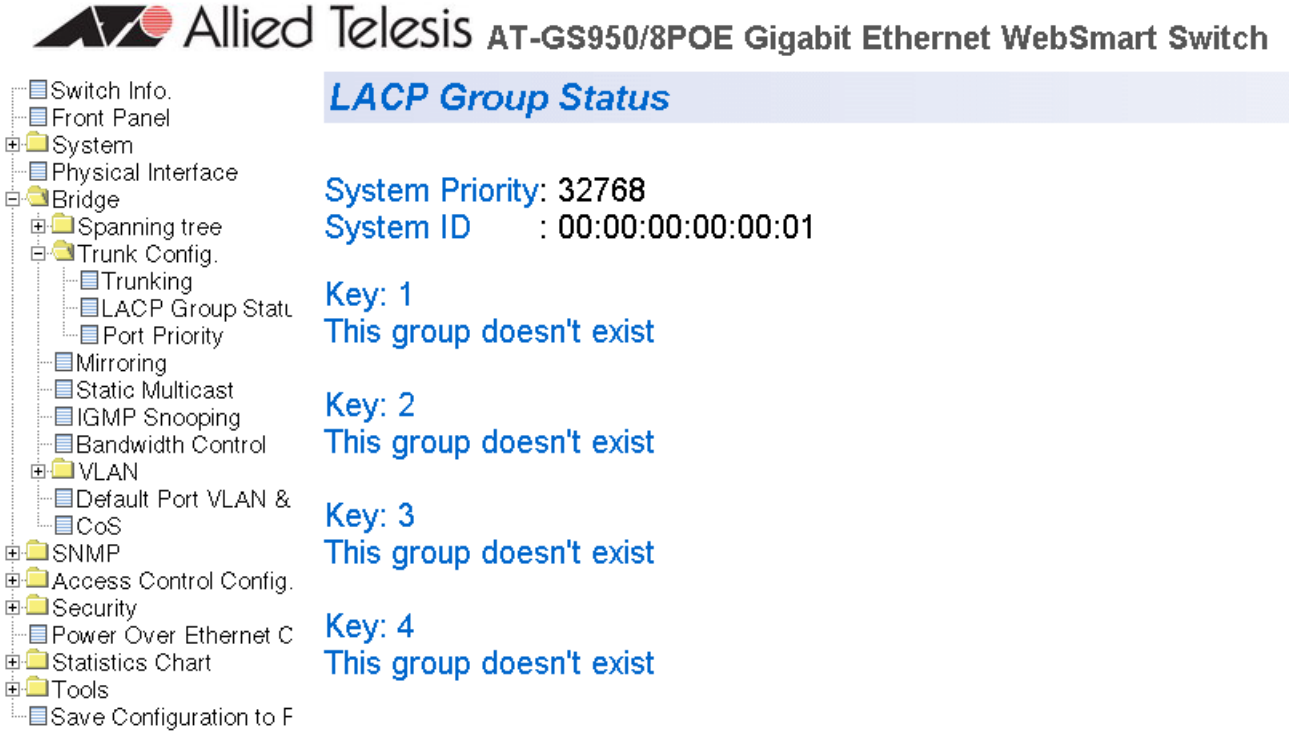

Figure 28. LACP Group Status Page

Note
You cannot change the parameters on this page:

The **System Priority** is a preassigned value that you cannot alter. This value applies to the switch. See “LACP System Priority” on page 90.

The **System ID** is a MAC address value assigned to the switch. You cannot change this value.

Key 1 - Key 4
Indicates the ID number of the trunk (aggregation group). See "Key Parameter" on page 90 for more information.

4. If you use the Trunking Page to configure port trunk ID 1, the LACP Group Status Page is updated. An example of these updates is shown in Figure 29.

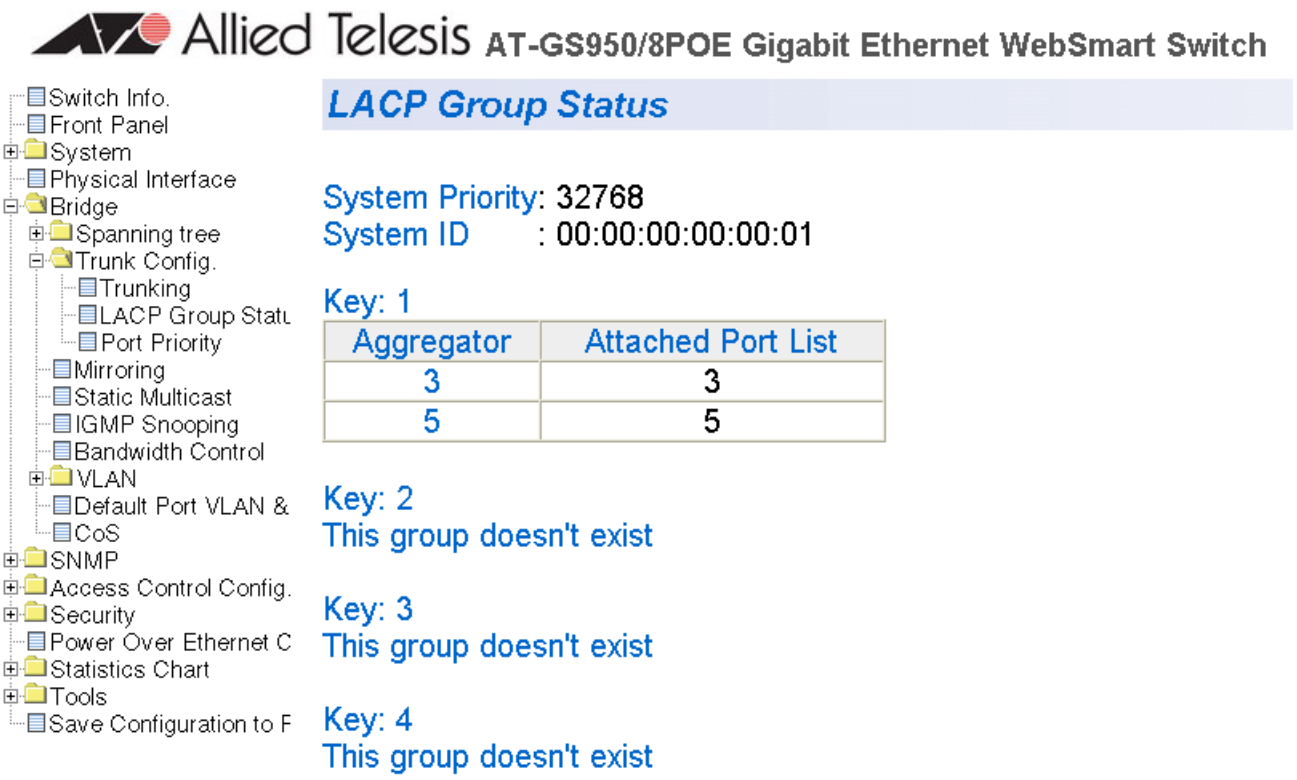

Figure 29. LACP Group Status Page with Key 1

Selecting Port Priority

To select port priority, perform the following procedure:

1. Select the **Bridge** folder.

   The **Bridge** folder expands.

2. From the **Bridge** folder**,** select the **Trunk Config.** folder.

   The **Trunk Config.** folder expands.

3. From the **Trunk Config.** folder, select **Port Priority**.

   The Port Priority Page is displayed. See Figure 30.

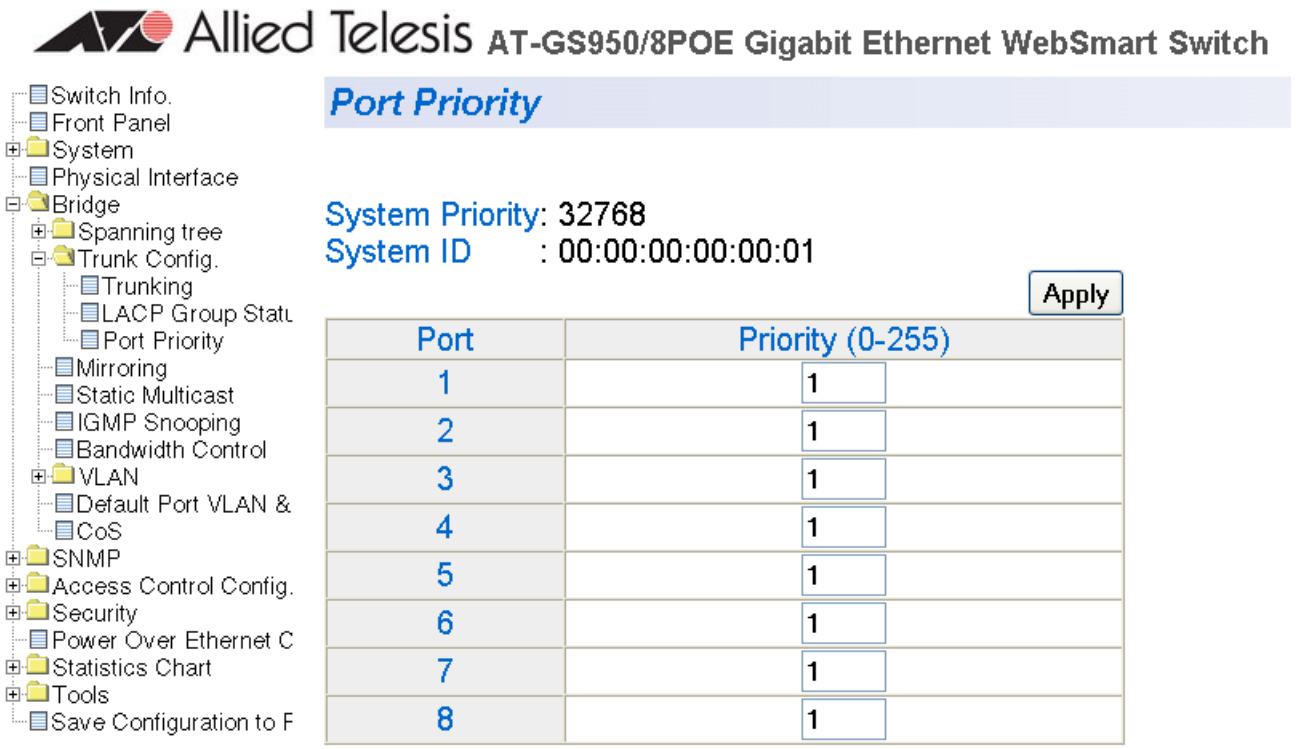

Figure 30. Port Priority Page

The **System Priority** is a preassigned value that you cannot alter. This value applies to the switch. See “LACP System Priority” on page 90.

The **System ID** is a MAC address value assigned to the switch. You cannot change this value.

4. To set the port priority, select a value from 0 to 255 in the Priority column for the port you want to alter.

   For more information, see "LACP Port Priority Value" on page 90

5. Select **Apply**.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 8

Simple Network Management Protocol (SNMP)

This chapter contains a description of SNMP and procedures for working with this protocol. This chapter contains the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

SNMP Overview

You can manage a switch by viewing and changing the management information base (MIB) objects on the device with the Simple Network Management Program (SNMP). The AT-S101 Management Software supports SNMPv1 and SNMPv2c.

To manage a switch using an SNMP application program, you must do the following:

- Activate SNMP management on the switch. The default setting for SNMP management is disabled.
- Load the Allied Telesis MIBs for the switch onto your management workstation containing the SNMP application program. The MIBs are available from the Allied Telesis web site at www.alliedtelesis.com.

To manage a switch using SNMP, you need the IP address of the switch and at least one of the switch's community strings.

Traps

A trap is a message sent by the switch to a management workstation or server to signal an operating event, such as when the device is reset.

An authentication failure trap is similar to other the traps. It too signals an operating event on the switch. But this trap is somewhat special because it relates to SNMP management. A switch that sends this trap could be indicating an attempt by someone to gain unauthorized management access using an SNMP application program to the switch. There are two events that can cause a switch to send this trap:

- An SNMP management station attempts to access the switch using an incorrect or invalid community name.
- An SNMP management station tried to access a closed access community string, to which its IP address is not assigned.

Given the importance of this trap to the protection of your switch, the management software allows you to disable and enable it separately from the other traps. If you enable it, the switch sends this trap if either of the above events occur. If you disable it, the switch does not send this trap. The default is disabled.

If you enable this trap, be sure to add one or more IP addresses of trap receivers to the community strings so that the switch will know where to send the trap if it needs to.

Community String Attributes

A community string has attributes for controlling who can use the string and what the string will allow a network management to do on the switch. The community string attributes are defined below.

Community String Name

A community string must have a name of one to eight alphanumeric characters. Spaces are allowed.

Access Mode

This attribute defines the permissions of a community string. There are two access modes: Read and Read/Write. A community string with an access mode of Read can only be used to view but not change the MIB objects on a switch. A community string with a Read/Write access can be used to both view the MIB objects and change them.

Operating Status

A community string can be enabled or disabled. When disabled, no one can use it to access the switch. You might disable a community string if you suspect someone is using it for unauthorized access to the device. When a community string is enabled, then it is available for use.

Open or Closed Access Status

This feature controls which management stations on your network can use a community string. An open access status permits any network manager who knows the community string to use it. A closed access status restricts the string to those network managers who work at particular workstations, identified by their IP addresses. You specify the workstations by assigning the IP addresses of the workstations to the community string. A closed community string can have up to eight IP addresses of management workstations.

If you decide to activate SNMP management on the switch, it is a good idea to assign a closed status to all community strings that have a Read/Write access mode and then assign the IP addresses of your management workstations to those strings. This helps reduce the chance of someone gaining management access to a switch through a community string and making unauthorized configuration changes.

Trap Receivers

A trap is a signal sent to one or more management workstations by the switch to indicate the occurrence of a particular operating event on the device. There are numerous operating events that can trigger a trap. For instance, resetting the switch or the failure of a cooling fan are two examples of occurrences that cause a switch to send a trap to the management workstations. You can use traps to monitor activities on the switch.

Trap receivers are the devices, typically management workstations or servers, that you want to receive the traps sent by the switch. You specify the trap receivers by their IP addresses. You assign the IP addresses to the community strings.

Each community string can have up to eight trap IP addresses.

It does not matter which community strings you assign to the trap receivers. When the switch sends a trap, it looks at all the community strings and sends the trap to all trap receivers on all community strings. This is true even for community strings that have a access mode of Read only.

If you are not interested in receiving traps, then you do not need to enter the IP addresses of trap receivers.

Default SNMP Community Strings

The AT-S101 Management Software provides two default community strings: public and private. The public string has an access mode of Read-Only and the private string has an access mode of Read/Write. If you activate SNMP management on the switch, you should delete or disable the private community string, which is a standard community string in the industry. Or, change the status of the community string from open to closed to prevent unauthorized changes to the switch.

Creating an SNMP Community

This procedure explains how to create an SNMP community.

To create an SNMP community, perform the following procedure:

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Community Table**.

   The Community Table Page is shown in Figure 31.

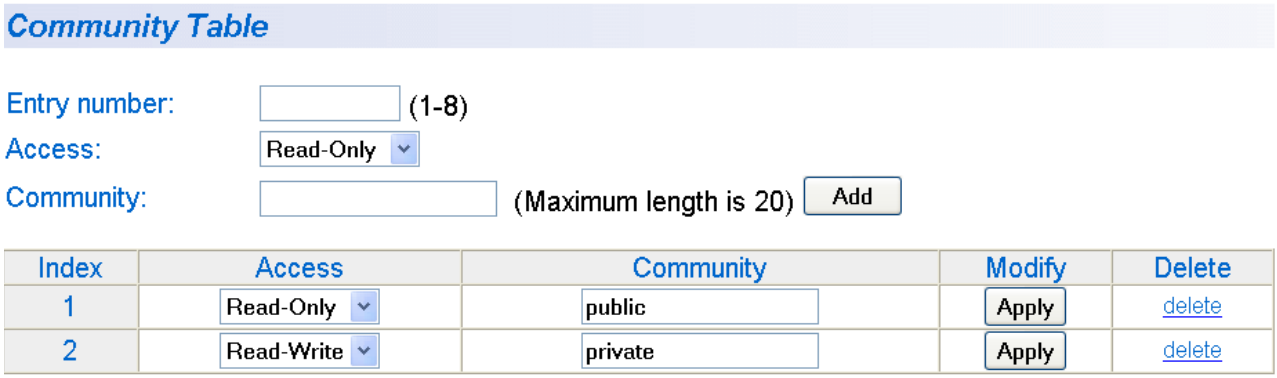

Figure 31. Community Table Page

3. Type an available entry number from 1 through 8 next to the Entry number field.

4. To select the read/write access for the community, use the pull-down menu next to the Access field to select Read-Only access or Read-Write access.

5. Type the name of the new SNMP community in the Community field. Then click **Add**.

   Enter a name between 1 and 20 characters in length.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Modifying an SNMP Community

Use the following procedure to modify the access level or a community name of an SNMP community in the Community Table.

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Community Table**.

   The Community Table Page is shown in Figure 31 on page 104.

3. To change the access level of an SNMP community, select the pull-down menu under the Access column in the Community table for the community you want to modify.

4. Select Read-Only access or Read-Write access.

5. To change the community name, type over an existing community name. Then click **Apply**.

   Note
   You cannot change the index number of an SNMP community.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting an SNMP Community

Use the following procedure to delete an existing SNMP community in the Community Table.

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Community Table**.

   The Community Table page is shown in Figure 31 on page 104.

3. To delete a community, select **delete** in the Community Table next to the community that you want to remove.

   The Community Table Page is updated.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Creating a Host Table

Use the following procedure to create a Host Table.

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Host Table**.

   The Host Table Page is shown in Figure 32.

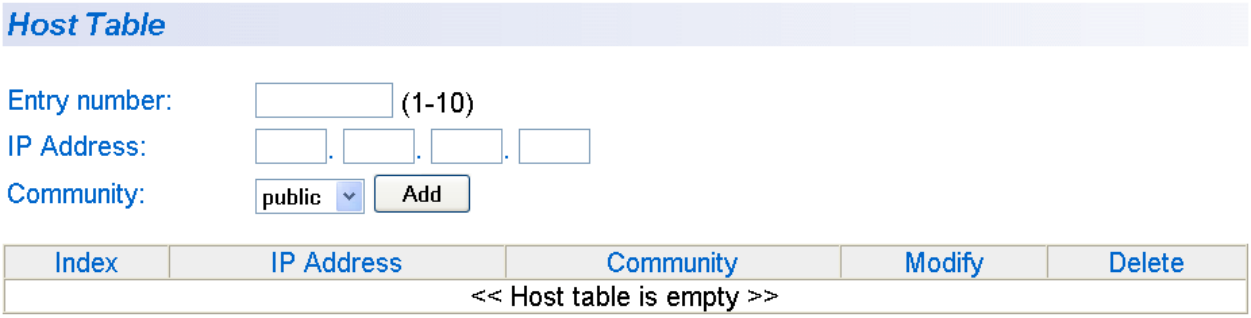

Figure 32. Host Table Page

3. To specify an entry number, type a value between 1 and 10 in the Entry number field.

4. For an SNMP community that you previously defined in the Community Table page, enter an IP address.

   The IP address must be in the xxx.xxx.xxx.xxx format.

5. Select a community name from the pull-down menu next to the Community Name field. Then click **Add**.

   The new host is added to the table.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Modifying a Host Table Entry

To modify the IP address or community name of an entry in the Host Table, use the following procedure:

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Host Table**.

   The Host Table Page is shown in Figure 32 on page 107.

3. To change an IP Address in the table, replace the old IP address with a new one.

4. To change the community name, use the pull-down menu to select a new community name in the Host Table.

5. To activate your changes on the switch, click **Apply** next to the entry that you want to modify.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting a Host Table Entry

Use the following procedure to delete a Host Table entry.

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Host Table**.

   The Host Table Page is shown in Figure 32 on page 107.

3. To delete an entry in the host table, click **delete** next to the entry in the table that you want to remove.

   The Host Table entry is removed from the table. No confirmation message is displayed.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Enabling or Disabling Traps

To enable or disable a trap for an SNMP community, perform the following procedure:

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Trap Setting**.

   The Trap Setting Page is shown in Figure 33.

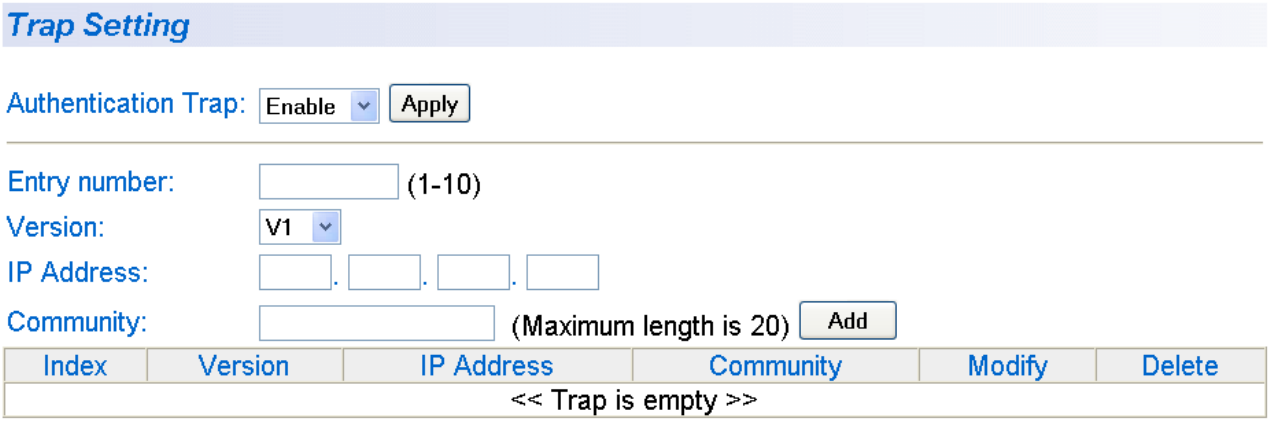

Figure 33. Trap Setting Page

3. Type a trap number between 1 and 10 in the Entry number field.

4. Select the SNMP version of the trap by selecting **V1** for SNMP version 1 or **V2c** for SNMP version 2vc in the Version field.

5. Enter an IP address, in the xxx.xxx.xxx.xxx format, in the IP Address field.

6. Enter a previously defined community name in the Community field. Then click **Add**.

   A new trap is displayed in the Trap Setting table.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Modifying Traps

To modify the SNMP version, IP address, or community name of a trap, perform the following procedure:

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Trap Setting**.

   The Trap Setting Page is shown in Figure 33 on page 110.

3. Within the Trap Setting table, select a pull-down menu in the Version column to the change the SNMP version of a trap that you want to modify.

   Select the SNMP version of the trap by selecting **V1** for SNMP version 1 or **V2c** for SNMP version 2vc.

4. Change an IP address by typing in the new IP address for a particular community within the Trap Setting table.

   Use the IP address format: xxx.xxx.xxx.xxx

5. Change the Community Name by replacing the old name with the new one.

6. To activate your changes on the switch click **Apply**.

   The Trap Setting Page is updated.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting Traps

To delete a trap from an SNMP community, perform the following procedure:

1. From the menu on the left side of the page, select the **SNMP** folder.

   The **SNMP** folder expands.

2. From the **SNMP** folder**,** select **Trap Setting**.

   The Trap Setting Page is shown in Figure 33 on page 110.

3. In the Trap table, click delete next to the trap you want to delete from the table.

   The trap is removed from the Trap Setting Page. A warning message is not displayed.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 9

IGMP Snooping

This chapter contains the following procedures for working with IGMP Snooping in the web interface. Sections in the chapter include:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

IGMP enables IPv4 routers to create lists of nodes that are members of multicast groups. (A multicast group is a group of end nodes that want to receive multicast packets from a multicast application.) The router creates a multicast membership list by periodically sending out queries to the local area networks connected to its ports.

A node that wants to become a member of a multicast group responds to a query by sending a *report* which indicates an end node's desire to become a member of a multicast group. Nodes that join a multicast group are referred to as *host nodes*. After becoming a member of a multicast group, a host node must continue to periodically issue reports to remain a member.

After the router has received a report from a host node, it notes the multicast group that the host node wants to join and the port on the router where the node is located. Any multicast packets belonging to that multicast group are then forwarded by the router out the port. If a particular port on the router has no nodes that want to be members of multicast groups, the router does not send multicast packets from the port. This improves network performance by restricting multicast packets only to router ports where host nodes are located.

There are three versions of IGMP — versions 1, 2, and 3. One of the differences between the versions is how a host node signals that it no longer wants to be a member of a multicast group. In version 1 it stops sending reports. If a router does not receive a report from a host node after a predefined length of time, referred to as a *time-out value*, it assumes that the host node no longer wants to receive multicast frames, and removes it from the membership list of the multicast group.

In version 2, a host node exits from a multicast group by sending a *leave request*. After receiving a leave request from a host node, the router removes the node from appropriate membership list. The router also stops sending multicast packets from the port to which the node is connected if it determines there are no further host nodes on the port.

Version 3 adds the ability of host nodes to join or leave specific sources in a multicast group.

The IGMP snooping feature on the AT-GS950/8POE switches support IGMP versions 1 and 2. The switch monitors the flow of queries from a router and reports and leave messages from host nodes to build its own multicast membership lists. It uses the lists to forward multicast packets only to ports where there are host nodes that are members of multicast groups. This improves switch performance and network security by restricting the flow of multicast packets only to those ports connected to host nodes.

Without IGMP snooping, a switch would have to flood multicast packets from all of its ports, except the port on which it received the packet. Such flooding of packets can negatively impact network performance.

The AT-GS950/8POE switches maintain a list of multicast groups through an adjustable timeout value, which controls how frequently it expects to see reports from end nodes that want to remain members of multicast groups, and by processing leave requests.

Note
By default, IGMP snooping is disabled on the switch.

Configuring IGMP Snooping

This procedure explains how to set IGMP snooping on the switch and set the IGMP Snooping age-out timer.

To configure IGMP snooping, perform the following procedure:

1. From the menu on the left side of the page, select the **Bridge** folder.

   The **Bridge** folder expands.

2. From the **Bridge** folder**,** select **IGMP Snooping**.

   The IGMP Snooping Page is shown in Figure 34.

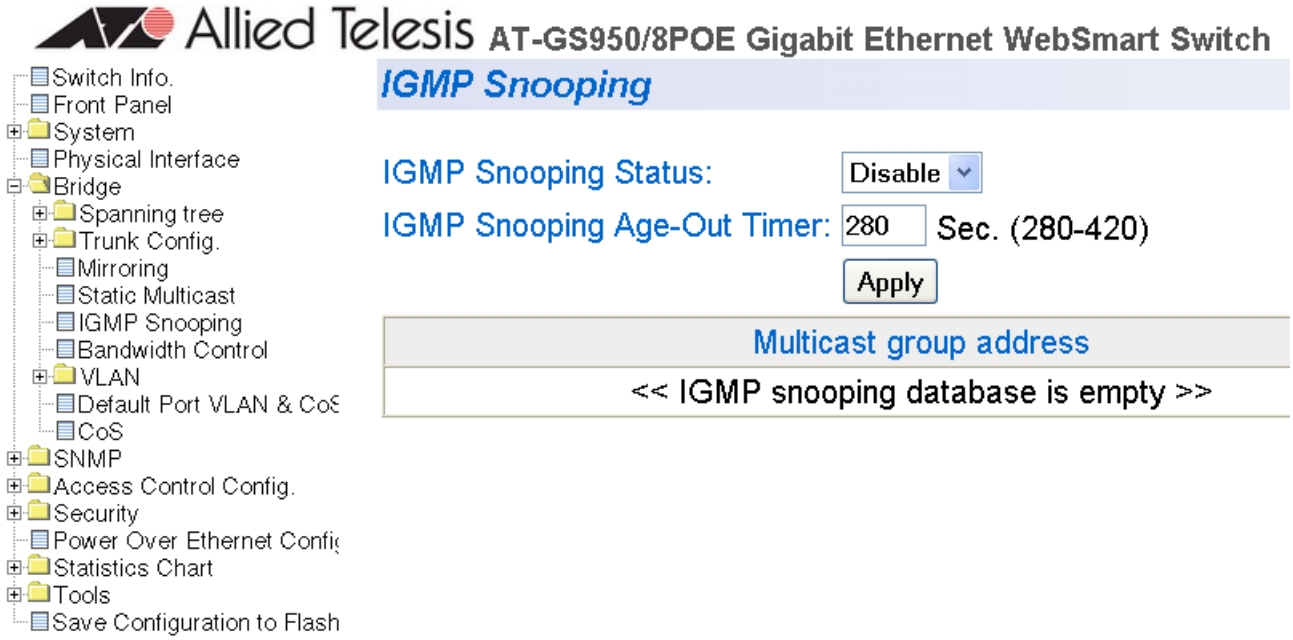

Figure 34. IGMP Snooping Page

3. To enable or disable IGMP Snooping on the switch, select **Enable** or **Disable**. Then press **Apply**.

   By default, IGMP is disabled.

4. To set the age-out timer, type the number of seconds you want the switch to wait before it purges an inactive dynamic MAC address. Then press **Apply**.

   The Set Age-Out Timer field is set to 280 seconds by default. The range of this parameter is between 280 to 420 seconds.

Note
The **Multicast Group Address** table contains MAC addresses of nodes that are members of multicast groups. To set a Multicast Group Address, see "Setting a Static Multicast Address" on page 129.

5. After you have configured a Group MAC Address on the Static Multicast Address Page, the IGMP Snooping Page is updated with the Multicast Group address. See Figure 35.

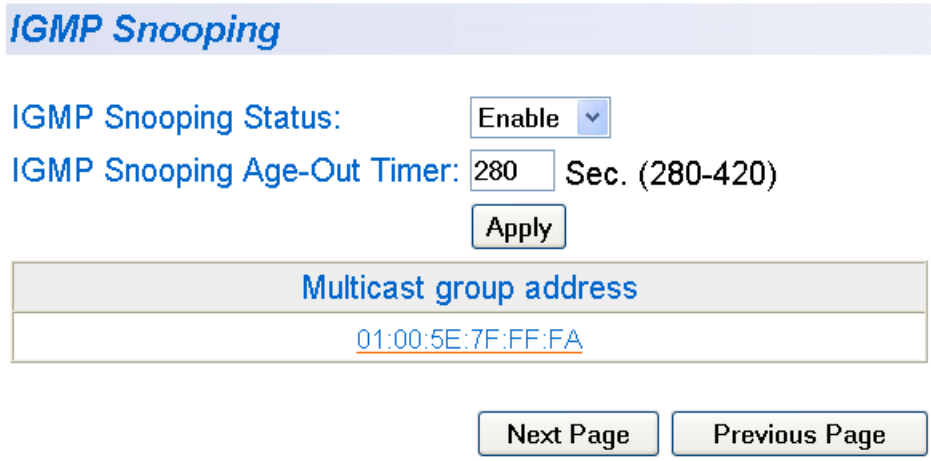

Figure 35. IGMP Snooping Page with MAC Address

6. To display more information about the multicast group address, click on the MAC address.

   The IGMP - Group Members Page is displayed. See Figure 36 on page 117.

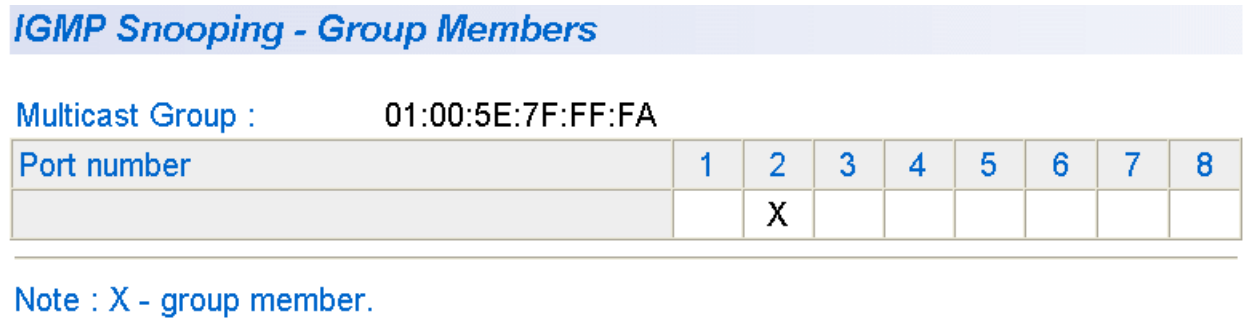

Figure 36. IGMP Snooping —Group Members Page

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 10

Bandwidth Control

This chapter contains a description of the bandwidth features and the procedure for configuring them. This chapter includes the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

The bandwidth control feature allows you regulate the reception rate of broadcast, multicast, and destination lookup failure packets. The AT-S101 software allows you to set separate limits for each port beyond which each of the different packet types are discarded. Each setting can be configured on individual ports or on all of the ports of the AT-GS950/8POE switch. Traffic is measured in packets per second. See the following definitions for more information about these settings.

- Destination Lookup Failure-- The Destination Lookup Failure (DFL) setting is concerned with comparing the destination MAC address of a packet received by the switch to the forwarding database. When an AT-GS950/8POE switch receives a packet, it scans the forwarding database and looks for a match to the destination MAC address in the received packet. If the MAC address is not present in the forwarding database, then the packet is flooded according to the VLAN ingress and egress rules. By default, this setting is disabled on the switch which means that all DLF packets are automatically forwarded according to the VLAN ingress and egress rules.
- Broadcast Setting-- The broadcast setting applies to allowing or denying broadcast packets on each port.
- Multicast Setting-- The multicast setting applies to allowing or denying multicast packets on each port.
- Threshold Level-- In regards to Bandwidth control, the threshold level is the number of DFL, broadcast, and multicast packets that are sent by or received from a port. This value is measured in packets per second. You can set the threshold level to low, medium, or high.

Setting Bandwidth Control

This procedure explains how to set DLF, broadcast, multicast, and threshold levels for each port on the AT-GS950/8POE switch.

To change the default settings, perform the following procedure:

1. From the menu on the left side of the page, select the **Bridge** folder.

   The **Bridge** folder expands to display several folders including the **Trunk Config**. folder.

2. From the **Trunk Config.** folder, select **Bandwidth Control**.

   The Bandwidth Control Page is shown in Figure 37.

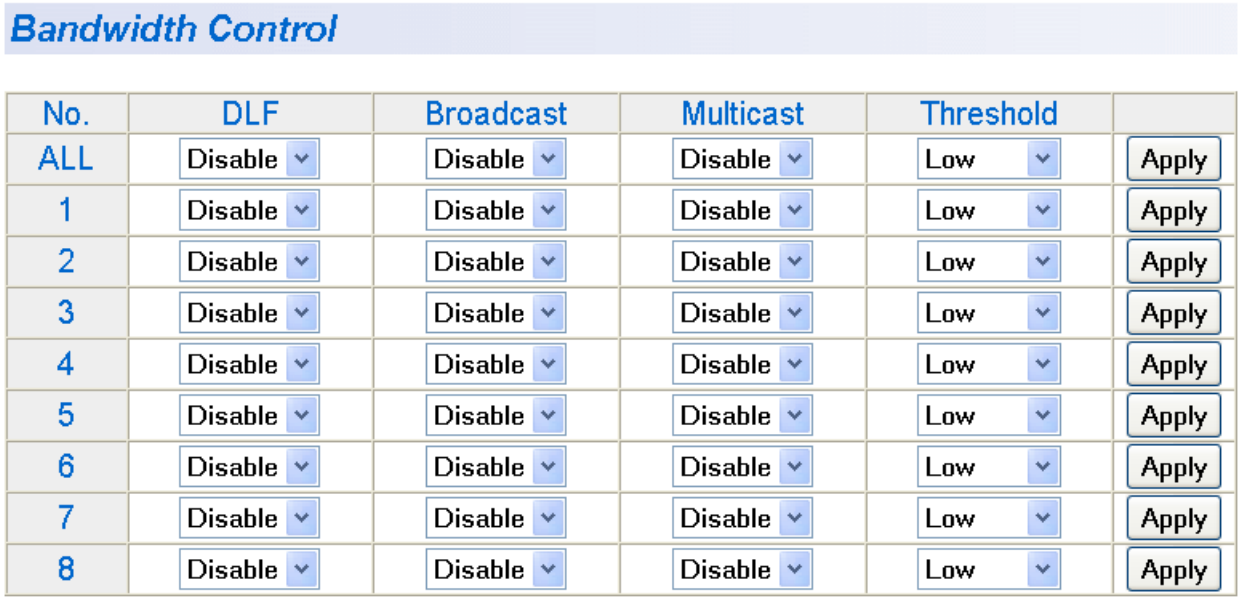

Figure 37. Bandwidth Control Page

3. To enable or disable the DLF field, select **Enable** or **Disable** from the DLF pull-down menu next to the port that you want to change. Then click **Apply**.

   The default is **Disable**. You can use the option next to the ALL row to set all of the ports to the same setting.

4. To enable or disable ingress and egress Broadcast packets, select **Enable** or **Disable** from the Broadcast pull-down menu next to the port that you want to change. Then click **Apply**.

The default is **Disable**. You can use the option next to the ALL row to set all of the ports to the same setting.

5. To enable or disable ingress and egress Multicast packets, select **Enable** or **Disable** from the Multicast pull-down menu next to the port that you want to change. Then click **Apply**.

   The default is **Disable**. You can use the option next to the ALL row to set all of the ports to the same setting.

6. To set the **Threshold** field, use the pull-down menu next to the port that you want to change. Select Low, Medium, or High which correspond to the following values:

   Low - Specifies 450 to 550 packets per second.

   Medium - Specifies 880 to 1,000 packets per second.

   High - Specifies 2,200 to 2,500 packets per second.

   The default is **Low**. You can use the option next to the ALL row to set all of the ports to the same setting.

7. Then click **Apply**.

8. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 11

Port Mirroring

This chapter contains the procedure for setting up port mirroring. Port mirroring allows you to unobtrusively monitor the ingress and egress traffic on a port by having the traffic copied to another port. This chapter contains the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

The port mirroring feature allows you to unobtrusively monitor the traffic received and transmitted on one or more ports by copying the traffic to another switch port. You can connect a network analyzer to the port where the traffic is copied and monitor the traffic on the other ports without impacting network performance or speed.

A port mirror has two component ports. The port or ports whose traffic you want to mirror is called the *source port(s)*. The port where the traffic will be copied to is called the *mirroring port*.

Observe the following guidelines when you create a port mirror:

- You can select more than one source port at a time. However, the more ports you mirror, the less likely the monitor port is able to handle all the traffic. For example, if you mirror the traffic of six heavily active ports, the destination port is likely to drop packets, meaning that it does not provide an accurate mirror of the traffic of the six source ports.
- The source and mirror ports must be located on the same switch.
- You can mirror the ingress or egress traffic of the source ports or both.

Configuring Port Mirroring

To set up port mirroring, perform the following procedure:

1. Select the **Bridge** folder**.**

   The Bridge folder expands.

2. From the **Bridge** folder**,** select **Mirroring**.

   The Mirroring Page is shown in Figure 38.

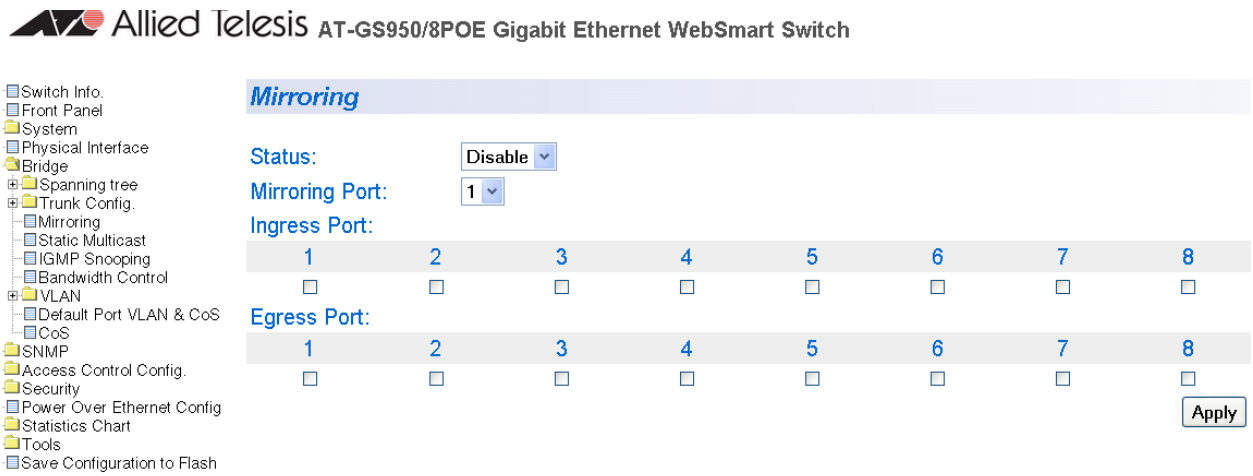

Figure 38. Mirroring Page

3. Click **Mirroring Port** and from the pull-down menu select the port where the network analyzer is connected.

4. For the mirrored port, select the port whose ingress, egress, or both ingress and egress traffic you want to monitor.

5. Click **Apply** on the right-hand side of the page.

   Port mirroring is immediately enabled on the switch. You can now connect a data analyzer to the mirroring port to monitor the traffic on the other port.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Disabling Port Mirroring

To disable port mirroring, perform the following procedure:

1. Select the **Bridge** folder**.**

   The Bridge folder expands.

2. From the **Bridge** folder**,** select **Mirroring**.

   The Mirroring page is shown in Figure 38 on page 125.

3. From the Mirroring Status list, select **Disable** and click **Apply**.

   Port mirroring is immediately disabled on the switch. You can now use the mirroring port for regular network operations.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 12

Static Multicast MAC Address

This chapter contains a description of the static multicast MAC address feature and the procedure for configuring it. This chapter includes the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

The AT-GS950/8POE switch has a MAC address table with a storage capacity of 8,000 entries. The table stores the MAC addresses of the network nodes connected to its ports and the port number where each address was learned.

There are two types of MAC addresses—dynamic and static.

Dynamic MAC addresses are addresses that the switch learns by examining the source MAC addresses of the frames received by the ports. This type of MAC address is not stored indefinitely in the MAC address table. The switch deletes a dynamic MAC address from the table if it does not receive any frames from the node after a specified period of time. The switch assumes that the node is no longer active and that its MAC address can be purged from the table. This prevents the MAC address table from becoming filled with addresses of nodes that are no longer active.

The MAC address table can also store a *static MAC address* which is a MAC address of an end node that you assign to a switch port manually. A static MAC address remains in the table indefinitely and is never deleted, even when the end node is inactive.

There are two reasons to enter static MAC addresses. You may want to enter end nodes the switch does not learn in its normal dynamic learning process. Or, you want a MAC address to remain permanently in the table, even when the end node is inactive.

Static multicast addresses are a subset of the static MAC addresses. With the Static Multicast Address feature, you can add a static multicast address to the MAC address table. Then you assign a static MAC address to a port or ports which are called Group members in the AT-S101 software interface. Each port has a maximum limit of 32 static multicast addresses.

In some network environments that are confined to one LAN (such as an industrial application with a server, a switch and many controllers), there may be various multicast streams that need to be distributed to some network nodes, but not others. If the data sent in these streams is time-sensitive and cannot be delayed because of the configuration time associated with the IGMP Snooping feature, then static multicast addresses may be the solution.

If a multicast address and its associated ports of the switch are predefined within the network design and they will not change over time, then they can be manually entered as static entries into the MAC address table. This allows the multicast stream to be forwarded immediately to those predefined ports entered in the MAC table without any configuration delays or loss of data.

Setting a Static Multicast Address

This procedure explains how to set the static multicast feature for each port on the AT-GS950/8POE switch.

To add a static MAC address to the switch, perform the following procedure:

1. From the menu on the left side of the page, select the **Bridge** folder.

   The **Bridge** folder expands to display several folders including the **Trunk Config**. folder.

2. From the **Trunk Config.** folder, select **Static Multicast**.

   The Static Multicast Address Table Page is displayed. See Figure 39.

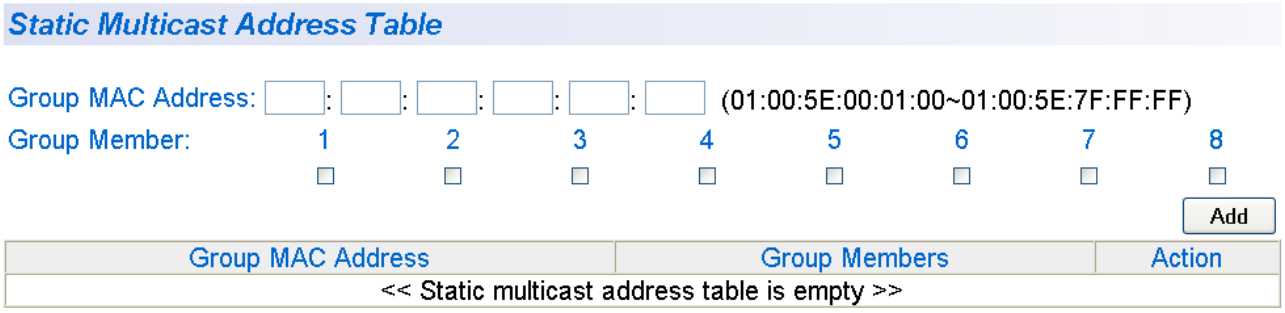

Figure 39. Static Multicast Address Table Page

3. In the Group MAC Address field, enter a MAC address.

   The range is from 01:00:5E:00:01:00 to 01:00:5E:7F:FF:FF.

4. Assign the MAC address a Group Member (or members) for selecting the checkbox below each group member. Then click **Add**.

   Note
   Each group member corresponds to a port number. In addition, you can assign a port a maximum limit of 32 static multicast addresses.

   The Static Multicast Address Table is updated with the new Group MAC Address.

Note
The Group MAC Address values that you enter on the Static Multicast Address Table Page are also displayed on the IGMP Snooping Page. For more information, see “Configuring IGMP Snooping” on page 116.

5. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Modifying a Static Multicast Address

To modify the port assignment of a multicast MAC address in the MAC address table, perform the following procedure:

1. From the menu on the left side of the page, select the **Bridge** folder.

   The **Bridge** folder expands to display several folders including the **Trunk Config**. folder.

2. From the **Trunk Config.** folder, select **Static Multicast**.

   The Static Multicast Address Table Page is displayed. See Figure 39 on page 129.

3. Select modify next to the static MAC address that you want to modify.

   The Modify Static Multicast Address Page is displayed. See Figure 40.

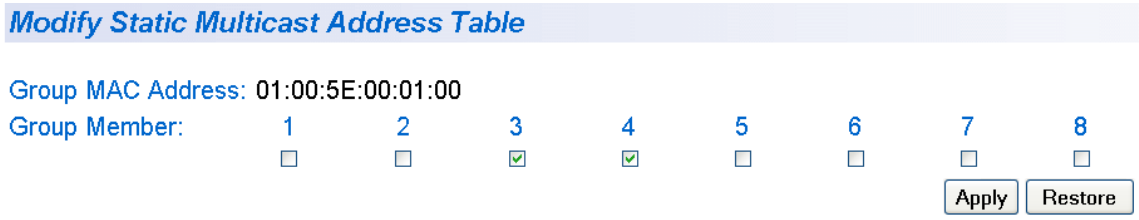

Figure 40. Modify Static Multicast Address Page

4. In the Group Member row, select the ports that you want to include in the group MAC address.

   Selected ports are indicated with a check mark. Each port has a maximum limit of 32 static multicast addresses.

5. Then click **Apply** to apply your changes.

   Note
   To restore the original group member ports, click **Restore**.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting a Static Multicast Address

To delete a multicast MAC address from the MAC address table, perform the following procedure:

1. From the menu on the left side of the page, select the **Bridge** folder.

   The **Bridge** folder expands to display several folders including the **Trunk Config**. folder.

2. From the **Trunk Config.** folder, select **Static Multicast**.

   The Static Multicast Address Table Page is displayed. See Figure 39 on page 129.

3. Select **delete** next to the static multicast address that you want to remove.

   The static multicast address is removed from the Static Multicast Address Table Page.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 13

Spanning Tree and Rapid Spanning Tree Protocols

This chapter provides background information about the Spanning Tree Protocol (STP) and the Rapid Spanning Tree Protocol (RSTP). In addition, there are procedures to configure STP and RSTP. The sections in the chapter include:

For detailed information about STP, refer to IEEE Std 802.1D. For detailed information about RSTP, refer to IEEE Std 802.1w.

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

The performance of a Ethernet network can be negatively impacted by the formation of a data loop in the network topology. A data loop exists when two or more nodes on a network can transmit data to each other over more than one data path. The problem that data loops pose is that data packets can become caught in repeating cycles, referred to as broadcast storms, that needlessly consume network bandwidth and can significantly reduce network performance.

STP and RSTP prevent data loops from forming by ensuring that only one path exists between the end nodes in your network. Where multiple paths exist, these protocols place the extra paths in a standby or blocking mode, leaving only one main active path.

In addition, STP and RSTP can activate a redundant path if the main path goes down. So not only do these protocols guard against multiple links between segments and the risk of broadcast storms, but they can also maintain network connectivity by activating a backup redundant path in case a main link fails.

Where the two protocols differ is in the time each takes to complete the process referred to as *convergence*. When a change is made to the network topology, such as the addition of a new bridge, a spanning tree protocol must determine whether there are redundant paths that must be blocked to prevent data loops, or activated to maintain communications between the various network segments. This is the process of convergence.

With STP, convergence can take up to a minute to complete in a large network. This can result in the loss of communication between various parts of the network during the convergence process, and the subsequent lost of data packets.

RSTP is much faster. It can complete a convergence in seconds, and so greatly diminish the possible impact the process can have on your network.

Only one spanning tree can be active on the switch at a time. The default protocol is RSTP.

The STP implementation on the AT-S101 Management Software complies with the IEEE 802.1d standard. The RSTP implementation complies with the IEEE 802.1w standard. The following subsections provide a basic overview on how STP and RSTP operate and define the different parameters that you can adjust.

Bridge Priority and the Root Bridge

The first task that bridges perform when a spanning tree protocol is activated on a network is the selection of a *root bridge*. A root bridge distributes network topology information to the other network bridges and is used by the other bridges to determine if there are redundant paths in the network.

A root bridge is selected by the *bridge priority* number, also referred to as the bridge identifier, and sometimes the bridge's MAC address. The bridge with the lowest bridge priority number in the network is selected as the root bridge. If two or more bridges have the same bridge priority number, of those bridges the one with the lowest MAC address is designated as the root bridge.

You can change the bridge priority number in the AT-S101 Management Software. You can designate which switch on your network as the root bridge by giving it the lowest bridge priority number. You may also consider which bridge should function as the backup root bridge in the event you need to take the primary root bridge offline and assign that bridge the second lowest bridge identifier number.

The bridge priority has a range 0 to 61440 in increments of 4096. To make this easier for you, the AT-S101 Management Software divides the range into increments. You specify the increment that represents the desired bridge priority value. The range is divided into sixteen increments, as shown in Table 2.

Table 2. Bridge Priority Value Increments

Increment	Bridge Priority	Increment	Bridge Priority
0	0	8	32768
1	4096	9	36864
2	8192	10	40960
3	12288	11	45056
4	16384	12	49152
5	20480	13	53248
6	24576	14	57344
7	28672	15	61440

Path Costs and Port Costs

After the root bridge has been selected, the bridges determine if the network contains redundant paths and, if one is found, select a preferred path while placing the redundant paths in a backup or blocking state.

Where there is only one path between a bridge and the root bridge, the bridge is referred to as the *designated bridge* and the port through which the bridge is communicating with the root bridge is referred to as the *root port*.

If redundant paths exist, the bridges that are a part of the paths must determine which path is the primary, active path, and which path(s) are placed in the standby, blocking mode. This is accomplished by an determination of *path costs*. The path offering the lowest cost to the root bridge becomes the primary path and all other redundant paths are placed into blocking state.

Path cost is determined by evaluating *port costs*. Every port on a bridge participating in STP has a cost associated with it. The cost of a port on a bridge is typically based on port speed. The faster the port, the lower the port cost. The exception to this is the ports on the root bridge, where all ports have a port cost of 0.

Path cost is simply the sum of the port costs between a bridge and the root bridge.

The port cost of a port on the AT-GS950/8POE switch is adjustable through the AT-S101 Management Software. For STP and RSTP, the range is from 0 to 20,000,000.

Port Priority

If two paths have the same port cost, the bridges must select a preferred path. In some instances this can involve the use of the *port priority* parameter which is used as a tie breaker when two paths have the same cost.

The range for port priority is 0 to 240. As with bridge priority, this range is broken into increments, in this case multiples of 16. To select a port priority for a port, you enter the increment of the desired value. Table 3 lists the values and increments. The default value is 128, which is increment 8.

Table 3. Port Priority Value Increments

Increment	Bridge Priority	Increment	Bridge Priority
0	0	8	128
1	16	9	144
2	32	10	160

Table 3. Port Priority Value Increments (Continued)

Increment	Bridge Priority	Increment	Bridge Priority
3	48	11	176
4	64	12	192
5	80	13	208
6	96	14	224
7	112	15	240

Forwarding Delay and Topology Changes

If there is a change in the network topology due to a failure, removal, or addition of any active components, the active topology also changes. This may trigger a change in the state of some blocked ports. However, a change in a port state is not activated immediately.

It may take time for the root bridge to notify all bridges that a topology change has occurred, especially if it is a large network. A temporary data loop could occur if a topology change is made before all bridges have been notified and that could adversely impact network performance.

To forestall the formation of temporary data loops during topology changes, a port designated to change from blocking to forwarding passes through two additional states—listening and learning—before it begins to forward frames. The amount of time a port spends in these states is set by the forwarding *delay* value. This value states the amount of time that a port spends in the listening and learning states prior to changing to the forwarding state.

The forwarding delay value is adjustable in the AT-S101 Management Software. The appropriate value for this parameter depends on a number of variables; the size of your network is a primary factor. For large networks, you should specify a value large enough to allow the root bridge sufficient time to propagate a topology change throughout the entire network. For small networks, you should not specify a value so large that a topology change is unnecessarily delayed, which could result in the delay or loss of some data packets.

Note
The forwarding delay parameter applies only to ports on the switch that are operating STP-compatible mode.

Hello Time and Bridge Protocol Data Units (BPDU)

The bridges that are part of a spanning tree domain communicate with each other using a bridge broadcast frame that contains a special section devoted to carrying STP or RSTP information. This portion of the frame is referred to as the bridge protocol data unit (BPDU). When a bridge is brought online, it issues a BPDU in order to determine whether a root bridge has already been selected on the network, and if not, whether it has the lowest bridge priority number of all the bridges and should therefore become the root bridge.

The root bridge periodically transmits a BPDU to determine whether there have been any changes to the network topology and to inform other bridges of topology changes. The frequency with which the root bridge sends out a BPDU is called the *hello time*. This is a value that you can set in the AT-S101 Management Software. The interval is measured in

seconds and the default is two seconds. Consequently, if the AT-GS950/8POE switch is selected as the root bridge of a spanning tree domain, it transmits a BPDU every two seconds.

Point-to-Point and Edge Ports

Note
This section applies only to RSTP.

Part of the task of configuring RSTP is defining the port types on the bridge. This relates to the device(s) connected to the port. With the port types defined, RSTP can reconfigure a network much quicker than STP when a change in network topology is detected.

There are two possible selections:

- Point-to-point port
- Edge port

If a bridge port is operating in full-duplex mode, than the port is functioning as a point-to-point port. Figure 41 illustrates two AT-GS950/8POE switches that are connected with one data link. With the link operating in full-duplex, the ports are point-to-point ports.

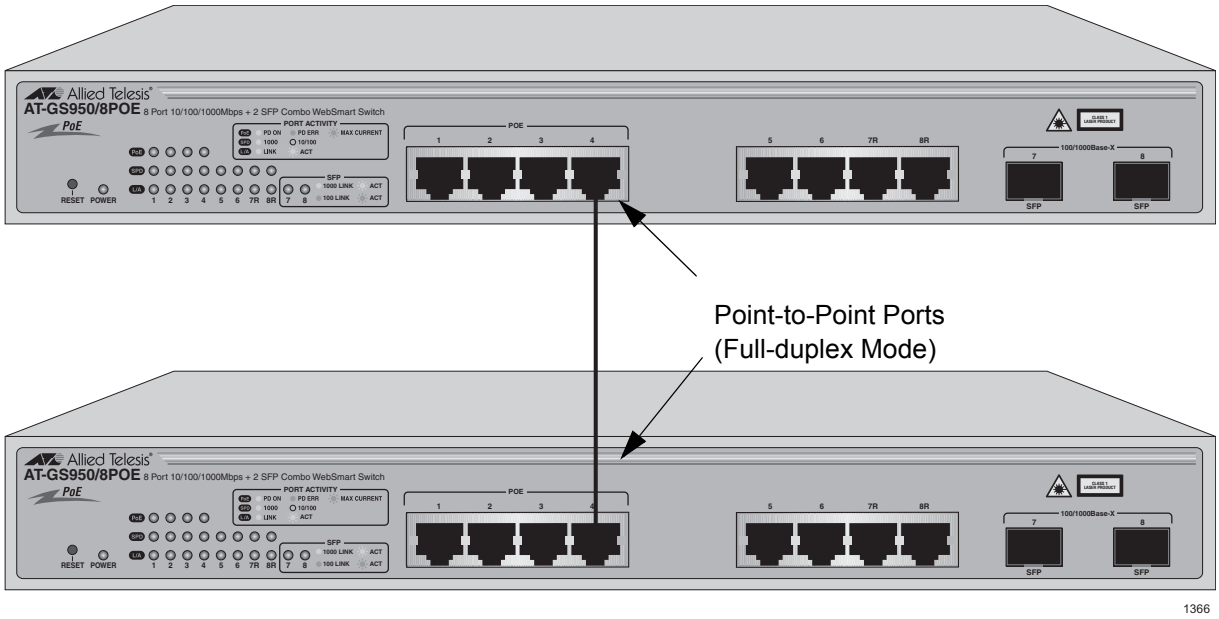

Figure 41. Point-to-Point Ports

If a port is operating in half-duplex mode and is not connected to any further bridges participating in STP or RSTP, then the port is an edge port. Figure 42 on page 140 illustrates an edge port on an AT-GS950/8POE

switch. The port is connected to an Ethernet hub, which in turn is connected to a series of Ethernet workstations. This is an edge port because it is connected to a device operating at half-duplex mode and there are no participating STP or RSTP devices connected to it.

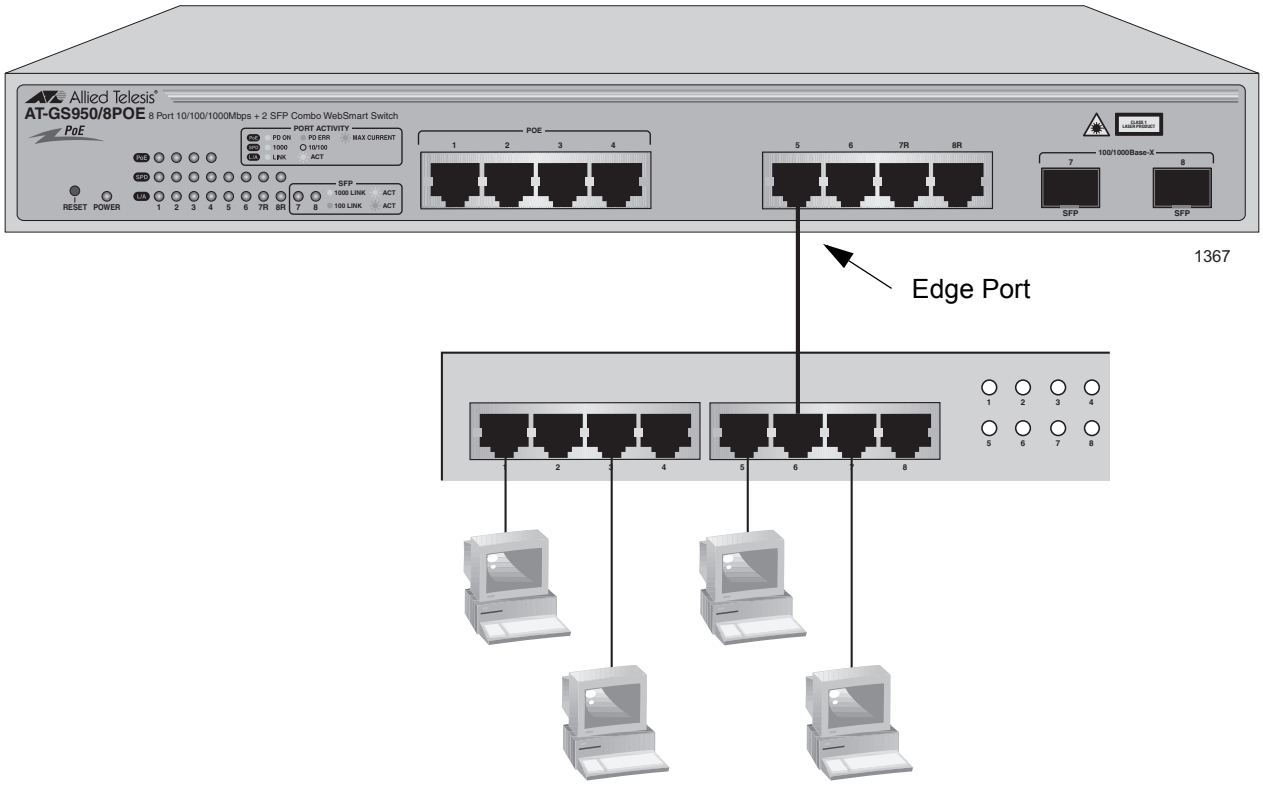

Figure 42. Edge Port

A port can be both a point-to-point and an edge port at the same time. It operates in full-duplex and has no STP or RSTP devices connected to it. Figure 43 on page 141 illustrates a port functioning as both a point-to-point and edge port.

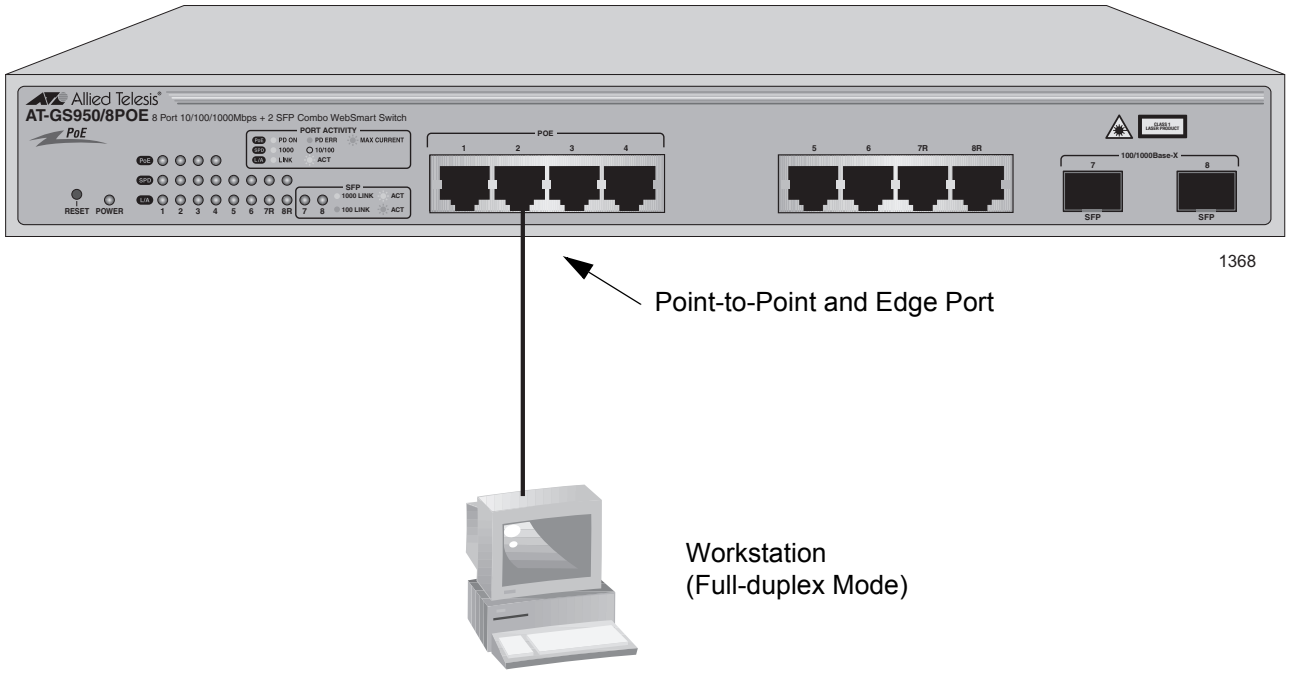

Figure 43. Point-to-Point and Edge Port

Determining whether a bridge port is point-to-point, edge, or both, can be a bit confusing. For that reason, do not change the default values for this RSTP feature unless you have a good grasp of the concepts. In most cases, the default values work well.

Mixed STP and RSTP Networks

RSTP IEEE 802.1w is fully compliant with STP IEEE 802.1d. Your network can consist of bridges running both protocols. STP and RSTP in the same network can operate together to create a single spanning tree domain.

If you decide to activate spanning tree on the switch, Allied Telesis recommends RSTP instead of STP even when all of other switches in the network are running STP. The AT-GS950/8POE switch can combine its RSTP with the STP of the other switches. The AT-GS950/8POE switch monitors the traffic on each port for BPDU packets. Ports that receive RSTP BPDU packets operates in RSTP mode while ports receiving STP BPDU packets operate in STP mode.

Spanning Tree and VLANs

The spanning tree implementation in the AT-S101 Management Software is a single-instance spanning tree. The switch supports just one spanning tree. You cannot define multiple spanning trees.

The single spanning tree encompasses all ports on the switch. If the ports are divided into different VLANs, the spanning tree crosses the VLAN boundaries. This point can pose a problem in networks containing multiple VLANs that span different switches and are connected with untagged ports. In this situation, STP blocks a data link because it detects a data loop. This can cause fragmentation of your VLANs.

This issue is illustrated in Figure 44. Two VLANs, Sales and Production, span two AT-GS950/8POE switches. Two links consisting of untagged ports connect the separate parts of each VLAN. If STP or RSTP is activated on the switches, one of the links is disabled. In the example, the port on the top switch that links the two parts of the Production VLAN is changed to the block state. This leaves the two parts of the Production VLAN unable to communicate with each other.

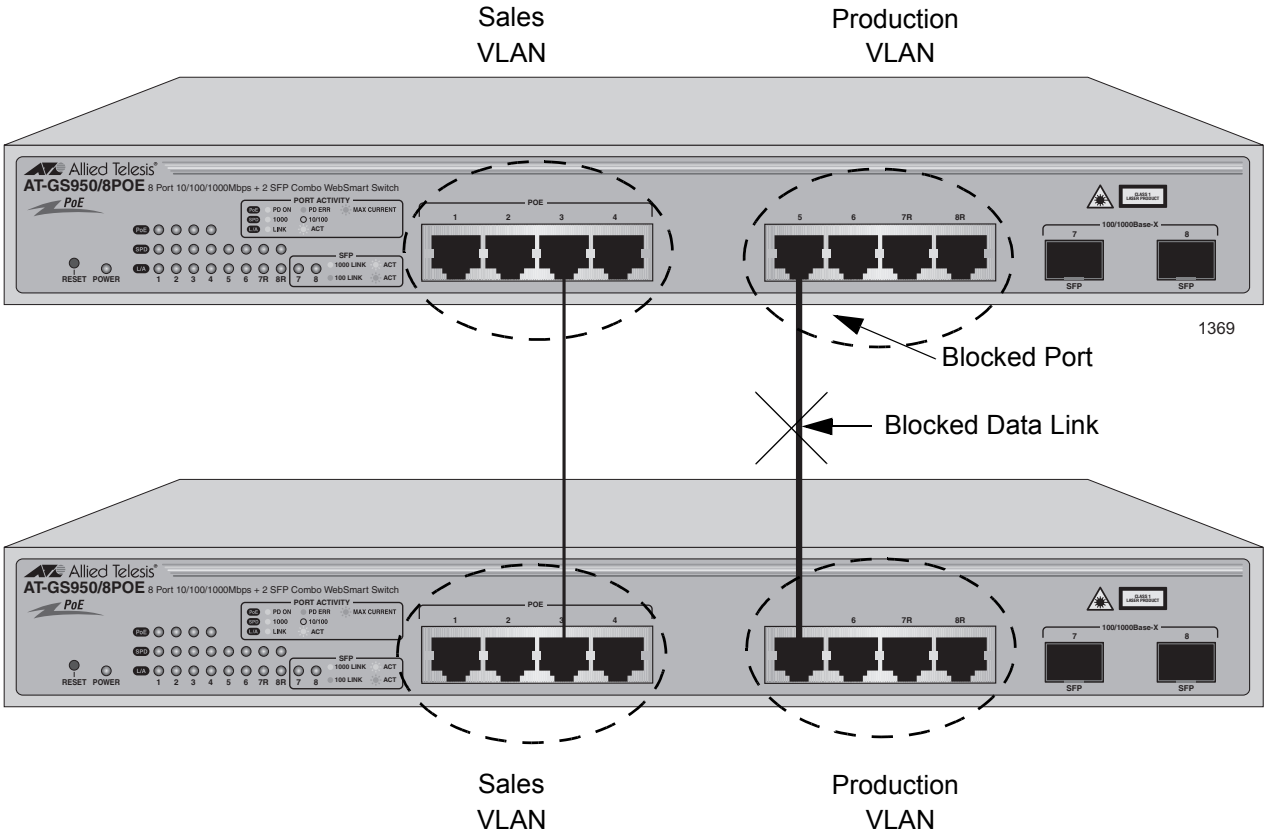

Figure 44. VLAN Fragmentation

You can avoid this problem by not activating spanning tree or by connecting VLANs using tagged instead of untagged ports. (For information on tagged and untagged ports, refer to Chapter 22, “Port-based and Tagged VLANs” on page 247.)

Basic STP and RSTP Configuration

To configure the basic STP and RSTP settings, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The **Spanning Tree** folder is displayed.

2. From the **Bridge** folder, select the **Spanning tree** folder.

3. From the **Spanning tree** folder, select **RSTP**.

   The Rapid Spanning Tree Configuration Page is shown in Figure 45.

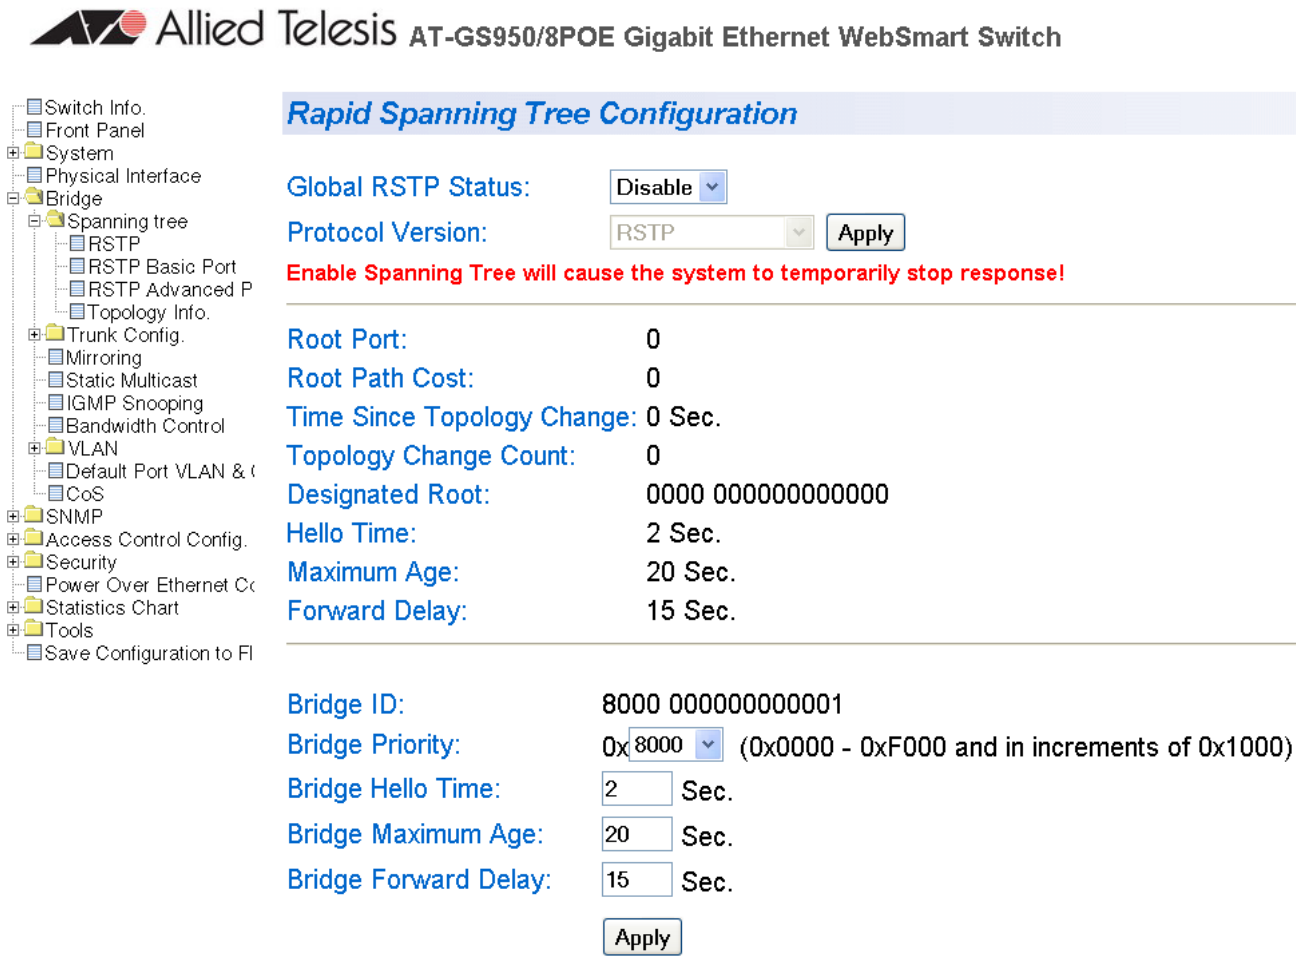

Figure 45. Rapid Spanning Tree Configuration Page

The RSTP Configuration page allows you to configure RSTP as well as to view the current settings. In the upper portion of the page, you can set the following parameters:

Global RSTP Status
Set this field to enable to activate RSTP on the switch. The default is disable.

Protocol Version
Set this field to enable to activate RSTP or STP on the switch. This field is greyed out until you set the Global RSTP Status to enable. To activate this field click **Apply** and then select RSTP or STP-compatible. The default value is RSTP.

This section contains the following items in the middle portion of the web page. You cannot change these fields.

Root Port
The active port on the switch that is communicating with the root bridge. If the switch is the root bridge for the LAN, then there is no root port and the root port parameter is set to 0.

Root Path Cost
The sum of all the root port costs of all the bridges between the switch's root port and the root bridge including the switch's root port cost.

Time Since Topology Change
The time in seconds since the last topology change took place. When RSTP detects a change to the LAN's topology or when the switch is rebooted, this parameter is reset to 0 seconds and begins incrementing until the next topology change is detected.

Topology Change Count
An integer that reflects the number of times RSTP has detected a topology change on the LAN since the switch was initially powered on or rebooted.

The following parameters refer to the designated root bridge. You cannot change these fields.

Designated Root
This parameter includes two fields: the root bridge priority and the MAC address of the root bridge. For example, 1000 00C08F1211BB shows the root bridge priority as 1000, and 00C08F1211BB as the MAC address.

Hello Time
The hello time. See "Hello Time and Bridge Protocol Data Units (BPDU)" on page 138. This parameter affects only the root bridge.

Maximum Age
The maximum amount of time that BPDUs are stored before being deleted on the root bridge.

Forward Delay
The time interval between generating and sending configuration messages by the root bridge.

The bottom section of the web page provides information about the bridge. The following parameters appear in the bottom third of the web page:

Bridge ID
The MAC address of the bridge. The bridge identifier is use as a tie breaker in the selection of the root bridge when two or more bridges have the same bridge priority. You cannot change this setting.

Bridge Priority
The priority number for the bridge, in hexadecimal format. This number is used to determine the root bridge for RSTP. The bridge with the lowest priority number is selected as the root bridge. If two or more bridges have the same priority value, that is, the lowest value of all the other bridges, then the bridge with the numerically lowest MAC address becomes the root bridge. When a root bridge goes offline, the bridge with the lowest priority number automatically takes over as the root bridge. This parameter can be from 0X0000 to 0XF000, with 0XF000 being the highest priority.

Bridge Hello Time
This is the time interval between generating and sending configuration messages by the bridge. This parameter is active only when the switch is the root bridge.

Bridge Maximum Age
The length of time after which stored bridge protocol data units (BPDUs) are deleted by the bridge.

Bridge Forward Delay
This is the time interval between generating and sending configuration messages by the bridge.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Configuring RSTP Port Settings

This section contains the following topics:

Configuring the Basic RSTP Port Settings

To configure the basic RSTP port settings, perform the following procedure:

From the menu on the left side of the page, select **Bridge**.

The **Bridge** folder expands.

5. From the **Bridge** folder, select the **Spanning tree** folder.
6. From the **Spanning tree** folder, select the **RSTP Basic Port** folder.

The RSTP Basic Port Configuration Page is shown in Figure 46.

RSTP Basic Port Configuration

Port	Trunk	Link Status	Port State	Role	STP Status	Priority	Path Cost	
All	-	-	-	-	Enable			Apply
1	---	Up	Forwarding	Disabled	Enable	128	200000	Apply
2	---	Down	Forwarding	Disabled	Enable	128	20000	Apply
3	---	Down	Forwarding	Disabled	Enable	128	20000	Apply
4	---	Down	Forwarding	Disabled	Enable	128	20000	Apply
5	---	Down	Forwarding	Disabled	Enable	128	20000	Apply
6	---	Down	Forwarding	Disabled	Enable	128	20000	Apply
7	---	Down	Forwarding	Disabled	Enable	128	20000	Apply
8	---	Down	Forwarding	Disabled	Enable	128	20000	Apply

Figure 46. RSTP Basic Port Configuration Page

This page displays the following information about the ports:

Port
Indicates ports 1 through 8. Use the All row to apply the same settings to the STP Status, Priority, and Path Cost fields to ports 1 through 8.

Trunk
Indicates the trunk assignment of a port.

Link Status
Indicates if the port is connected to (Up) or disconnected from (Down) another network device.

Port State
Indicates one of the following port states:

- Blocking— A port that would cause a switching loop. In this state, no user data is sent or received by the port. The port may go into the forwarding state if the other links in use failed and the spanning tree algorithm determines the port may transition to the forwarding state. BPDU data is received in the blocking state.
- Listening— The port processes BPDUs and awaits new information that would cause the port to return to the blocking state.
- Learning— While the port does not yet forward frames (packets), in this state the port does learn source addresses from frames received and adds them to the filtering (switching) database.
- Forwarding— A port that both receives and sends data. This indicates normal operation. STP continues to monitor the port for incoming BPDUs that indicate the port should return to the blocking state to prevent a loop.
- Disabled— This state is not strictly part of STP. However, a network administrator can manually disable a port.

Role
Indicates one of the following port roles:

- Disabled— The Disabled Port role is assigned if the port is not operational or is excluded from the active topology by management or it is a network access port (IEEE Std 802.1X) and it is Unauthorized, or its Administrative Bridge Port state is Disabled.
- Root— If the least cost path to the root is through this port, then it becomes the root port for this bridge.
- Designated— If this is the designated bridge for the LAN and if the root path cost information received on this port is greater than the root port's path cost and less than any other port's received information, then this port becomes the designated port.
- Backup— Any operational Bridge Port that is not a Root or Designated Port is a Backup Port if the Bridge is the Designated Bridge for the attached LAN.
- Alternate— Any operational Bridge Port that is not a Root or a Designated Port is an Alternate Port if that Bridge is *not* the Designated Bridge for the attached LAN.

7. In the STP Status column for the port you want to configure, select the STP status from the list, either Enable or Disable.

8. In the Priority column for the port you want to configure, type a number for the port priority.

   Port priority is described in “Port Priority” on page 136.

9. In the Path Cost column for the port you want to configure, type a number for the Path Cost.

   For STP, the range is from 0 to 65,535. For RSTP, the range is from 0 to 20,000,000. For both protocols, the default value is 128. The Path cost is described in “Path Costs and Port Costs” on page 136.

10. Click **Apply**.

11. To configure all of the ports to the same settings, in the All row, configure one, two, or all of the following settings: STP Status, Priority, and Path Cost. Then click **Apply**.

12. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Configuring the Advanced RSTP Port Settings

To configure the advanced RSTP port settings, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   The Spanning Tree folder is displayed.

2. From the **Bridge** folder, select the **Spanning tree** folder.

3. From the **Spanning tree** folder, select **RSTP Advanced Port** folder.

The RSTP Advanced Port Configuration Page is shown in Figure 47.

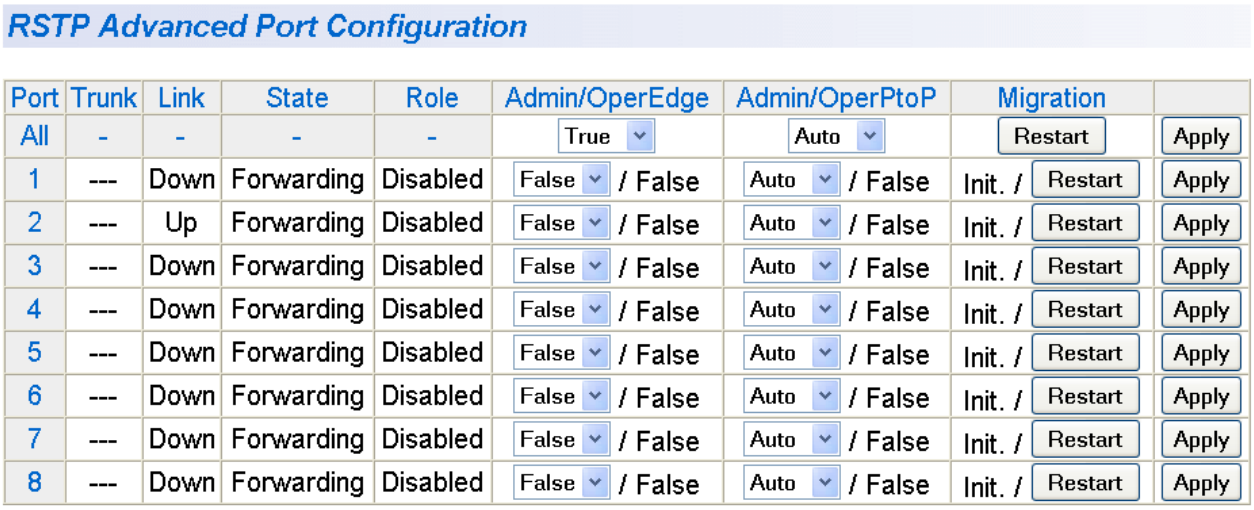

Figure 47. RSTP Advanced Port Configuration Page

This page displays the following information about the ports:

Port
Indicates ports 1 through 8. Use the All row to apply the same settings to the STP Status, Priority, and Path Cost fields to ports 1 through 8.

Trunk
Indicates the trunk assignment of a port.

Link
Indicates if the port is connected to (Up) or disconnected from (Down) another network device.

State
Indicates one of the following port states:

- Blocking— A port that would cause a switching loop. In this state, no user data is sent or received by the port. The port may go into the forwarding state if the other links in use failed and the spanning tree algorithm determines the port may transition to the forwarding state. BPDU data is received in the blocking state.
- Listening— The port processes BPDUs and awaits new information that would cause the port to return to the blocking state.
- Learning— While the port does not yet forward frames (packets), in this state the port does learn source addresses from frames received and adds them to the filtering (switching) database.
- Forwarding— A port that both receives and sends data. This indicates normal operation. STP continues to monitor the port for

incoming BPDUs that indicate the port should return to the blocking state to prevent a loop.

- ❒ Disabled—This state is not strictly part of STP. However, a network administrator can manually disable a port.

Role
Indicates one of the following port roles:

- ❒ Disabled—The Disabled Port role is assigned if the port is not operational or is excluded from the active topology by management or it is a network access port (IEEE Std 802.1X) and it is Unauthorized, or its Administrative Bridge Port state is Disabled.
- ❒ Root— If the least cost path to the root is through this port, then it becomes the root port for this bridge.
- ❒ Designated— If this is the designated bridge for the LAN and if the root path cost information received on this port is greater than the root port's path cost and less than any other port's received information, then this port becomes the designated port.
- ❒ Backup— Any operational Bridge Port that is not a Root or Designated Port is a Backup Port if the Bridge is the Designated Bridge for the attached LAN.
- ❒ Alternate— Any operational Bridge Port that is not a Root or a Designated Port is an Alternate Port if that Bridge is *not* the Designated Bridge for the attached LAN.

4. In the Admin/OperEdge column for the port you want to configure, choose True or False to set whether or not the port will operate as an edge port.

5. In the Admin/OperPtoP column for the port you want to configure, choose a setting based on the information in Table 4.

Table 4. RSTP Point-to-Point Status

Admin	Operation	Port Duplex Operation
Auto	True	Full
	False	Half
True	True	Full or Half
False	False	Full or Half

6. In the Migration column for the port you want to configure, click **Restart** to reset the port.

7. Click **Apply**.

8. To configure all of the ports to the same settings, in the All row, configure one, two, or all of the following settings: Admin/OperEdge, Admin/OperPtoP, and Migration. Then click **Apply**.

9. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Viewing the Spanning Tree Topology

To view the current spanning tree topology, perform the following procedure:

1. From the menu on the left side of the page, select **Bridge**.

   This folder expands.

2. From the **Bridge** folder, select the **Spanning tree** folder.

3. From the **Spanning tree** folder, select **Topology Info**.

   The Designated Topology Information Page is shown in Figure 48.

Designated Topology Information

Port	Trunk	Link Status	Designated Root	Designated Cost	Designated Bridge	Designated Port
1	---	Down	0000 000000000000	0	0000 000000000000	00 00
2	---	Up	0000 000000000000	0	0000 000000000000	00 00
3	---	Down	0000 000000000000	0	0000 000000000000	00 00
4	---	Down	0000 000000000000	0	0000 000000000000	00 00
5	---	Down	0000 000000000000	0	0000 000000000000	00 00
6	---	Down	0000 000000000000	0	0000 000000000000	00 00
7	---	Down	0000 000000000000	0	0000 000000000000	00 00
8	---	Down	0000 000000000000	0	0000 000000000000	00 00

Figure 48. Designated Topology Information Page

This page displays the following information about the ports:

Port
Indicates ports 1 through 8 on the AT-GS590/8POE switch.

Port Trunk
The trunk of which the port is a member.

Link Status
Whether the link on the port is up or down.

Designated Root
The designated root bridge to which the switch's root port is actively connected.

Designated Cost
The sum of all the root port costs on all bridges, including the switch, between the switch and the root bridge.

Designated Bridge
An adjacent bridge to which the root port of the switch is actively connected.

Designated Port
The root bridge to which the root port of the switch is actively connected.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 14

802.1x Port-based Network Access Control

This chapter contains information about the 802.1x Port-based Network Access Control feature. This chapter includes the following sections:

Note
To activate 802.1x port authentication, you must also configure the RADIUS feature. See Chapter 15, “RADIUS Authentication Protocol” on page 169.

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

802.1x Port-based Network Access Control (IEEE 802.1x) is used to control who can send traffic through and receive traffic from a switch port. With this feature, the switch does not allow an end node to send or receive traffic through a port until the user of the node logs on by entering a username and password.

This feature can prevent an unauthorized individual from connecting a computer to a port or using an unattended workstation to access your network resources. Only those users to whom you have assigned a username and password are able to use the switch to access the network.

This feature must be used with the RADIUS authentication protocol and requires that there is a RADIUS server on your network. The RADIUS server performs the authentication of the username and password combinations.

Note
RADIUS with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server for this feature.

Following are several terms to keep in mind when using this feature.

- Supplicant - A supplicant is an end user or end node that wants to access the network through a switch port. A supplicant is also referred to as a client.
- Authenticator - The authenticator is a port on the switch that prohibits network access by a supplicant until the network user has entered a valid username and password.
- Authentication server - The authentication server is the network device that has the RADIUS server software installed. This is the device that does the actual authenticating of the user names and passwords from the supplicants.

The AT-GS950/8POE switch does not authenticate the usernames and passwords from the end users. Rather, the switch acts as an intermediary between a supplicant and the authentication server during the authentication process.

Authentication Process

Below is a brief overview of the authentication process that occurs between a supplicant, authenticator, and authentication server. For further details, refer to the IEEE 802.1x standard.

- Either the authenticator (that is, a switch port) or the supplicant can initiate an authentication prompt exchange. The switch initiates an exchange when it detects a change in the status of a port (such as when the port transitions from no link to valid link), or if it receives a packet on the port with a source MAC address not in the MAC address table.
- An authenticator starts the exchange by sending an EAP-Request/Identity packet. A supplicant starts the exchange with an EAPOL-Start packet, to which the authenticator responds with a EAP-Request/Identity packet.
- The supplicant responds with an EAP-Response/Identity packet to the authentication server via the authenticator.
- The authentication server responds with an EAP-Request packet to the supplicant via the authenticator.
- The supplicant responds with an EAP-Response/MDS packet containing a username and password.
- The authentication server sends either an EAP-Success packet or EAP-Reject packet to the supplicant.
- Upon successful authorization of the supplicant by the authentication server, the switch adds the supplicant's MAC address to the MAC address as an authorized address and begins forwarding network traffic to and from the port.
- When the supplicant sends an EAPOL-Logoff prompt, the switch removes the supplicant's MAC address from the MAC address table, preventing the supplicant from sending or receiving any further traffic from the port.

Authenticator Ports

All of the ports on the AT-GS950/8POE switch are authenticator ports. An authenticator port can have one of three settings. These settings are referred to as the port control settings. The settings are:

- Auto - Activates 802.1x port-based authentication. An authenticator port with this setting does not forward network traffic to or from the end node until the client has entered a username and password that the authentication server must validate. The port begins in the unauthorized state, sending and receiving only EAPOL frames. All other frames, including multicast and broadcast frames, are discarded. The authentication process begins when the link state of the port changes or the port receives an EAPOL-Start packet from a supplicant. The switch requests the identity of the client and begins relaying authentication prompts between the client and the authentication server. Each client that attempts to access the network is uniquely identified by the switch using the client's MAC address.

- Force-unauthorized - Places the port in the unauthorized state, ignoring all attempts by the client to authenticate. This port control setting blocks all users from accessing the network through the port and is similar to disabling a port and can be used to secure a port from use. The port continues to forward EAPOL packets, but discards all other packets, including multicast and broadcast packets.
- Force-authorized - Disables IEEE 802.1x port-based authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1x-based authentication of the client. This is the default setting. Use this port control setting for those ports that are connected to network devices that are not to be authenticated.

Figure 49 illustrates the concept of the authenticator port control settings.

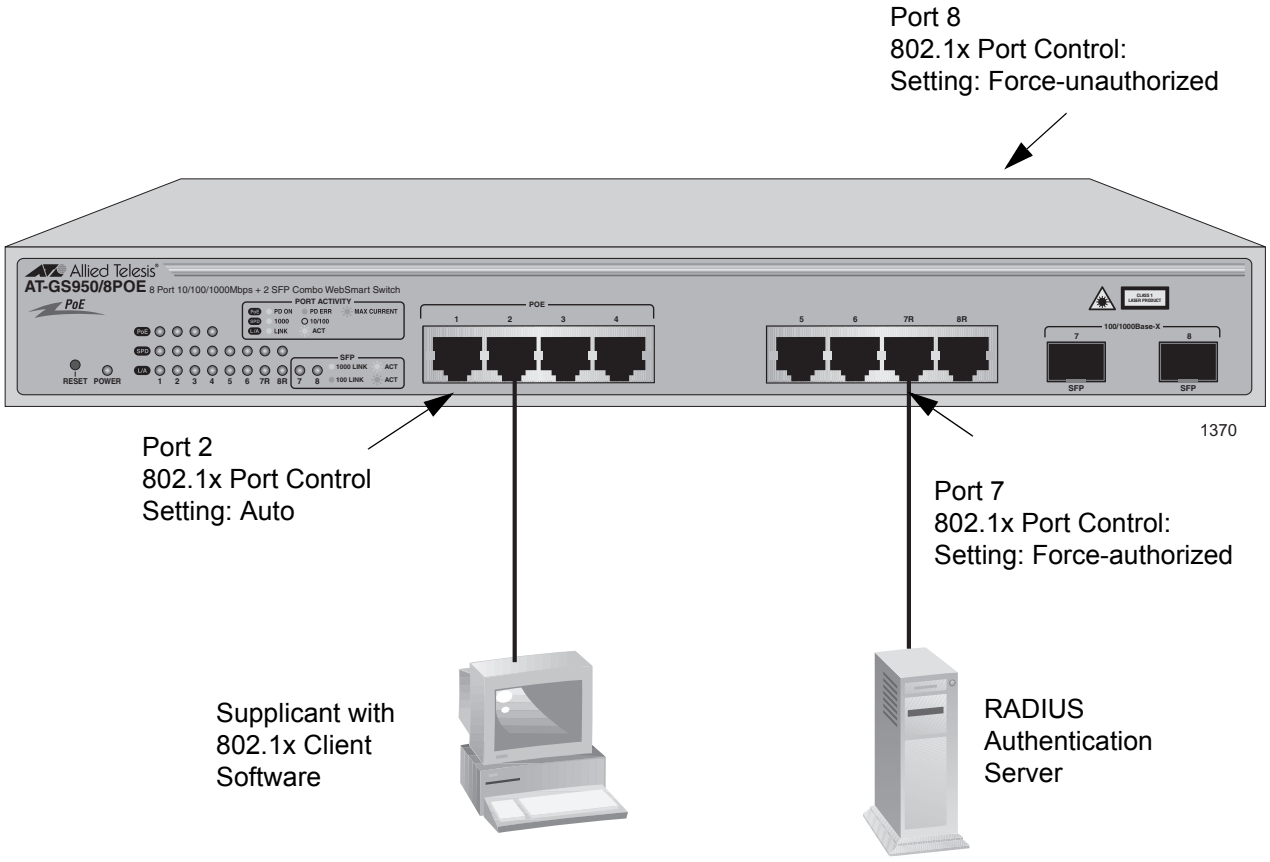

Figure 49. Example of the Authenticator Role

- Port 2 is set to Auto. The end node connected to the port must use its 802.1x client software and provide a username and password to send or receive traffic from the switch.
- Port 7 is set to the Force-authorized setting so that the end node connected to the port does not have to provide a user name or password to send or receive traffic from the switch. In the example, the

node is the RADIUS authentication server. Since the server cannot authenticate itself, its port must be set to Force-authorized in order for it to pass traffic through the port.

- Port 8 is set to Force-unauthorized to prevent anyone from using the port.

As mentioned earlier, the switch does not authenticate the user names and passwords from the clients. That is the responsibility of the authentication server, which contains the RADIUS server software. Instead, a switch acts as an intermediary for the authentication server by denying access to the network by the client until the client has provided a valid username and password, which the authentication server validates.

General Steps

Following are the general steps to implementing 802.1x Port-based Network Access Control:

1. You must install RADIUS server software on one or more of your network servers or management stations. Authentication protocol server software is not available from Allied Telesis. Consult the vendor's documentation for server installation instructions.

2. Install 802.1x client software on those workstations that will act as supplicants.

3. You must configure and activate the RADIUS client software in the AT-S101 Management Software. The default setting for the authentication protocol is disabled. You need to provide the following information:

   - The IP address of a RADIUS servers.
   - The encryption key used by the authentication server.

   For instructions, refer to Chapter 18, "RADIUS Authentication Protocol" on page 207.

4. Configure the authenticator port settings, as explained in "Configuring 802.1x Port-based Network Access Control" on page 165 in this chapter.

Port-based Network Access Control Guidelines

Following are the guidelines for using this feature:

- Ports set to Auto do not support port trunking or dynamic MAC address learning.
- The appropriate setting for a port on an AT-GS950/8POE switch connected to an authentication server is Force-authorized, the default setting. This is because an authentication server cannot authenticate itself.

- The authentication server must be a member of the Default VLAN by communicating with the switch through a port that is an untagged member of the Default VLAN.
- Allied Telesis does not support connecting more than one supplicant to an authenticator port on the switch. The switch allows only one supplicant to log on per port.

Note
Connecting multiple supplicants to a port set to the Auto setting does not conform to the IEEE 802.1x standard. This can introduce security risks and can result in undesirable switch behavior. To avoid this, Allied Telesis recommends use the Force-authorized setting of the Port Control feature on ports that are connected to more than one end node, such as a port connected to another switch or to a hub.

- A username and password combination is not tied to the MAC address of an end node. This allows end users to use the same username and password when working at different workstations.
- After a supplicant has successfully logged on, the MAC address of the end node is added to the switch's MAC address table as an authenticated address. It remains in the table until the end user logs off the network. The address is not timed out, even if the end node becomes inactive.

Note
End users of port-based access control should be instructed to always log off when they are finished with a work session. This prevents unauthorized individuals from accessing the network through unattended network workstations.

- There should be only one port in the authenticator port control setting of Auto between a client and the authentication server.
- Ports used to interconnect switches should be set to the port control setting of Force-authorized. This is illustrated in Figure 50 on page 163.

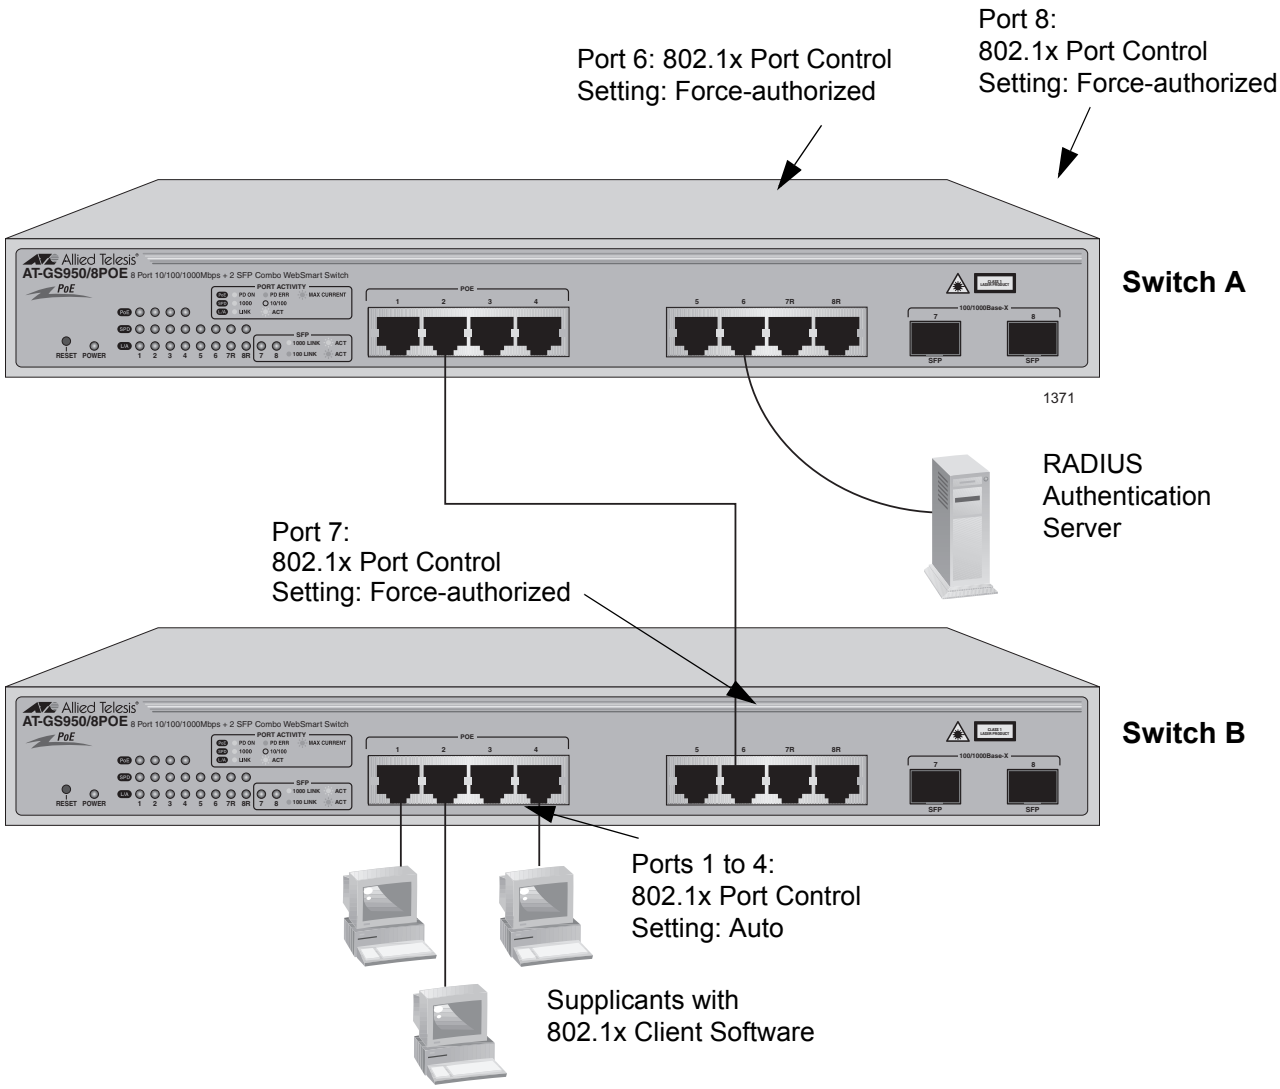

Figure 50. Port-based Authentication Across Multiple Switches

Guest VLANs

An authenticator port in the unauthorized state typically accepts and transmits only 802.1x packets while waiting to authenticate a supplicant. However, you can configure an authenticator port to be a member of a Guest VLAN when no supplicant is logged on. Any client using the port is not required to log on and has full access to the resources of the Guest VLAN.

If the switch receives 802.1x packets on the port, signalling that a supplicant is logging on, it moves the port to its predefined VLAN and places it in the unauthorized state. The port remains in the unauthorized state until the log on process between the supplicant and the RADIUS server is completed. When the supplicant logs off, the port automatically returns to the Guest VLAN.

Note
The Guest VLAN feature is only supported on an authenticator port in the single operating mode.

Configuring 802.1x Port-based Network Access Control

To configure 802.1x port-based network access control, perform the following procedure:

1. Select the **Security** folder from the menu on the left side of the page.

   The **Security** folder expands.

2. From the **Security** folder, select **Port Access Control**.

   The 802.1x Access Control Configuration Page is shown in Figure 51.

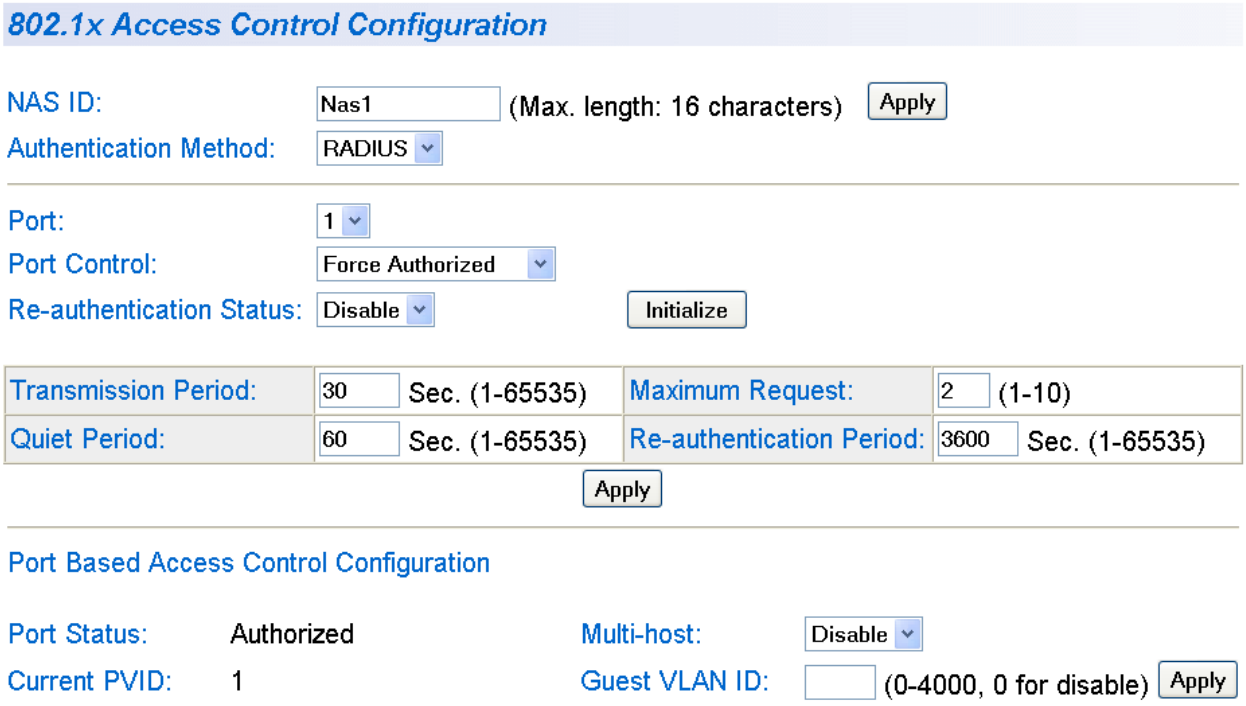

Figure 51. 802.1x Access Control Configuration Page

3. To select a port, click **Port** and select the port you want to configure from the pull-down menu. You can configure only one port at a time.

   The current settings for the selected port are displayed.

4. Configure the following parameters as needed. The parameters are defined here:

   NAS ID
   This parameter assigns an 802.1x identifier to the switch that applies to all ports. The NAS ID can be up to sixteen characters. Valid characters are 0 to 9, a to z, and A to Z. Spaces are allowed. Specifying an NAS ID is optional.

   Authentication Method
   This parameter indicates the authentication method used by the switch. The options are RADIUS or local. The default setting is RADIUS.

   Port Control
   Sets the 802.1x port control setting. The possible settings are:

   Auto - Enables 802.1x port-based authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port changes or the port receives an EAPOL-Start packet from a supplicant. The switch requests the identity of the client and begins relaying authentication prompts between the client and the authentication server.

   Force-unauthorized - Places the port in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.

   Force-authorized - Disables IEEE 802.1x port-based authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1x-based authentication of the client. This is the default setting

   Re-authentication Status
   Specifies if reauthentication should occur according to the reauthentication period. The options are Enabled or Disabled.

   Transmission Period
   Sets the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before retransmitting the request. The default value is 30 seconds. The range is 1 to 65,535 seconds.

   Initialize
   Pressing this button ends the 802.1x session and connectivity is lost during re-authentication on an 802.1x enabled port. In addition, the value of the **Port Status** parameter is changed to “Unauthorized” if you press the Initialize button on an 802.1x enabled port.

Quiet Period
Sets the number of seconds that the port remains in the quiet state following a failed authentication exchange with the client. The default value is 60 seconds. The range is 0 to 65,535 seconds.

Maximum Request
Sets the maximum number of times that the switch retransmits an EAP Request packet to the client before it times out the authentication session. The default value for this parameter is 2 retransmissions. The range is 1 to 10 retransmissions.

Re-authentication Period
Specifies the time period between periodic reauthentication of the client. The default value is 3600 seconds. The range is 1 to 65,535 seconds.

Port Status
Displays the current 802.1 status of the port as either authorized or unauthorized. This is not an adjustable parameter. See the description of the **Initialize** parameter.

Current PVID
Displays the Port VLAN identifier (PVID) of the port. This is not an adjustable parameter. For more information about this field, see "Port VLAN Identifier" on page 49.

Multi-host
Enables multiple hosts to a single 802.1x enabled port. The options are Enable or Disable. The default setting is Disable.

The Disable setting is appropriate when there is only one host node connected to each port on the switch. This setting causes the switch to immediately stop sending multicast packets from a switch port when a host node signals its desire to leave a multicast group by sending a leave request or when the host node stops sending reports and times out. The switch forwards the leave request to the router and simultaneously ceases transmission of any further multicast packets from the port where the host node is connected.

The Enable setting is appropriate if there is more than one host node connected to a switch port, such as when a port is connected to an Ethernet hub to which multiple host nodes are connected. With this setting selected, the switch continues sending multicast packets from a port even after it receives a leave request from a host node on the port. This ensures that the remaining active host nodes on the port continue to receive the multicast packets. Only after all of the host nodes connected to a switch port have transmitted leave requests (or have timed out) does the switch stop sending multicast packets from the port.

Guest VLAN ID
Specifies the guest VLAN ID. This feature is only supported on an

authentication port in the single operating mode. Choose a value between 0 and 4,000. Then click **Apply**. There is no default value. For more information, see “Guest VLANs” on page 164.

5. When you are finished configuring the parameters, click **Apply** at the bottom of the 802.1x Configuration page.

6. If the port control setting is Auto and you want to return the EAPOL machine state on the port to the initialized state, select **Yes** for the Initialize parameter and click **Apply**.

7. If the port control setting is Auto and you want the node connected to the port to reauthenticate with the RADIUS server, select **Yes** for the Re-auth Initialize parameter and click **Apply**.

8. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 15

RADIUS Authentication Protocol

This chapter explains how to configure the RADIUS client on the switch. You can use the RADIUS client with 802.1x port-based network access control to control who can forward packets through the switch. This chapter contains the following sections:

Note
To activate the RADIUS feature, you must also configure the 802.1x port-network access control feature. See Chapter 14, “802.1x Port-based Network Access Control” on page 157.

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

RADIUS (Remote Authentication Dial In User Services) is an authentication protocol for enhancing the security of your network. The protocol transfers the task of authenticating network access from a network device to an authentication protocol server.

The AT-S101 Management Software comes with RADIUS client software. You can use the client software together with 802.1x port-based network access control. See Chapter 17, “802.1x Port-based Network Access Control” on page 191, to control which end users and end nodes can send packets through the switch.

RADIUS Implementation Guidelines

What do you need to use the RADIUS protocol? Following are the main points.

- You must install RADIUS server software on a network server or management station. Authentication protocol server software is not available from Allied Telesis.
- The RADIUS server must communicate with the switch through a port that is an untagged member of the Default VLAN.
- If the RADIUS server is on a different subnet from switch, be sure to specify a default gateway in the IP Setup Page, so that the switch and server can communicate with each other. See “Configuring an IP Address, Subnet Mask and Gateway Address” on page 22.
- You need to configure the RADIUS server software on the authentication server by specifying the username and password combinations. The maximum length of a username or password is 12 alphanumeric characters.

  Note
  This manual does not explain how to configure RADIUS server software. Refer to the documentation that came with the software for instructions.

- You must activate the RADIUS client software on the switch using the AT-S101 Management Software and configure the settings. This is explained in “Configuring the RADIUS Client” on page 171. By default, authentication protocol is disabled.

  Note
  For more information on the RADIUS authentication protocol, refer to the RFC 2865 standard.

Configuring the RADIUS Client

To configure the RADIUS client, perform the following procedure:

1. From the menu on the left side of the page, select the **Security** folder.

   The Security folder expands.

2. From the **Security** folder, select **RADIUS**.

   The RADIUS Page is shown in Figure 52.

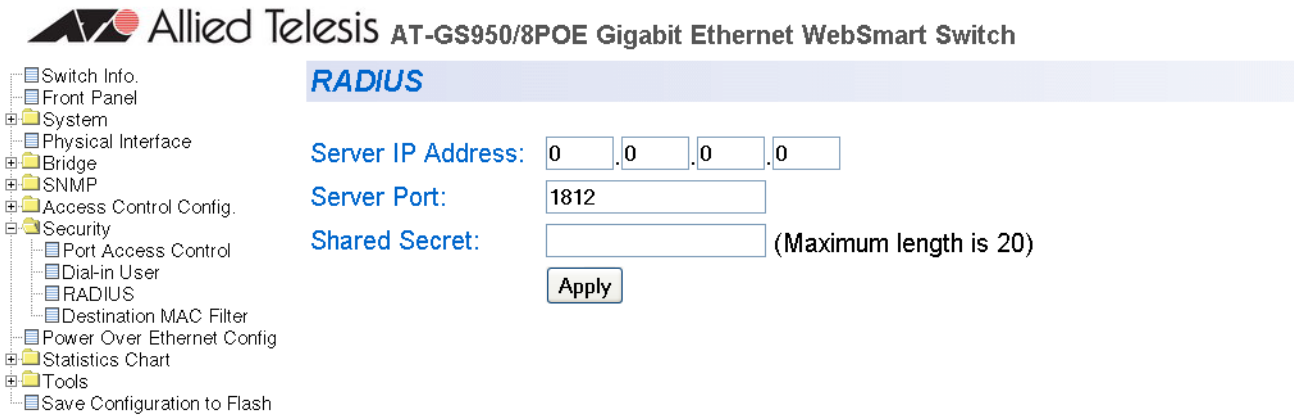

Figure 52. RADIUS Page

3. To enter the RADIUS server's IP address, enter the address in the **Server IP Address** field.

4. To specify the server's encryption key, click the **Shared Secret** field and enter the encryption key.

5. To select the port number that you want to assign to UDP, type in the port number in the **Server Port** field.

   You may only assign one port number to this parameter. The default value is 1812.

6. Click **Apply** to save your changes.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 16

Destination MAC Filter

This chapter contains an explanation of the Destination MAC Filter feature as well a procedure for configuring it. This chapter includes the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

The Destination MAC Filter feature prevents the AT-GS950/8POE switch from forwarding packets to a specified device. On the Destination MAC Filter Page of the AT-S101 software, you enter the MAC address of the device that you want to filter.

After the switch receives a packet, it examines the destination MAC address of the packet. If the destination MAC address matches a MAC address set in the filter, then the software prevents the switch from forwarding it. Instead, the switch drops the packet.

You may want to block access to a device within your organization. For instance, you may not want users on the Sales group switch to have access to a server on the Accounting group switch. You can enter the MAC address of the Accounting server as a destination MAC address filter on the Sales group switch. When a packet destined for the Accounting server is received by the Sales group switch, the switch drops the packet.

The Destination MAC Filter is a subset of the static MAC address. For more information about MAC addresses, see Chapter 12, “Static Multicast MAC Address” on page 127.

Configuring a Destination MAC Filter

To set MAC address in the Destination MAC Filter, perform the following procedure:

1. From the menu on the left side of the page, select the **Security** folder.

   The **Security** folder expands.

2. From the **Security** folder, select **Destination MAC Filter**.

   The Destination MAC Filter Page is shown in Figure 53.

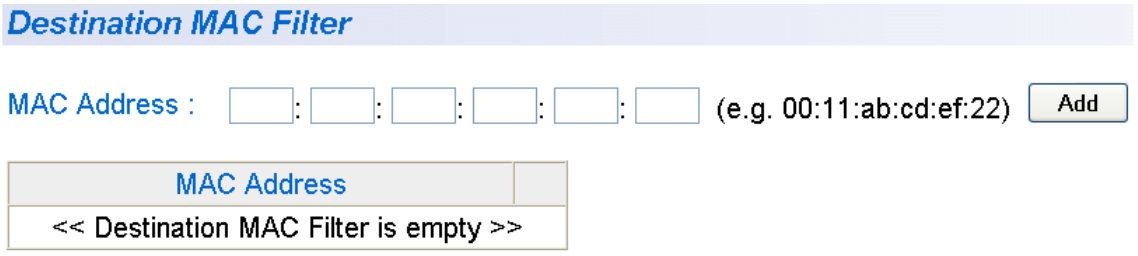

Figure 53. Destination MAC Filter Page

3. After you have configured a destination MAC address, the Destination MAC Filter Page is updated with the MAC address. See Figure 54.

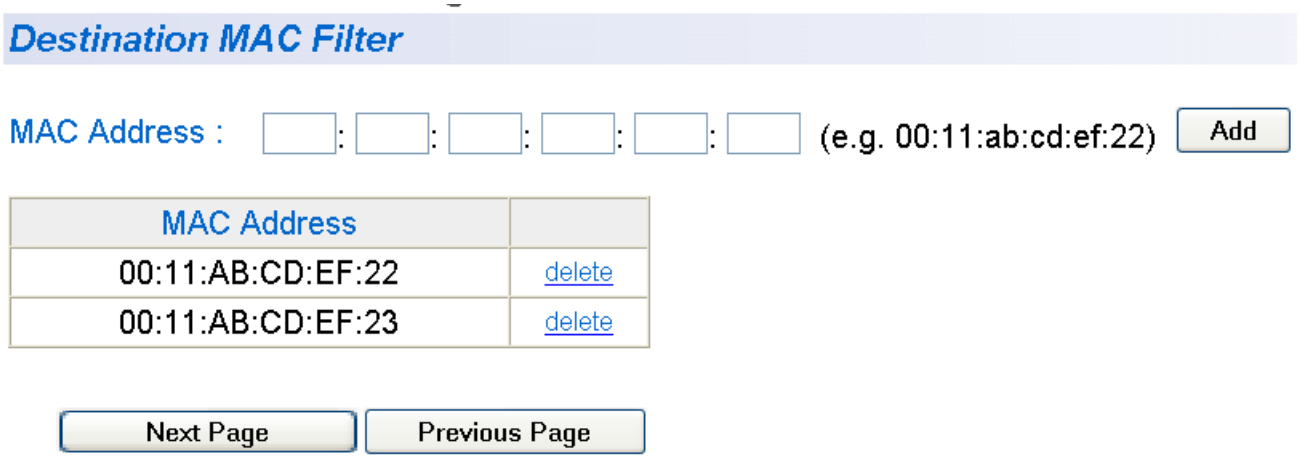

Figure 54. Updated Destination MAC Filter Page

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Deleting a Destination MAC Filter

To delete a MAC address from the Destination MAC Filter, perform the following procedure:

1. From the menu on the left side of the page, select the **Security** folder.

   The **Security** folder expands.

2. From the **Security** folder**,** select **Destination MAC Filter**.

   The Destination MAC Filter Page is shown in Figure 53 on page 175

3. Select **delete** next to the MAC address that you want to delete.

   The MAC address is removed from the MAC address table.

4. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 17

Power over Ethernet (PoE)

This chapter contains a description of the Power over Ethernet (PoE) feature and the procedure for configuring it. This chapter includes the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

The four of the twisted pair ports on the AT-GS950/8POE switch feature Power over Ethernet (PoE) which is a mechanism for supplying power to network devices over the same twisted pair cables used to carry network traffic. This feature can simplify network installation and maintenance by allowing you to use the switch as a central power source for other network devices.

A device that receives its power over an Ethernet cable is called a *powered device*. Examples of such devices can be wireless access points, IP telephones, web cams, and even other Ethernet switches. A powered device connected to a port on the switch receives both network traffic and power over the same twisted pair cable. In order to use the PoE feature, you must connect a powered device to one of the four PoE ports on the switch.

There are several advantages that the PoE feature of the AT-GS950/8POE switch adds to the installation and maintenance of your network. First of all, the switch acts as the central power source for your powered devices. Adding an uninterruptible power source (UPS) to the switch increases the protection to the switch from possible power source problems and to all of the powered devices connected to the switch. In addition, the PoE feature can increase the reliability of your network by minimizing the impact to network operations from a power failure.

A port on the switch connected to a powered device can supply up to 15.4 Watts of power to the device and, simultaneously, furnishing standard 10/100 Mbps Ethernet functionality. A PoE port that is connected to a network node that is not a powered device (that is, a device that receives its power from another power source) functions as a regular Ethernet port, without PoE. In this case, the PoE feature remains enabled on the port, but no power is delivered to the device.

Power Budgeting

The AT-GS950/8POE switch provides a maximum of 15.4 W of power per port on four of the eight ports for a total power consumption of 62 W, while at the same time furnishing standard 10/100 Mbps Ethernet functionality.

The switch power management functionality supports any combination of Ethernet ports (1-4) that supply power for IEEE 802.3af Class 0, 1, 2, or 3 powered devices up to a maximum of 62 watts, as described in Table 5.

Note
Power is supplied to the powered devices in the order that the ports are connected or on a first-come-first-served basis until the 62 watt limit is reached. If the switch is power cycled after the PoE devices are connected to the switch ports, the power is supplied to ports 1 through 4 in ascending order.

Table 5. IEEE 802.3af Class vs. Power Levels

Class	Usage	Minimum Power Levels Output at the PSE	Maximum Power Levels Output at the PD
0	Default	15.4W	0.44W to 12.95W
1	Optional	4.0W	0.44W to 3.84W
2	Optional	7.0W	3.84W to 6.49W
3	Optional	15.4W	6.49W to 12.95W

Port Power Priority

When the power budget for the switch is reached, each port is allotted power according to its priority level. You can set the power priority of each PoE port. There are three levels: Low, High, and Critical. By default, all of the PoE ports are set to Low. Naturally, you want to assign a Critical level to ports that are connected to the most important powered devices.

If you set all of the priority levels to the same value (or maintain the default value of Low on all ports), then the port whose port number is the lowest has the highest port power priority. For instance, if you assign PoE ports 1 through 4 with a priority level of High and the power budget is reached, then port 1 has the highest power priority level followed by port 2, etc.

Setting Power over Ethernet

This procedure explains how to set the POE feature on the AT-GS950/8POE switch.

To set the POE feature, perform the following procedure:

1. From the menu on the left side of the page, select the **Power Over Ethernet** Page.

   The Power Over Ethernet Configuration Page is shown in Figure 55.

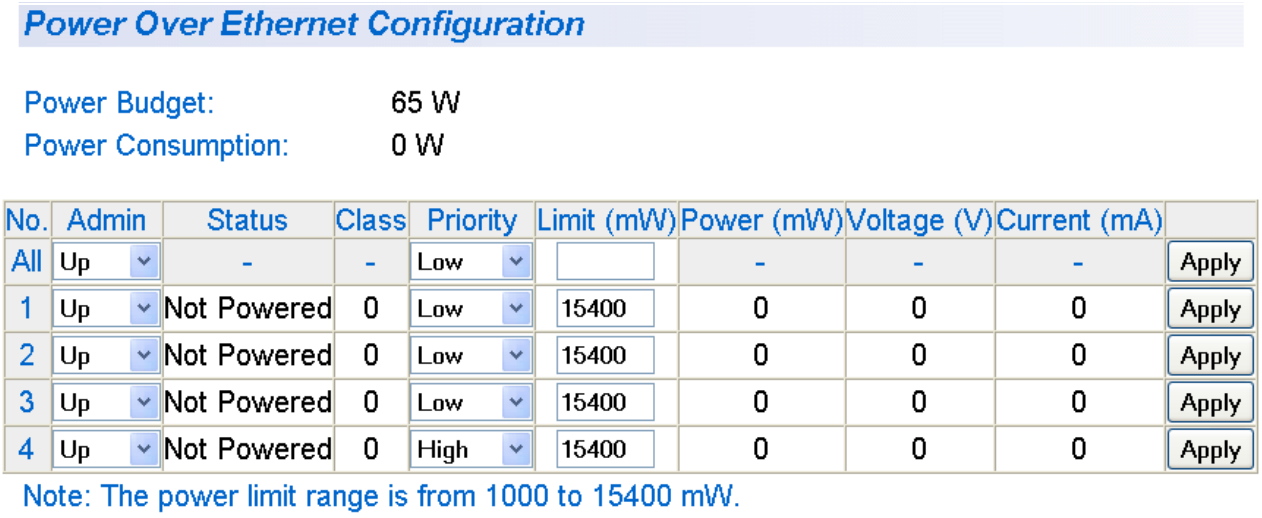

Figure 55. Power Over Ethernet Configuration Page

2. To change the setting of the **Admin** field, select the pull-down menu next to ports 1 through 4. Choose from the following options:

   Up
   Indicates the PoE port is available to be connected to a powered device. This is the default setting.

   Down
   Indicates the PoE port is not available to be connected to a powered device.

   Then click **Apply**. You can use the option next to the ALL row to set all of the ports to the same setting.

3. To change the setting of the **Priority** field, use the pull-down menu to select one of the following power priority values:

   - Low - Indicates the power priority of the PoE port is low. This is the default value.
   - High - Indicates the power priority of the PoE port is high.
   - Critical - Indicates the power priority of the PoE port is critical.

   Then click **Apply**. You can use the option next to the ALL row to set all of the ports to the same setting.

   Note
   For more information about setting this field, see "Port Power Priority" on page 181.

4. To select the amount of power a PoE port can received, enter a value in the **Limit** field. The default value is 15400 mW. The range is from 1000 to 15400 mW.

   The following fields are for display only:

   Power Budget
   Displays the maximum power budget of PoE ports 1 through 4.

   Power Consumption
   Displays the current power consumption of PoE ports 1 through 4.

   Status
   The Status field displays the POE status of each port. The status is either Not Powered (off) or Powered (on).

   Class
   The Class field displays IEEE 802.3af class assignment of a port. See "Power Budgeting" on page 181 for more information. The default value is Class 0.

   Power
   The Power field displays the power consumption of each PoE port.

   Voltage
   The Voltage field displays the voltage of each PoE port.

   Current
   The Current field displays the current of each PoE port.

5. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 18

Classifiers

This chapter explains the concepts of classifiers which are used to describe traffic flow for Access Control Policies. A procedure for configuring classifiers is also provided.

This chapter contains the following sections:

Note
For information about Access Control Policies, see Chapter 19, “Access Control Policies” on page 195.

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

A classifier defines a *traffic flow* which consists of packets that share one or more characteristics. You can define a traffic flow can broadly or narrowly. An example of the a broad definition is all IP traffic while an example of a narrow definition is packets with specified source and destination MAC addresses.

A classifier contains a set of criteria for defining a traffic flow. Examples of the variables include source and destination MAC addresses, source and destination IP addresses, IP protocols, source and destination TCP and UDP ports numbers, and so on. You can also specify more than one criteria within a classifier to make the definition of the traffic flow more specific. Some of the variables you can mix-and-match, but there are restrictions, as explained later in this section in the descriptions of the individual variables.

By itself, a classifier does not perform any action or produce any result because it lacks instructions on what a port should do when it receives a packet that belongs to the defined traffic flow. Rather, the action is established outside the classifier. As a result, you never use a classifier by itself.

The switch uses classifiers to help define Access Control Policies (ACP).

An ACP filters ingress packets on a port by controlling which packets a port accepts and rejects. You can use this feature to improve the security of your switch or enhance switch performance by changing the precedence.

When you create an ACP you must specify the traffic flow you want the ACP to control. You do that by creating one or more classifiers and adding the classifiers to the ACP. The action that the port takes when an ingress packet matches the traffic flow specified by a classifier is contained in the ACP itself. The action is to either accept packets of the traffic flow or discard them.

In summary, a classifier is a list of variables that define a traffic flow. ACP uses a classifier to determine which packets it will manipulate.

Classifier Criteria

The components of a classifier are defined in the following subsections.

Source MAC Address
Destination MAC Address

You can identify a traffic flow by specifying a source and/or destination MAC address. For instance, you might create a classifier for a traffic flow destined to a particular destination node, or from a specific source node to a specific destination node, all identified by their MAC addresses.

The management software does not support a classifier based on a range of MAC addresses. Different MAC addresses must be considered separate traffic flows, with their own classifiers.

Ethernet 802.2 and Ethernet II Frame Types

You can create a classifier that filters packets based on Ethernet frame type and whether a packet is tagged or untagged within a frame type. (A tagged Ethernet frame contains within it a field that specifies the ID number of the VLAN to which the frame belongs. Untagged packets lack this field.) Options are:

- Ethernet II tagged packets
- Ethernet II untagged packets
- Ethernet 802.2 tagged packets
- Ethernet 802.2 untagged packets

802.1p Priority Level

A tagged Ethernet frame, as explained in “Tagged VLAN Overview” on page 48, contains within it a field that specifies its VLAN membership. Such frames also contain a user priority level used by the switch to determine the Quality of Service to apply to the frame and which egress queue on the egress port a packet should be stored in. The three bit binary number represents eight priority levels, 0 to 7, with 0 the lowest priority and 7 the highest. Figure 56 on page 188 illustrates the location of the user priority field within an Ethernet frame.

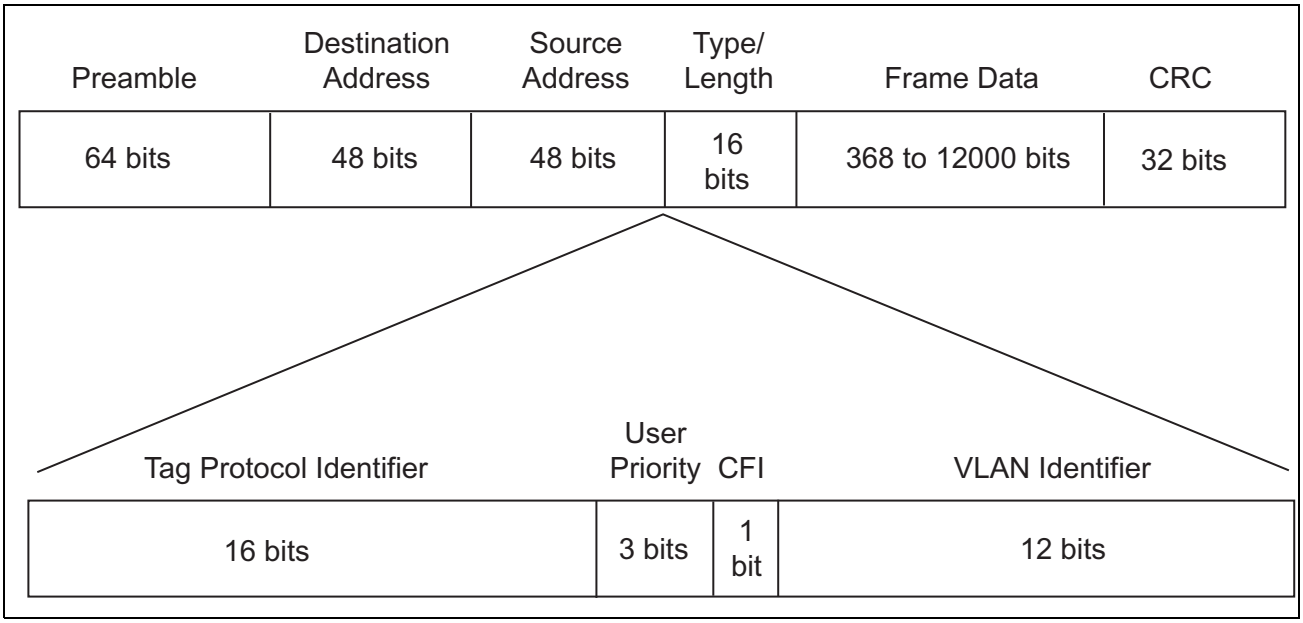

Figure 56. User Priority and VLAN Fields within an Ethernet Frame

You can identify a traffic flow of tagged packets using the user priority value. A classifier for this type of traffic flow instructs a port to watch for tagged packets containing the specified user priority level.

The priority level criteria can contain only one value, and the value must be from 0 (zero) to 7. Multiple classifiers are required if a port is to watch for several different traffic flows of different priority levels.

VLAN ID

A tagged Ethernet frame also contains within it a field of 12 bits that specifies the ID number of the VLAN to which the frame belongs. The field, illustrated in Figure 56, can be used to identify a traffic flow.

A classifier can contain only one VLAN ID. To create a port policy that applies to several different VLAN IDs, multiple classifiers are required.

Protocol (Layer 2)

Traffic flows can be identified by the protocol specified in the Ethertype field of the MAC header in an Ethernet II frame. Possible values are:

- IP
- ARP
- RARP
- Protocol Number

Observe the following guidelines when using this variable:

- When selecting a Layer 4 variable, this variable must be left blank or set to IP.
- If you choose to specify a protocol by its number, you can enter the value in decimal or hexadecimal format. If you choose the latter, precede the number with the prefix “0x.”
- The range for the protocol number is from 1536 (0x600) to 65535 (0xFFFF).

DSCP (DiffServ Code Point)

The Differentiated Services Code Point (DSCP) tag indicates the class of service to which packets belong. The DSCP value is written into the ToS field of the IP header, as shown in Figure 57. Routers within the network use this DSCP value to classify packets and assign QoS appropriately.

The location of this value is shown in Figure 57.

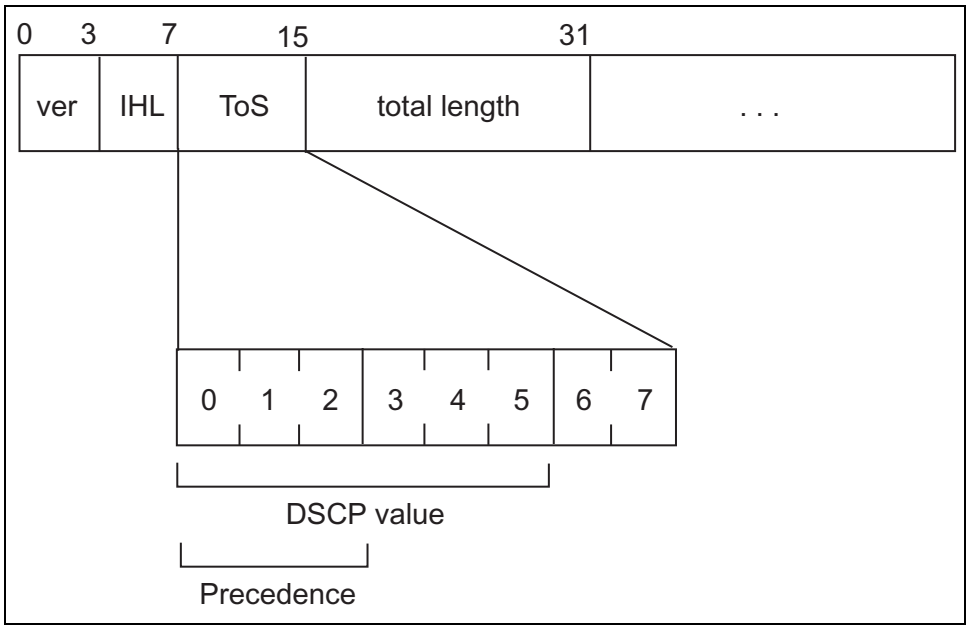

Figure 57. DSCP value in an IP Header

Protocol

You can define a traffic flow by the following Layer 3 protocols:

- TCP
- UDP
- ICMP

- ❒ IGMP
- ❒ IP protocol number

If you choose to specify the protocol by its number, you can enter the value in decimal or hexadecimal format. It the latter, include the prefix “0x.” The range for the protocol number is 0 (0x0) to 255 (0xFF).

Source IP Addresses (Layer 3)
Source IP Mask (Layer 3)

You can define a traffic flow by the source IP address contained in IP packets. The address can be of a subnet or a specific end node.

You do not need to enter a source IP mask if you are filtering on the IP address of a specific end node. A mask is required, however, when you filter on a subnet. A binary “1” indicates the switch should filter on the corresponding bit of the IP address, while a “0” indicates that it should not. For example, the subnet address 149.11.11.0 would have the mask “255.255.255.0.”

Observe this guideline when using these criteria:

- ❒ The Protocol variable must be left blank or set to IP.

Destination IP Addresses (Layer 3)
Destination IP Mask (Layer 3)

You can also define a traffic flow based on the destination IP address of a subnet or a specific end node.

You do not need to enter a destination IP mask for an IP address of a specific end node. A mask is required, however, when filtering on a subnet. A binary “1” indicates the switch should filter on the corresponding bit of the IP address while a “0” indicates that it should not. For example, the subnet address 149.11.11.0 would have the mask “255.255.255.0.”

Observe this guideline when using these criteria:

- ❒ The Protocol variable must be left blank or set to IP.

Source Ports (Layer 4)
Destination Ports (Layer 4)

A traffic flow can be identified by a source and/or destination TCP or UDP port number contained within the header of an IP frame. Observe the following guidelines when using these criteria:

- ❒ The Protocol variable must be left blank or set to IP.
- ❒ A classifier cannot contain criteria for both TCP and UDP ports. You may specify either TCP or UDP ports in a classifier, but not both.

Guidelines

Follow these guidelines when creating a classifier:

- Each classifier represents a separate traffic flow.
- The variables within a classifier are linked by AND. The more variables defined within a classifier, the more specific it becomes in terms of the flow it defines. For instance, specifying both a source IP address and a TCP destination port within the same classifier defines a traffic flow that relates to IP packets containing both the designated source IP address and TCP destination port. However, there are some restrictions on combining variables in the same classifier. See “Classifier Criteria” on page 187.
- You can apply the same classifier to more than one access control policy.
- A classifier without any defined variables applies to all packets.
- You cannot create two classifiers that have the same settings. There can be only one classifier for any given type of traffic flow.
- The switch can store up to 256 classifiers. However, the maximum number of classifiers that you can assign to active access control policies at any one time is 64.

Creating Classifiers

To create a classifier, perform the following procedure:

1. From the menu on the left side of the page, select the **Access Control Configuration** folder.

   The **Access Control Configuration** folder expands.

2. From the **Access Control Configuration** folder, select **Classifier**.

   The Create Classifier Page is displayed. See Figure 58

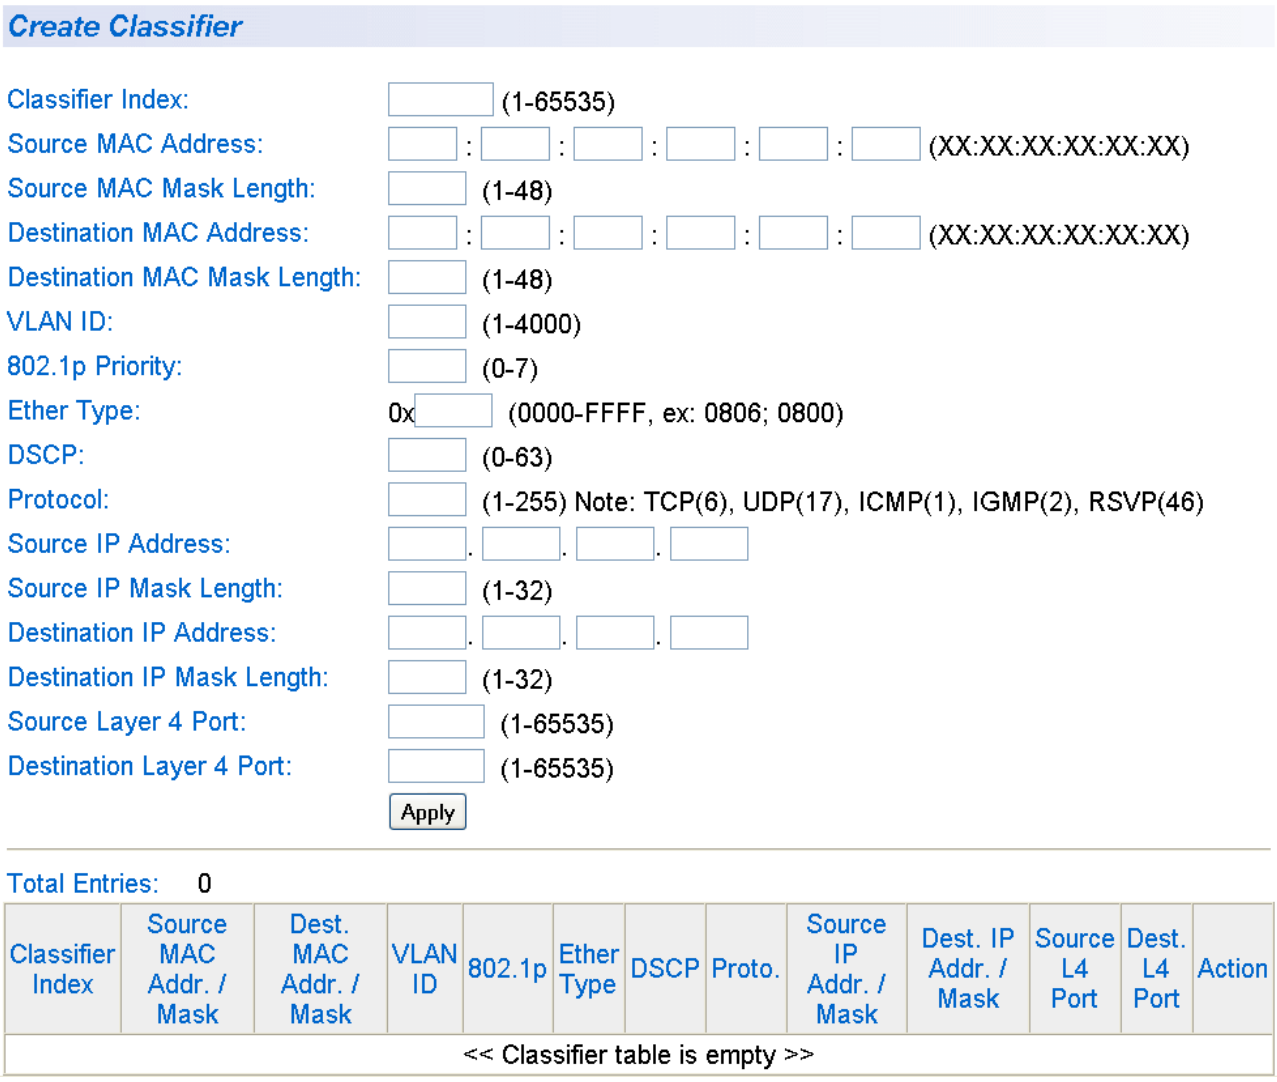

Figure 58. Create Classifier Page

3. Input the following fields:

Classifier Index
Indicates the identification number of the classifier. You must assign each classifier a unique index.

Source MAC Address
Indicates the source MAC address of the classifier.

Source MAC Mask Length
Indicates the source MAC mask length of the classifier.

Destination MAC Address
Indicates the destination MAC address of the classifier.

Destination MAC Mask Length
Indicates the destination MAC mask length of the classifier.

VLAN ID
Specifies the VLAN ID of the VLAN to which the frame belongs. Each classifier can contain only one VLAN ID. To create a policy that applies to several different VLAN IDs, you must define multiple classifiers.

802.1p Priority
Specify the 802.1p priority level. Enter a value between 0 and 7 with 0 as the lowest priority level and 7 as the highest priority level.

Ether Type
Specify an Ethernet frame for filtering. For example, to filter an ARP packet, type” 0806.”

DSCP
Indicates the Differentiated Services Code Point (DSCP) tag which specifies which class of service packets belong to. Specify a value between 0 to 63. For more information, see “DSCP (DiffServ Code Point)” on page 189.

Protocol
Specify a protocol to filter by the protocol field in the IP header. See “Protocol” on page 189 for more information.

Source IP Address
Specify the source IP address of the classifier.

Source IP Mask Length
Specify the mask length of the source IP address.

Destination IP Address
Specify the destination IP address of the classifier.

Destination IP Mask Length
Specify the destination IP mask length of the classifier.

Source Layer 4 Port
Specify a source port number in TCP or UDP to filter. See “Source Ports (Layer 4) Destination Ports (Layer 4)” on page 190.

Destination Layer 4 Port
Specify a filter destination port number in TCP or UDP to filter. “Source Ports (Layer 4) Destination Ports (Layer 4)” on page 190.

4. Click **Apply** to activate your changes.

5. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 19

Access Control Policies

This chapter describes access control policies (ACP) and how they can improve network security and performance. This chapter contains the following sections:

Note
Classifiers are a component of Access Control Policies. For information about classifiers, see Chapter 18, “Classifiers” on page 185.

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

An access control policy is a filter that controls the ingress traffic on a port. It defines a category of traffic and the action of the port when it receives packets of the category. The action is either to accept the defined packets or discard them. You can use this feature to increase network security by restricting access to certain areas or subnets or to enhance switch performance by forming network links dedicated to carrying specified types of traffic.

The heart of an ACP is a classifier which, as explained in “Overview” on page 186, defines packets that share a common trait. You can define a classifiers broadly, such as all IP packets, or specifically, such as packets from a specified end node destined for another specified node. You specify the traffic using criteria, such as source and destination MAC addresses or protocol.

When you create an ACP, you must specify the classifier that defines the traffic flow to permit or deny on a port.

There are two kinds of ACPs based on the two actions that an ACP can perform. One is called a permit ACP. Packets that meet the criteria in a permit ACP are accepted by a port. These packets can be modified by the policy in the DSCP (IP header) or CoS (Ethernet priority tag).

The second type of ACP is a deny ACP. This type of ACP denies entry to packets that meet the criteria of its classifiers.

Here is an overview of how the process works.

1. When an ingress packet arrives on a port, it is checked against the criteria in the classifiers of all the ACPs, both permit and deny, assigned to the port.

2. If the numeric sequence of the ACP determines its priority. If a deny ACP has a higher priority then a permit ACP, then the packet is discarded.

3. If the numeric sequence of the permit ACP is less than the deny ACP, then the packet is forwarded.

4. Finally, if a packet does not meet the criteria of any ACPs on a port, it is accepted by the port.

ACP Components

In order to create a policy, you must define the ACP components. A policy must contain a classifier which defines the traffic flow and a port list. For information about defining a classifier, see Chapter 18, "Classifiers" on page 185. To define a port list, see "Creating an Access Control Port List" on page 205.

In addition, you must define either a in-profile action or an out-profile action for each policy. Both may have a profile action associated with them that adjusts either the DSCP or CoS of a packet. Within the in-profile action and the out-profile action you define if the policy is a permit ACP or a deny ACP.

To manipulate the DSCP or CoS values in a packet, you may define a profile action. For a procedure to accomplish this optional task, see "Creating Profile Action" on page 199.

Guidelines

Here are guidelines for creating ACPs:

- A port can have multiple permit and deny ACPs.
- An ACP must have at least one classifier.
- You can assign an ACP to more than one port.
- An ACP filters ingress traffic, but not egress traffic.
- The action of a ACP can be either permit or deny. If a deny ACP has a lower sequence number than a permit ACP, then the deny ACP overrides the permit ACP.
- The order in which the ACPs are added to a port is not important since the packets are compared against all of a port's ACPs.
- Since classifiers cannot be assigned more than once to a port, ACPs that have the same classifier cannot be assigned to the same port.
- The switch can store up to 64 ACPs.

Creating Profile Action

The procedure in this section allows you to create a profile action for DSCP and CoS. This is an optional task. For more information about these parameters, see "ACP Components" on page 197.

To create a profile action, perform the following procedure:

1. From the menu on the left side of the page, select the **Access Control Configuration** folder.

   The **Access Control Configuration** folder expands.

2. From the **Access Control Configuration** folder, select **Profile Action**.

   The Create Profile Action Page is displayed. See Figure 59.

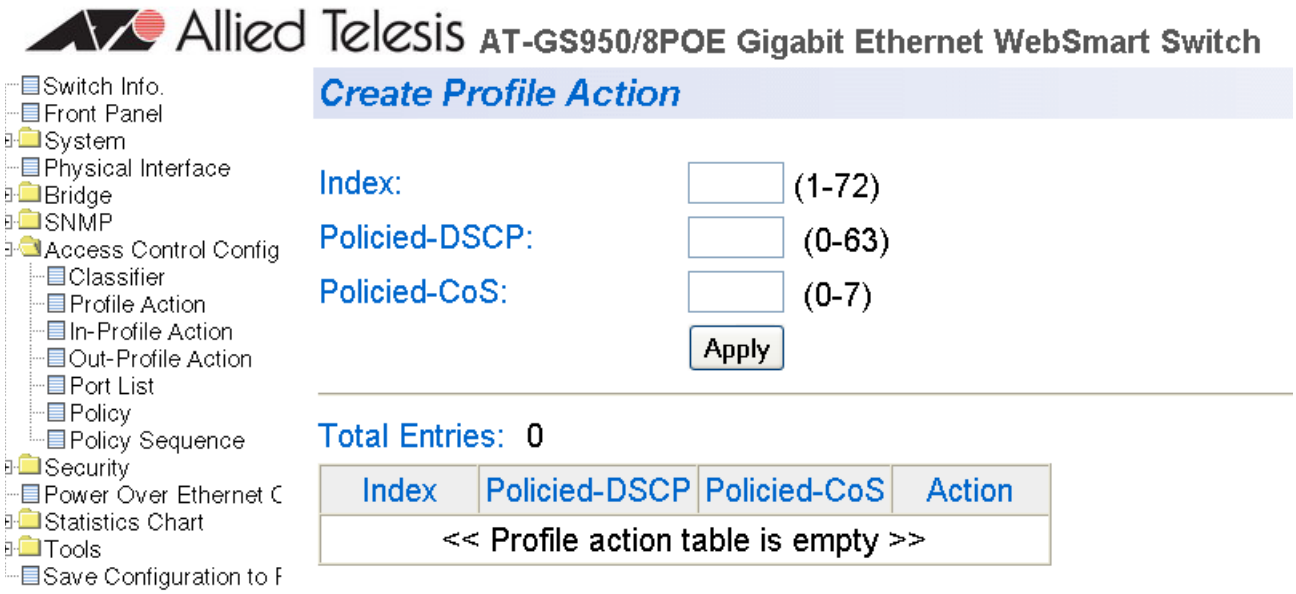

Figure 59. Create Profile Action Page

3. To set a unique numeric identifier for the profile action, select a value between 1 and 72 in the **Index** field.

4. To set the policied DSCP value for this profile, select a value between 0 and 63 in the **Policied-DSCP** field.

   This value indicates the class of service to which packets belong. For more information about DSCP, see "DSCP (DiffServ Code Point)" on page 189.

5. To set the policied CoS value for this profile, select a value between 0 and 7 in the **Policied-CoS** field.

6. Click **Apply** to activate your changes.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Creating an In-profile Action

The In-profile Action Page allows you to assign an action to a policy. For more information about these parameters, see "ACP Components" on page 197.

To create an In-profile action, perform the following procedure:

1. From the menu on the left side of the page, select the **Access Control Configuration** folder.

   The **Access Control Configuration** folder expands.

2. From the **Access Control Configuration** folder, select **In-Profile Action**.

   The Create In-Profile Action Page is displayed. See Figure 60.

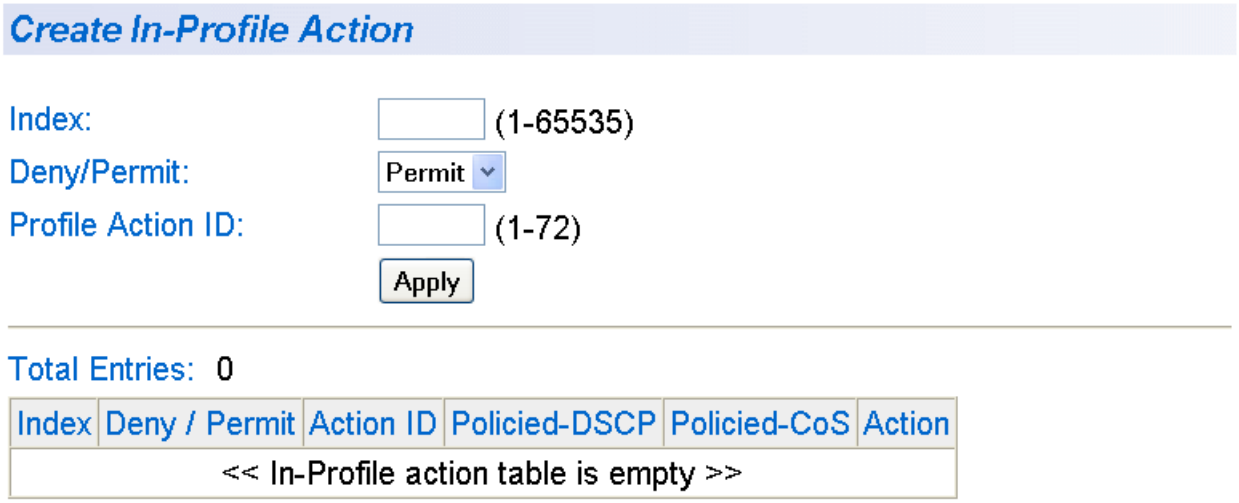

Figure 60. Create In-Profile Action Page

3. Enter a classifier index between 1 and 65535 in the **Index** field.

   For more information about the classifier index, see "Creating Classifiers" on page 192.

4. To permit or deny the profile, use the pull-down menu next to the **Deny/Permit** field.

   Choose either Permit or Deny. The default value is Permit.

5. To set an profile action ID, enter a value between 1 and 72 in the **Profile Action ID** field that was configured on the Create Profile Page.

   See "Creating Profile Action" on page 199.

6. Then click **Apply** to activate your changes.

   The In-profile Action table is updated.

7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Creating an Out-Profile Action

The Out-profile Action Page allows you to assign a classifier index to a profile action ID as well as a committed rate and a burst size. For information about how to create a profile action ID, see "Creating Profile Action" on page 199.

To create an Out-profile action, perform the following procedure:

1. From the menu on the left side of the page, select the **Access Control Configuration** folder.

   The **Access Control Configuration** folder expands.

2. From the **Access Control Configuration** folder, select **Out-Profile Action**.

   The Create Out-Profile Action Page is displayed. See Figure 61.

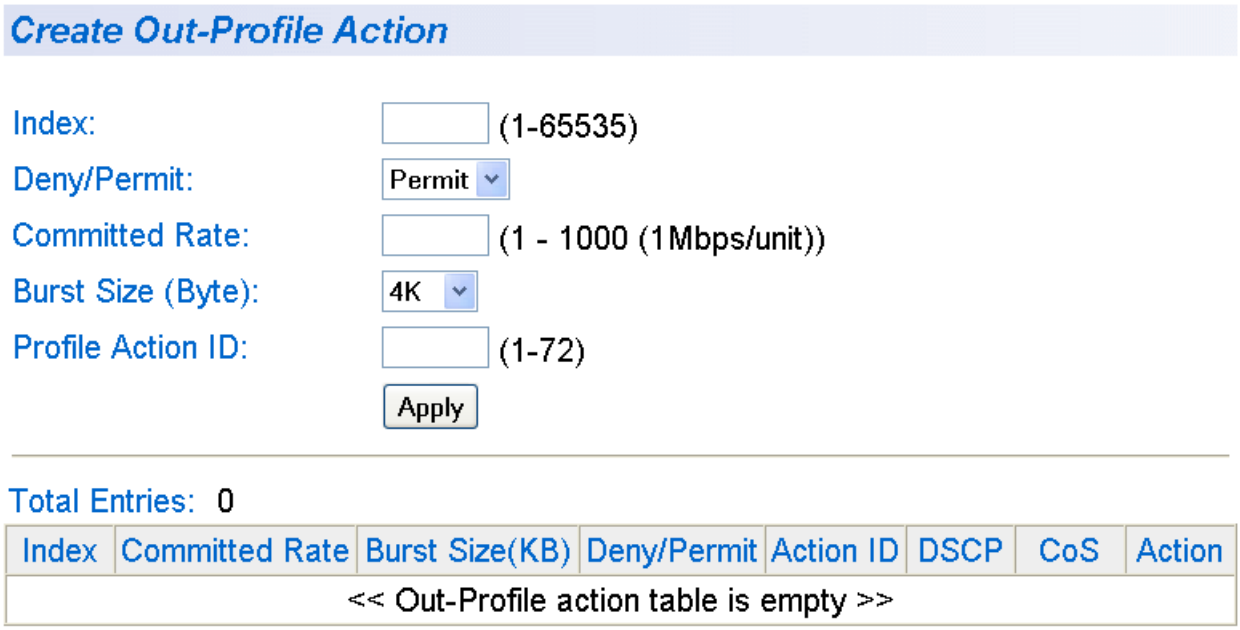

Figure 61. Create Out-Profile Action Page

3. Enter a classifier index between 1 and 65535 in the **Index** field.

   For more information about the classifier index, see "Creating Classifiers" on page 192.

4. To permit or deny the profile, use the pull-down menu next to the **Deny/Permit** field.

   Choose either Permit or Deny. The default value is Permit.

5. Enter a committed rate between 1 and 1,000 Megabits per second.
6. Enter a burst size which is the maximum amount of data, in bytes, that the switch agrees to transfer during a specified time interval by enter a value in the **Burst Size** field. Choose from the following values:
   - 4K
   - 8K
   - 16K
   - 32K
   - 64K
7. To set an profile action ID, enter a value between 1 and 72 in the **Profile Action ID** field that was configured on the Create Profile Page.

   See “Creating Profile Action” on page 199.
8. Then click **Apply** to activate your changes.

   The Out-profile Action table is updated.
9. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Creating an Access Control Port List

This section provides a procedure to assign an index to a list of ports.

To create a port list, perform the following procedure:

1. From the menu on the left side of the home page, select the **Access Control Configuration** folder.

   The **Access Control Configuration** folder expands.

2. From the **Access Control Configuration** folder, select **Port List**.

   The Port List Page is displayed. See Figure 62.

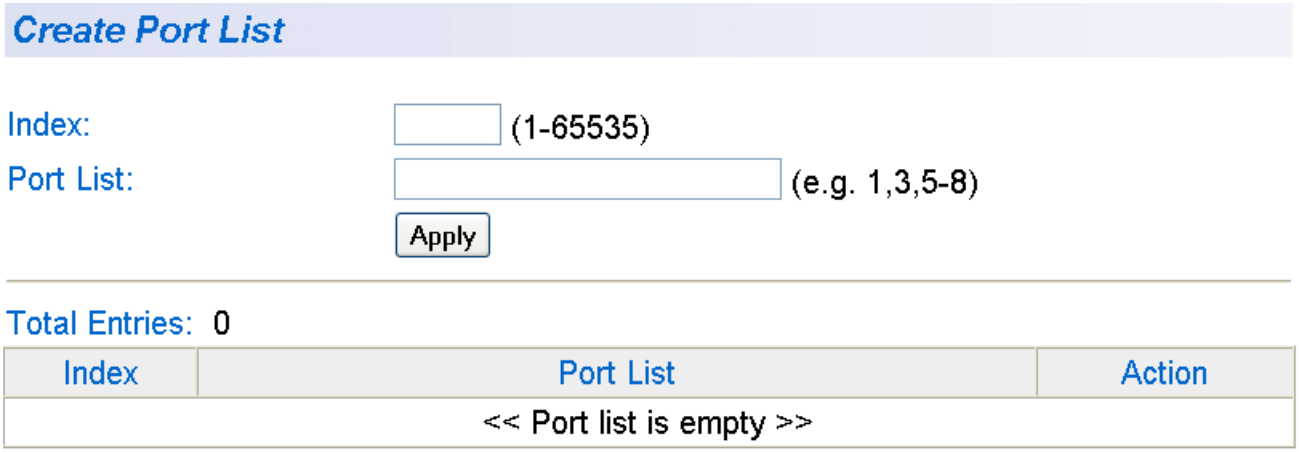

Figure 62. Create Port List Page

3. Enter a classifier index in the **Index** field.
4. Enter list of ports in the **Port List** field.

   Enter ports separated by commas or enter a range of ports separated by a hyphen.

5. Then click **Apply** to activate your changes.

   The Port List table is updated.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Creating a Policy

This section provides a procedure to create a policy. For information about classifiers, see “Creating Classifiers” on page 192. For procedures to create profile actions, see “Creating Profile Action” on page 199, “Creating an In-profile Action” on page 201, and “Creating an Out-Profile Action” on page 203. To create a port list, see “Creating an Access Control Port List” on page 205.

To create a policy, perform the following procedure:

1. From the menu on the left side of the home page, select the **Access Control Configuration** folder.

   The **Access Control Configuration** folder expands.

2. From the **Access Control Configuration** folder, select **Policy**.

   The Policy Page is displayed. See Figure 62.

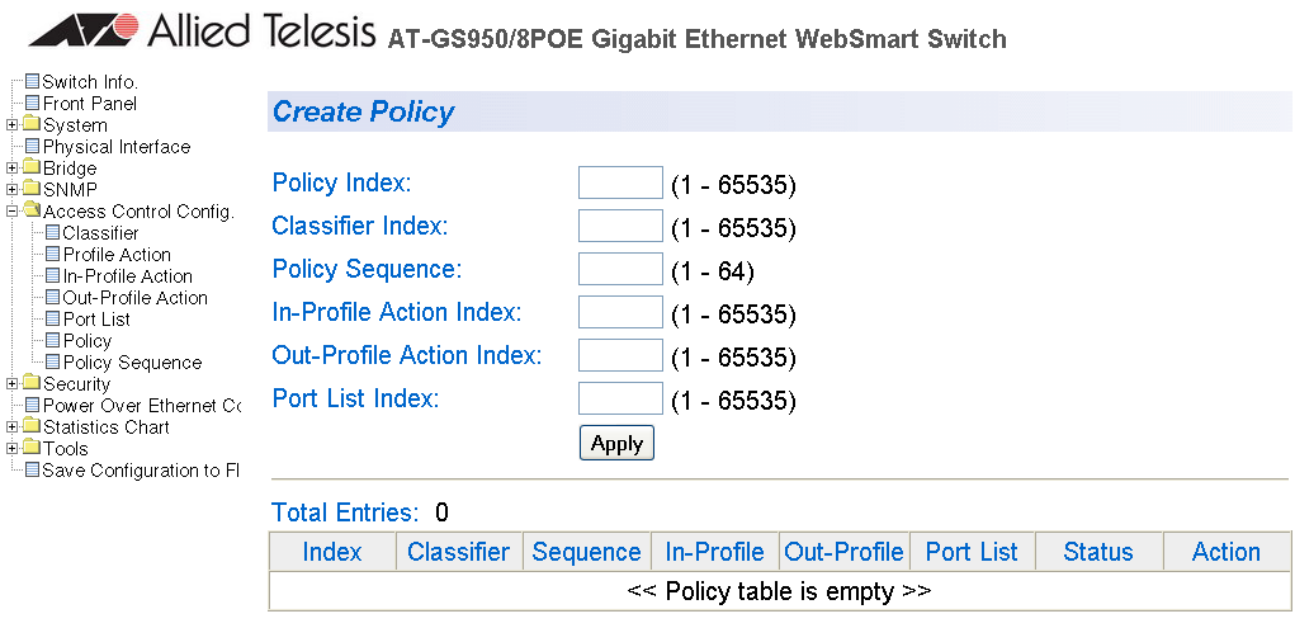

Figure 63. Policy Page

3. Input the following fields:

   Policy Index
   Indicates the index value of the policy. Specify a value between 1 and 65,536.

Classifier Index
Specific the identification number of a classifier. You must assign a classifier that you created on the Classifier Page.

Policy Sequence
Specify the priority of the policy where the lowest number of the policy sequence has the highest priority. Enter a value between 1 (highest priority) and 64 (lowest priority).

Note
If a deny ACP has a lower sequence number than a permit ACP, then the deny ACP overrides the permit ACP.

In-Profile Action Index
Specify an In-profile Action index that you created with on the Create In-Profile Action Page.

Out-Profile Action Index
Specify an Out-profile Action index that you created with on the Create Out-Profile Action Page.

Port List Index
Specify a port list that you created with the Create Port List Page.

4. Then click **Apply** to activate your changes.

5. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Displaying a Policy Sequence

To create a policy sequence, perform the following procedure:

1. From the menu on the left side of the home page, select the **Access Control Configuration** folder.

   The **Access Control Configuration** folder expands.

2. From the **Access Control Configuration** folder, select **Policy Sequence**.

   The Policy Sequence Page is displayed. See Figure 64.

Figure 64. Policy Sequence Page

3. Select a port with the pull-down menu next to the **Select Port** field.

4. Select **Display by Index order** to display the policy by the index number or select **Display by Sequence order** to display the policy by policy sequence.

   See “Creating a Policy” on page 206 for information about a policy index and a policy sequence.

5. Then click **Apply** to activate your changes.

6. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Chapter 20

Management Software Updates

This chapter explains the methods for upgrading the AT-S101 Management Software on the switch and saving configuration files. This chapter contains the following sections:

Note
For information on how to obtain new releases of the AT-S101 Management Software, refer to “Management Software Updates” on page 12.

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

You can use the Management Software Updates features to upgrade the AT-S101 Management Software to a new version, upload a configuration file from the AT-GS950/8POE switch onto an a PC, or download an image or configuration file from the switch onto an a PC.

There are two methods to upgrade the AT-S101 Management Software:

- using a web browser via HTTP
- using a TFTP server

To upgrade a firmware image using HTTP, you only need to have access to an Internet browser. However, to upgrade a firmware image using TFTP, you must have access to an TFTP server.

In addition, you can save a configuration file from one switch and load it onto another switch or onto all of your AT-GS950/8POE switches. This ensures identical configurations on all of your switches. In addition, loading an existing configuration saves time.

Upgrading a Firmware Image Using HTTP

This section describes how to upgrade an firmware image of the AT-S101 software using HTTP on an Internet server. Before downloading a new version of the AT-S101 Management Software onto the switch with HTTP, note the following:

- The current configuration of a switch is retained when a new AT-S101 software image is installed. To return a switch to its default configuration values, see "Returning the AT-S101 Management Software to the Factory Default Values" on page 44.
- On the switch that you are downloading the new image file to, assign an IP address and subnet mask. For instructions on how to set the IP address and subnetmask on a switch, see "Configuring an IP Address, Subnet Mask and Gateway Address" on page 22. To disable a DHCP client, see "Enabling and Disabling the DHCP Client" on page 27.

Caution
Downloading a new version of management software onto the switch causes the device to reset. Some network traffic may be lost during the reset process.

This procedure assumes that you have already obtained the software and have stored it on the computer from which you will be performing this procedure.

To download the AT-S101 image software onto the switch using HTTP, perform the following procedure:

1. From the menu on the left side of the home page, select the **Tools** folder.

   This folder expands to show the contents of the **Firmware Upgrade** folder.

2. From the **Firmware Upgrade** folder, select **via HTTP**.

The Firmware Upgrade via HTTP Page is displayed. See Figure 65.

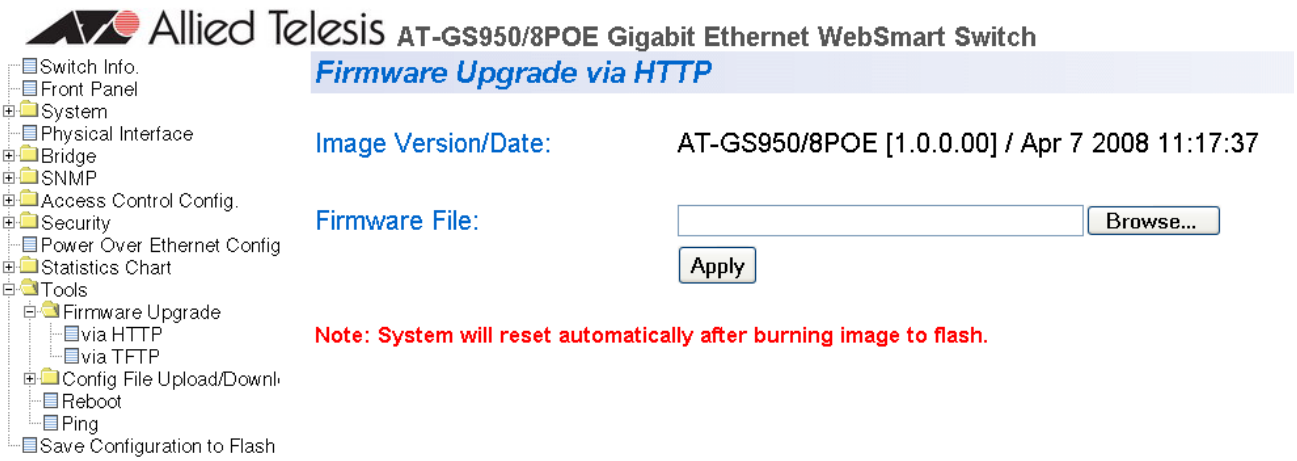

Figure 65. Firmware Upgrade via HTTP Page

3. Change the following parameter as necessary:

   Firmware File:
   Enter the path of the firmware file or click the **Browse** button and select the filename.

4. Click **Apply**.

   The software begins to download onto the switch immediately. This process takes a few minutes. After the software download is complete, the switch initializes the software and reboots. You will lose your web browser connection to the switch during the reboot process.

Upgrading a Firmware Image Using TFTP

This section describes how to upgrade an firmware image of the AT-S101 software using TFTP on an TFTP server. Before downloading a new version of the AT-S101 software onto the switch, note the following:

- The current configuration of a switch is retained when a new AT-S101 software image is installed. To return a switch to its default configuration values, see "Returning the AT-S101 Management Software to the Factory Default Values" on page 44.
- Your network must have a node with TFTP server software.
- You must store the new AT-S101 image file on the TFTP server.
- Start the TFTP server software *before* you begin the download procedure.

Caution
Downloading a new version of management software onto the switch causes the device to reset. Some network traffic may be lost during the reset process.

This procedure assumes that you have already obtained the software and have stored it on the computer from which you will be performing this procedure.

To download the AT-S101 image software onto the switch using a TFTP server, perform the following procedure:

1. From the menu on the left side of the home page, select the **Tools** folder.

   This folder expands to show the contents of the **Firmware Upgrade** folder.

2. From the **Firmware Upgrade** folder, select **via TFTP**.

The Firmware Upgrade via TFTP page is shown in Figure 66.

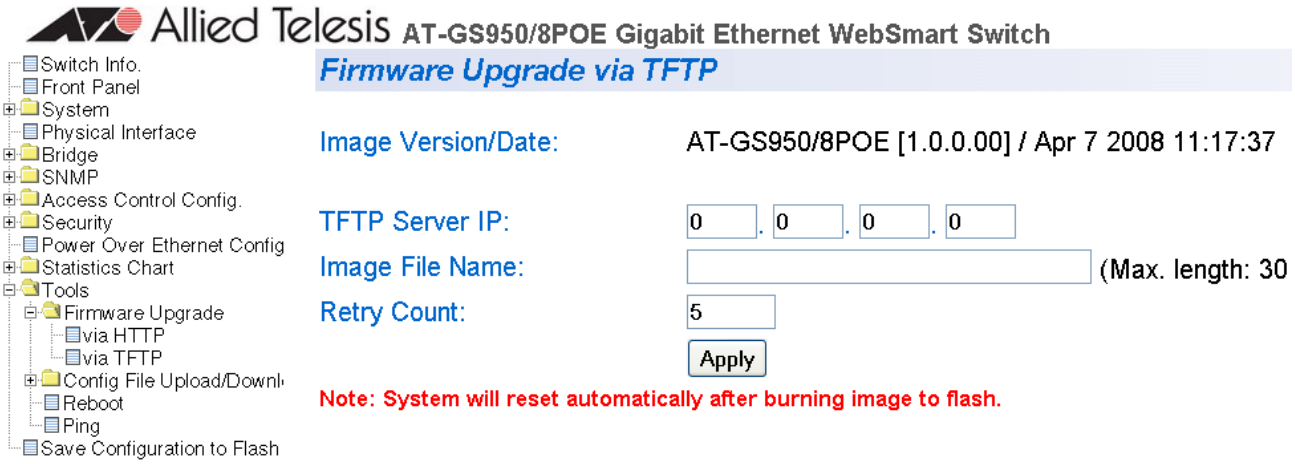

Figure 66. Firmware Upgrade via TFTP Page

The Image/Version Date shows the current version and date of software installed on the switch.

3. Change the following parameters as necessary:

   TFTP Server IP
   The IP address of the TFTP server from which you are downloading the new software.

   Image File Name
   The name of the AT-S101 file you are downloading.

   Retry Count:
   The number of times the firmware upgrade is retried. The default number of tries is 5. The range is 1 through 20.

4. Click **Apply**.

   The software immediately begins to download onto the switch. This process takes a few minutes. After the software download is complete, the switch initializes the software and reboots. You will lose your web browser connection to the switch during the reboot process.

Downloading or Uploading a Configuration File via HTTP

This section describes how to download or upload a configuration file using HTTP on an Internet server. Before you upload or download a configuration file via HTTP, note the following:

- You must be able to access the new AT-S101 image file from your PC.
- On the switch that you are downloading the new image file to, assign an IP address and subnet mask. For instructions on how to set the IP address on a switch, refer to "Configuring an IP Address, Subnet Mask and Gateway Address" on page 22. To disable a DHCP Client, see "Enabling and Disabling the DHCP Client" on page 27.

To download or upload an AT-S101 configuration file onto the switch using a web browser, perform the following procedure:

1. From the menu on the left side of the home page, select the **Tools** folder.

   This **Tools** folder expands.

2. From the **Tools** folder, select **Config File Upload/Down** folder.

   This **Config File Upload/Down** folder expands.

3. From the **Config File Upload/Down** folder, select **via HTTP**.

   The Configuration Upload/Download via HTTP Page is displayed. See Figure 67.

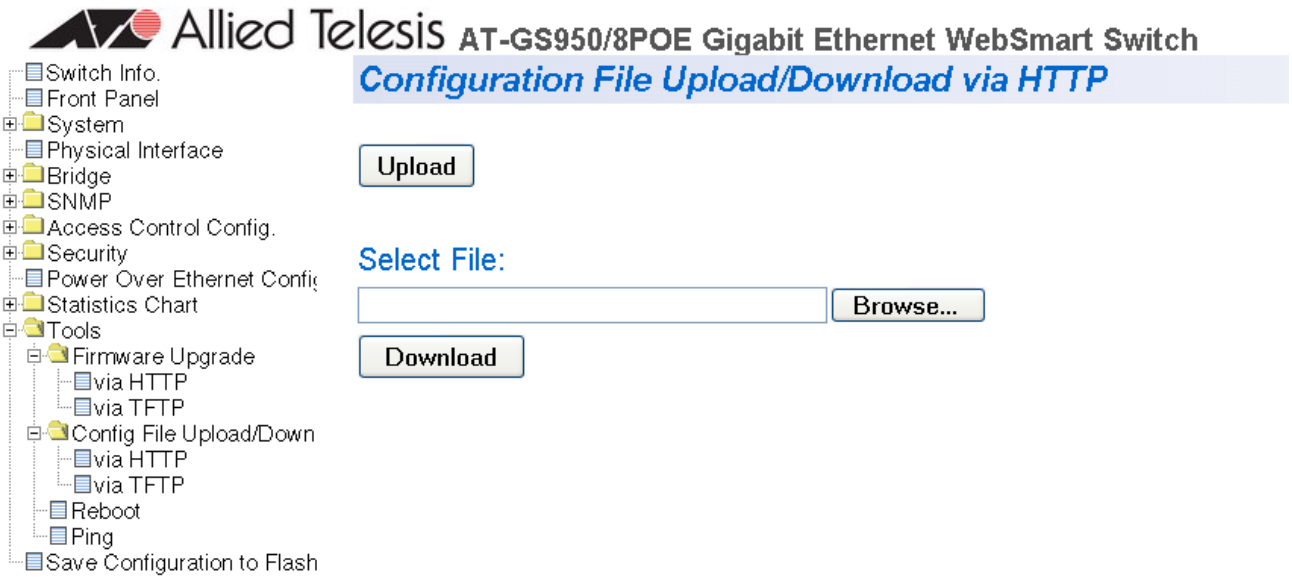

Figure 67. Configuration Upload/Download via HTTP Page

4. Change the following parameters as necessary:

   Select File:
   Enter the path of the firmware file or click the **Browse** button and select the filename.

5. Select one of the following:

   Upload
   Select this button to upload a configuration file onto the switch.

   The software immediately begins to upload from or download onto the switch.

   Download
   Select this button to download a configuration file from the switch onto your PC.

   The following window is displayed. See Figure 68.

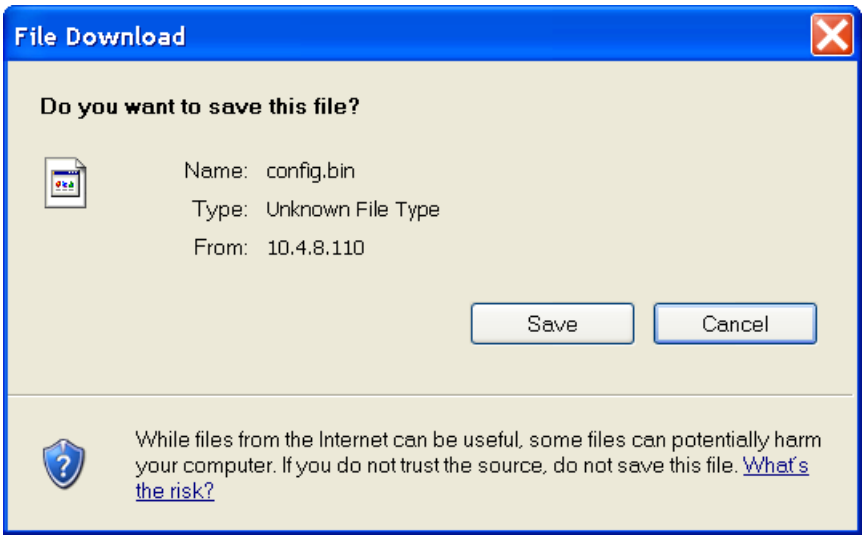

Figure 68. File Download with HTTP

6. Click Save to download the configuration file onto the switch.

   The Microsoft “Save As” window is displayed.

7. Save the file in the appropriate directory.

Downloading or Uploading a Configuration File via TFTP

This section describes how to download or upload a configuration file using TFTP on an TFTP server. Before uploading or downloading a configuration file onto the switch using TFTP, note the following:

- The current configuration of a switch is retained when a new AT-S101 software image is installed. To return a switch to its default configuration values, see "Returning the AT-S101 Management Software to the Factory Default Values" on page 44.
- Your network must have a node with TFTP server software.
- You must store the new AT-S101 image file on the TFTP server.
- Start the TFTP server software *before* you begin the download procedure.

To download or upload an AT-S101 configuration file onto the switch using a TFTP server, perform the following procedure:

1. From the menu on the left side of the home page, select the **Tools** folder.

   The **Tools** folder expands.

2. From the **Tools** folder, select the **Config File Upload/Download** folder.

   The **Config File Upload/Download** folder expands.

3. From the **Config File Upload/Down** folder, select **via TFTP**.

   The Configuration Upload/Download via TFTP Page is displayed. See Figure 69.

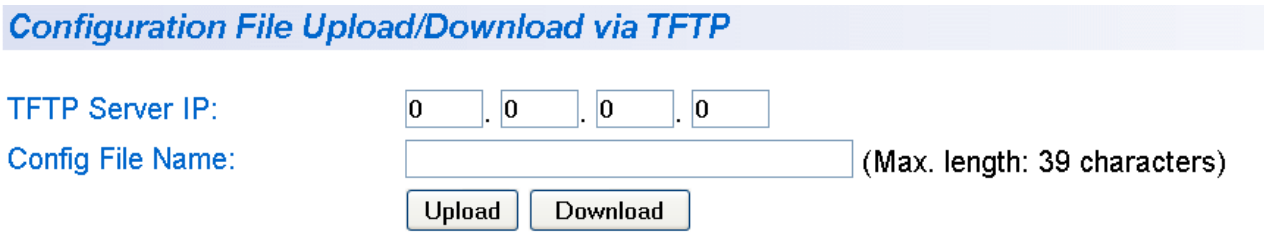

Figure 69. Configuration Upload/Download via TFTP Page

4. Change the following parameters as necessary:

   Enter the IP address of the TFTP server in the field next to the **TFTP Server IP** parameter.

   Enter the name of the configuration file in the field next to the **Config File Name** parameter.

5. Select one of the following:

   Upload
   Select this button to upload a configuration file onto the switch.

   Download
   Select this button to download a configuration file from the switch onto your PC.

   The software immediately begins to upload from or download onto the switch. This process takes a few minutes.

   If you are downloading software, the switch initializes the software and reboots after the software download is complete. You will lose your web browser connection to the switch during the reboot process.

Chapter 21

Statistics

The sections in this chapter explain how to display traffic, error, and history statistics about the AT-GS950/8POE switch and its ports. This chapter includes the following sections:

Note
To save your changes, select **Save Configuration to Flash** from the menu on the left side of the page.

Overview

Statistics provide important information for troubleshooting switch problems at the port level. The AT-S101 Management Software provides a versatile set of statistics charts that you can customize for your needs, including (depending upon the chart) the ports whose statistics you want to view and the color used to draw the chart.

There are three types of statistics charts:

- Traffic Comparison. This chart allows you to display a specified traffic statistic over all of the ports. You can select 24 statistics types and 12 colors for each port. The Traffic Comparison statistics chart is described in “Displaying Traffic Comparison Statistics” on page 221.
- Error Group. The Error Group chart displays the discard and error counts for a specified port and is described in “Displaying Error Group Statistics” on page 225.
- Historical Status. This chart allows you to select from 12 statistics to view for a selection of ports for however long this chart is running on the management workstation. The Historical Status chart is described in “Displaying Historical Status Charts” on page 227.

Displaying Traffic Comparison Statistics

To display traffic comparison statistics, perform the following procedure:

1. Select the **Statistics Chart** folder.

   The **Statistics Chart** folder expands.

2. From the **Statistics Chart** folder, select **Traffic Comparison**.

   The Traffic Comparison Page opens as shown in Figure 70.

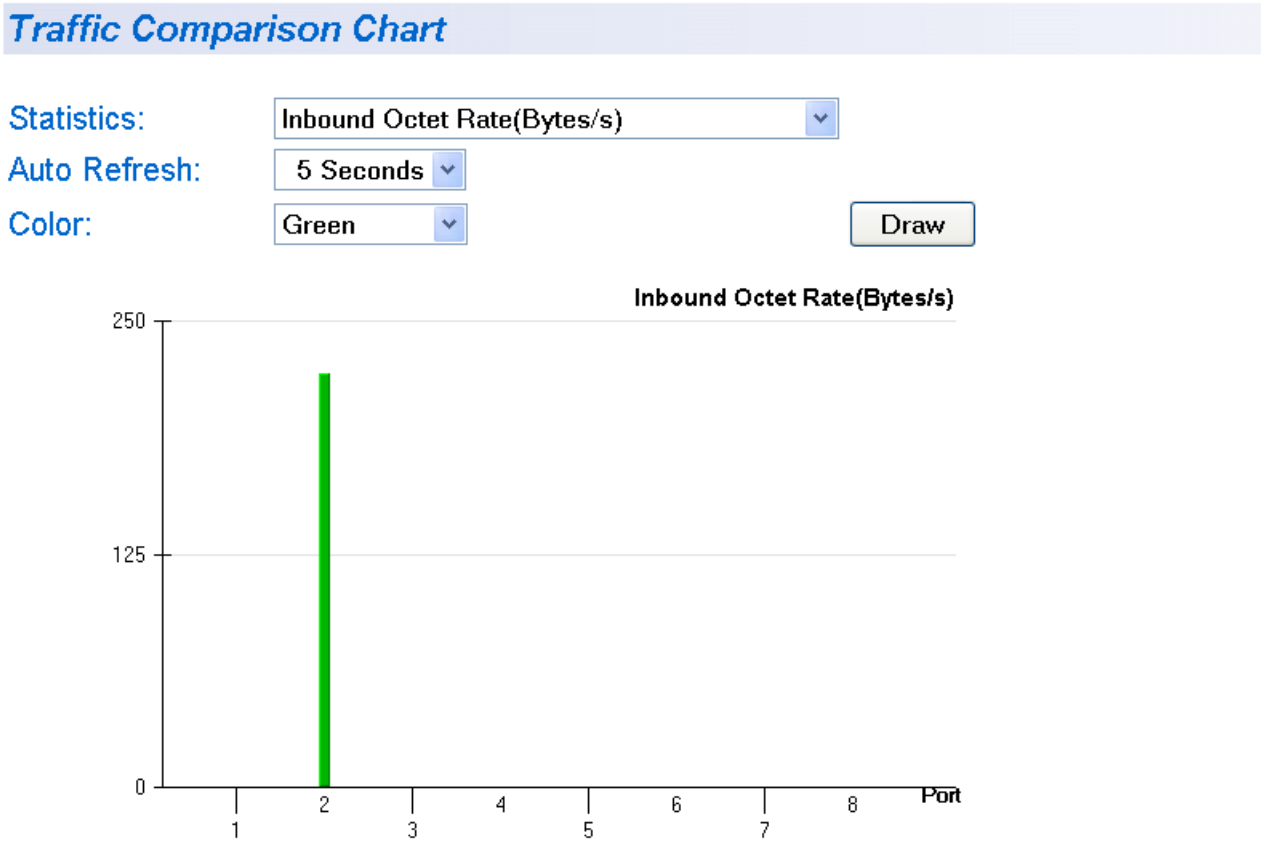

Figure 70. Traffic Comparison Page

3. To view traffic statistics, click on the arrow next to “Statistics” and select one of the options in Table 6.

Table 6 Traffic Comparison Options

Option	Definition
Inbound Octet Rate (Bytes/s)	Measures the rate of inbound octet bits in bytes per second.
Inbound Unicast Packet Rate (Pkts/s)	Measures the rate of inbound unicast packets in packets per second.
Inbound Non-unicast Packet Rate (Pkts/s)	Measures the rate of inbound non-unicast packets (such as broadcast and multicast packets) in packets per second.
Inbound Discard Rate (Pkts/s)	Measures the rate of inbound packets that are discarded. This is measured in packets per second.
Inbound Error Rate (Pkts/s)	Measures the number of inbound errors in packets per second.
Outbound Octet Rate (Bytes/s)	Measures the number of outbound octet bits in bytes per second.
Outbound Unicast Packet Rate (Pkts/s)	Measures the number of outbound unicast packets in packets per second.
Outbound Non-unicast Packet Rate (Pkts/s)	Measures the number of outbound non-unicast packets (such as broadcast and multicast packets) in packets per second.
Outbound Discard Rate (Pkts/s)	Measures the rate of outbound discarded packets in packets per second.
Outbound Error Rate (Pkts/s)	Measures the rate of outbound errors in packets per second.
Ethernet Undersize Packet Rate (Pkts/s)	Measures the rate of undersized Ethernet packets in packets per second.
Ethernet Oversize Packet Rate (Pkts/s)	Measures the rate of oversized Ethernet packets in packets per second.
Inbound Octets (Bytes/s)	Measures the number of inbound octet bits in bytes per second.
Inbound Unicast Packets (Pkts)	Measures the number of inbound unicast packets in packets per second.

Table 6 Traffic Comparison Options (Continued)

Option	Definition
Inbound Non-unicast Packets (Pkts)	Measures the number of inbound non-unicast packets (such as broadcast and multicast packets) in packets per second.
Inbound Discards (Pkts)	Measures the number of inbound discarded packets in packets per second.
Inbound Errors (Pkts/s)	Measures the number of inbound errors in packets per second.
Outbound Octets (Bytes/s)	Measures the rate of outbound octet bits in bytes per second.
Outbound Unicast Packets (Pkts)	Measures the number of outbound unicast packets in packets per second.
Outbound Non-unicast Packets (Pkts)	Measures the number of outbound non-unicast (such as broadcast and multicast packets) packets.
Outbound Discards (Pkts)	Measures the number of outbound discarded packets.
Outbound Errors (Pkts)	Measures the number of outbound error packets.
Ethernet Undersize Packets (Pkts)	Measures the number of undersized Ethernet packets.
Ethernet Oversize Packets (Pkts)	Measures the number of oversized Ethernet packets.

4. To select the amount of time before the screen is refreshed, click **Auto Refresh**. Choose from the following options:
   - 5 seconds
   - 10 seconds
   - 15 seconds
   - 30 seconds
5. To select the color of the traffic comparison graph, select **Color**. Choose one of the following colors:
   - Green (This is the default.)
   - Blue
   - Red
   - Purple

- Yellow
- Orange
- Gray
- Light Red
- Light Blue
- Light Green
- Light Yellow
- Light Gray

6. To create the traffic comparison graph, select **Draw**.
7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Displaying Error Group Statistics

To display error group statistics for a port, perform the following procedure:

1. Select the **Statistics Chart** folder.

   The **Statistics Chart** folder expands.

2. From the **Statistics Chart** folder, select **Error Group**.

   The Error Group Chart Page is displayed in Table 71.

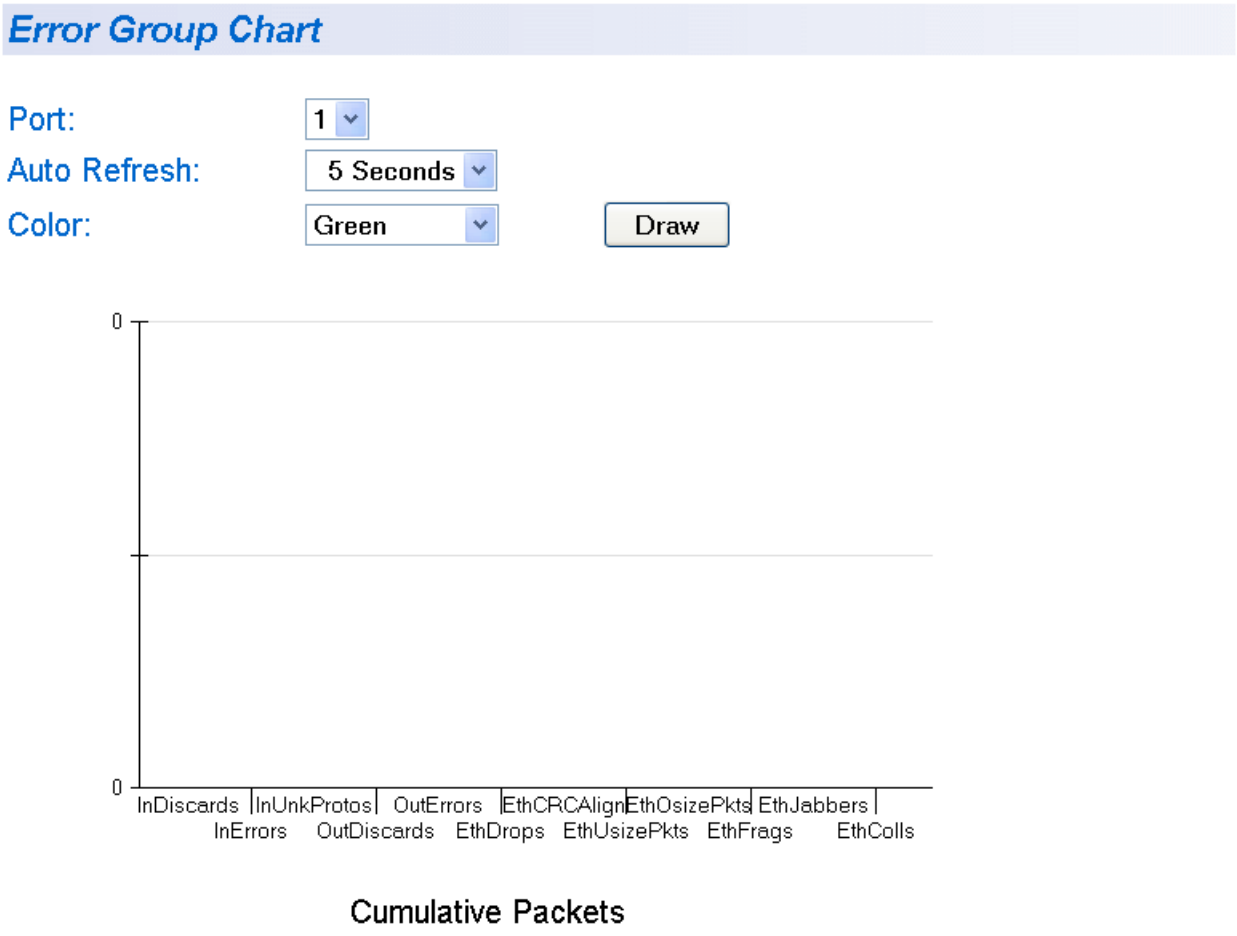

Figure 71. Error Group Chart Page

3. Select a port number from the pull down menu next to Port.

4. To select the amount of time before the screen is refreshed, click **Auto Refresh**. Choose from the following options:
   - 5 seconds
   - 10 seconds
   - 15 seconds
   - 30 seconds
5. To select the color of the traffic comparison graph, select **Color**. Choose one of the following colors:
   - Green (This is the default.)
   - Blue
   - Red
   - Purple
   - Yellow
   - Orange
   - Gray
   - Light Red
   - Light Blue
   - Light Green
   - Light Yellow
   - Light Gray
6. To create the Error Group Chart, select **Draw**.
7. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Displaying Historical Status Charts

To display historical status charts statistics for a port, perform the following procedure:

1. Select the **Statistics Chart** folder.

   The **Statistics Chart** folder expands.

2. From the **Statistics Chart** folder, select **Historical Status**.

   The Historical Status Chart Page is displayed in Table 72.

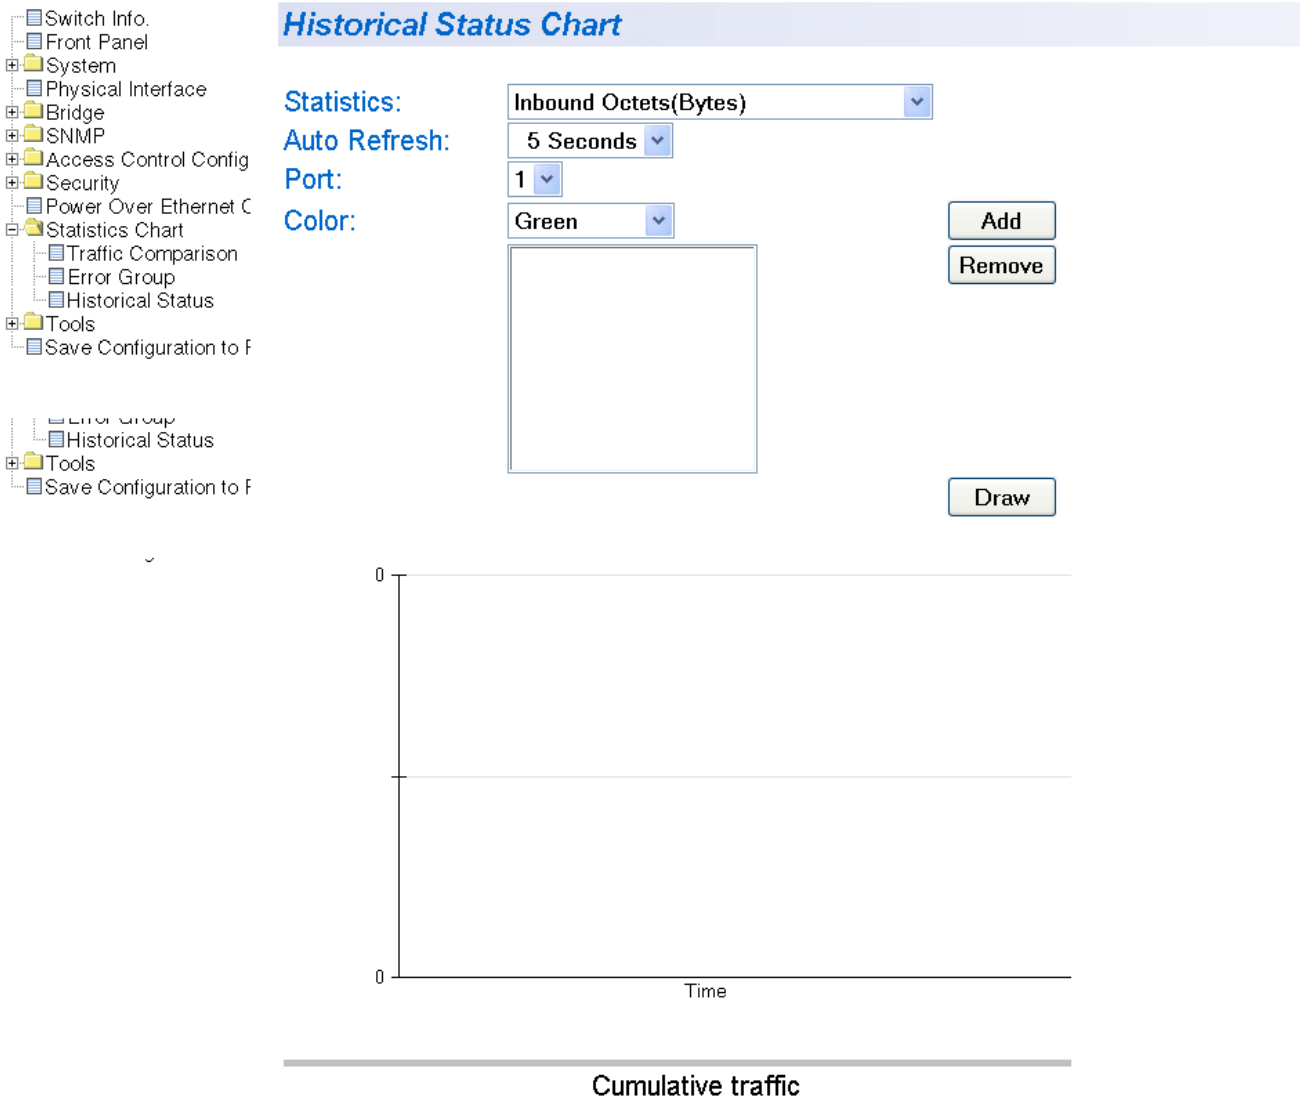

Figure 72. Historical Status Chart Page

3. To view historical statistics, click on the arrow next to “Statistics” and select one of the options in Table 7.

Table 7 Historical Status Options

Option	Definition
Inbound Octet Rate (Bytes)	Measures the rate of inbound octet bits in bytes per second.
Inbound Unicast Packet Rate (Pkts)	Measures the rate of inbound unicast packets in packets per second.
Inbound Non-unicast Packet Rate (Pkts)	Measures the rate of inbound non-unicast packets (such as broadcast and multicast packets) in packets per second.
Inbound Discards (Pkts)	Measures the number of inbound discarded packets in packets per second.
Inbound Errors (Pkts)	Measures the number of inbound errors in packets per second.
Outbound Octets (Bytes)	Measures the number of outbound octet bits in bytes per second.
Outbound Unicast Packets (Pkts)	Measures the number of outbound unicast packets in packets per second.
Outbound Non-unicast Packets (Pkts)	Measures the number of outbound non-unicast (such as broadcast and multicast packets) packets.
Outbound Discards (Pkts)	Measures the number of outbound discarded packets.
Outbound Errors (Pkts)	Measures the number of outbound error packets.
Ethernet Undersize Packets (Pkts)	Measures the number of undersized Ethernet packets.
Ethernet Oversize Packet Rate (Pkts)	Measures the number of oversized Ethernet packets.

4. To select the amount of time before the screen is refreshed, click **Auto Refresh**. Choose from the following options:

   – 5 seconds

   – 10 seconds

- 15 seconds
- 30 seconds

5. To select the color of the traffic comparison graph, select **Color**. Choose one of the following colors:
   - Green (This is the default.)
   - Blue
   - Red
   - Purple
   - Yellow
   - Orange
   - Gray
   - Light Red
   - Light Blue
   - Light Green
   - Light Yellow
   - Light Gray
6. To create the history group chart, select **Add**. Then click **Draw.**
7. To draw the historical group chart, select **Draw**.
8. From the menu on the left side of the page, select **Save Configuration to Flash** to save your changes.

Appendix A

AT-S101 Management Software Default Settings

Table 8 lists the factory default settings for the management software.

Table 8. AT-S101 Default Settings

Parameter	Default Setting
IP Configuration	
IP Address	0.0.0.0
Subnet Mask	0.0.0.0
Default Gateway Address	0.0.0.0
DHCP Client	Disabled
System Administration	
System Name	(blank)
System Location	(blank)
System Contact	(blank)
Manager Interface	
Manager Username	manager
Manager Password	friend
Console Idle Timeout	5 minutes
Web Server	Enabled
Ping Configuration	
Target IP Address	0.0.0.0
Number of Requests	10
Timeout	3 seconds
Port Configuration	
Link Status	Down
Speed	Auto-Negotiation
Admin. Status	Enabled

Table 8. AT-S101 Default Settings (Continued)

Parameter	Default Setting
(Duplex) Mode	Auto-Negotiation
Jumbo (frames)	Disabled
Flow Control (Full-duplex Mode)	Enabled
Back pressure (Half-duplex Mode)	Enabled (not adjustable)
Port Trunking	
Status	Disabled
LACP Port Trunking	
System Priority	32768
IGMP Snooping	
Status	Disabled
IGMP Snooping Age-Out Timer	280 seconds
Port Mirroring	
Status	Disabled
VLAN	
Name	Default VLAN
VID	1
Ports	All Ports (Untagged)
SNMP	
public community	Enabled, Read Only
private community	Enabled, Read Write
Quality of Service	
Status	Disabled
Mappings of IEEE 802.1p Priority Levels to Egress Port Priority Queues	See Table 1 on page 63.
Priority Override Status	Disabled
Priority Queue	Queue 0
STP and RSTP	
Status	Disabled
Protocol version	RSTP

Table 8. AT-S101 Default Settings (Continued)

Parameter	Default Setting
Hello Time	2 seconds
Maximum Age	20 seconds
Forward Delay	15 seconds
Bridge Hello Time	2 seconds
Bridge Maximum Age	20 seconds
Bridge Forward Delay	20 seconds
Bandwidth Control	
DLF Ingress Packet Status	Disabled
Broadcast/Multicast Packet Threshold	Low
Packet Threshold Mode	Broadcast/Multicast
IP Access List	
IP Restriction	Disabled
802.1x Port-based Network Access Control	
NAS ID	Nas1
Port Control	Force Authorized
Transmission Period	30 seconds
Supplicant Timeout	30 seconds
Server Timeout	30 seconds
Maximum Requests	2
Quiet Period	60 seconds
Re-authentication Period	3600 seconds
Re-authentication Status	Disabled
RADIUS Client	
Server IP Address	0.0.0.0
Shared Secret	(blank)
Response Time	10 seconds
Maximum Retransmissions	3

Table 8. AT-S101 Default Settings (Continued)

Parameter	Default Setting
Power over Ethernet	
Power budget	65 W
Priority	Low
Limit (mW)	15400
Access Control Policies	
Burst Size (Bytes)	4K
Upgrade Configuration	
TFTP Server IP Address	0.0.0.0
Image Filename	(blank)
Retry Count	5
Statistics	
Statistics	Inbound Octets (Bytes)
Auto Refresh	5 seconds
Port	1
Color	Green

Index