Management Software

AT-S62

Menus Interface User's Guide

AT-8500 Series Layer 2+ Fast Ethernet Switches

Version 1.4.0

613-000124 Rev. B

Allied Telesyn

Contents

Figures

Preface

This guide contains instructions on how to configure an AT-8500 Series Layer 2+ Fast Ethernet Switch using the menus interface in the AT-S62 management software.

For instructions on how to manage the switch from the web browser interface or the command line interface, refer to the *AT-S62 Web Browser Interface User's Guide* and the *AT-S62 Command Line Interface User's Guide*. These guides are available from the Allied Telesyn web site.

This preface contains the following sections:

Caution

The software described in this documentation contains certain cryptographic functionality and its export is restricted by U.S. law. As of this writing, it has been submitted for review as a "retail encryption item" in accordance with the Export Administration Regulations, 15 C.F.R. Part 730-772, promulgated by the U.S. Department of Commerce, and conditionally may be exported in accordance with the pertinent terms of License Exception ENC (described in 15 C.F.R. Part 740.17). In no case may it be exported to Cuba, Iran, Iraq, Libya, North Korea, Sudan, or Syria. If you wish to transfer this software outside the United States or Canada, please contact your local Allied Telesyn sales representative for current information on this product's export status.

How This Guide is Organized

This manual is divided into the following sections.

Section I: Basic Operations

The chapters in this section explain how to perform basic switch operations, such as setting port parameters, creating port trunks, and viewing the MAC address table.

Section II: Advanced Operations

The chapters in this section explain some of the more advanced operations, such as using the file system, downloading and uploading files, and configuring Quality of Service.

Section III: SNMPv3 Operations

The chapter in this section explains how to configure the switch for SNMPv3. (The instructions for SNMPv1 and SNMPv2 are in Section 1, Basic Operations.)

Section IV: Spanning Tree Protocols

The chapters in this section explain the Spanning Tree, Rapid Spanning Tree, and Multiple Spanning Tree Protocols.

Section V: Virtual LANs

The chapters in this section explain port-based and tagged VLANs, GVRP, multiple VLAN modes, and protected ports VLANs.

Section VI: Port Security

The chapters in this section explain MAC address-based port security and 802.1x port-based access control.

Section VII: Management Security

The chapters in this section explain the management security features, such as the Secure Sockets Layer (SSL) and the Secure Shell (SSH) protocols.

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 Telesyn products are available in Portable Document Format (PDF) from on our web site at **www.alliedtelesyn.com**. You can view the documents on-line or download them onto a local workstation or server.

Contacting Allied Telesyn

This section provides Allied Telesyn contact information for technical support as well as sales or corporate information.

Online Support

You can request technical support online by accessing the Allied Telesyn Knowledge Base from the following web site: **www.alliedtelesyn.com/kb**. 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 & Services section of the Allied Telesyn web site: **www.alliedtelesyn.com**.

Returning Products

Products for return or repair must first be assigned a Return Materials Authorization (RMA) number. A product sent to Allied Telesyn without a RMA number will be returned to the sender at the sender's expense.

To obtain a RMA number, contact Allied Telesyn's Technical Support at our web site: **www.alliedtelesyn.com**.

For Sales or Corporate Information

You can contact Allied Telesyn for sales or corporate information at our web site: **www.alliedtelesyn.com**. To find the contact information for your country, select Contact Us -> Worldwide Contacts.

Management Software Updates

You can download new releases of management software for our managed products from either of the following Internet sites:

- Allied Telesyn web site: **www.alliedtelesyn.com**
- Allied Telesyn FTP server: **ftp://ftp.alliedtelesyn.com**

To download new software from the Allied Telesyn FTP server using your workstation's command prompt, you need FTP client software and you must log in to the server. Enter "anonymous" as the user name and your email address for the password.

New Features History

The following subsection contains the new features in the AT-S62 management software.

Version 1.4.0

Table 1 lists the new features in version 1.4.0 of the AT-S62 management software.

Table 1. New Features in AT-S62 Version 1.4.0

Change	Chapter and Procedure
Fan Control Feature for the AT-8524POE Switch	
New feature.	Chapter 3, “Basic Switch Parameters” on page 49 New procedure: ❒ “Setting Fan Control” on page 69
Quality of Service - Flow Groups and Traffic Classes	
Added the following new parameters to QoS flow groups and traffic classes: ❒ ToS parameter for replacing the Type of Service field of IPv4 packets. ❒ Move ToS to Priority parameter for replacing the value in the 802.1p priority field with the value in the ToS priority field in IPv4 packets. ❒ Move Priority to ToS parameter for replacing the value in the ToS priority field with the 802.1p priority field in IPv4 packets.	Chapter 15, “Quality of Service” on page 267 Modified procedures: ❒ “Creating a Flow Group” on page 283 ❒ “Modifying a Flow Group” on page 285 ❒ “Creating a Traffic Class” on page 290 ❒ “Modifying a Traffic Class” on page 294
Quality of Service - Policies	
Added the following new parameters to QoS policies: ❒ ToS, Move ToS to Priority, and Move Priority to ToS, as defined above. ❒ Send to Mirror Port parameter for copying traffic to a destination mirror port.	Chapter 15, “Quality of Service” on page 267 Modified procedures: ❒ “Creating a Policy” on page 299 ❒ “Modifying a Policy” on page 302

Table 1. New Features in AT-S62 Version 1.4.0 (Continued)

Change	Chapter and Procedure
802.1x Port-based Network Access Control	
Added the following new features to 802.1x authenticator ports: ❒ Supplicant mode for supporting multiple supplicants on an authenticator port. For background information, see "Authenticator Ports with Single and Multiple Supplicants" on page 649. ❒ Guest VLAN. For background information, see "Guest VLAN" on page 657. ❒ VLAN Assignment and Secure VLAN for supporting dynamic VLAN assignments from a RADIUS authentication server for supplicant accounts. For background information, see "Supplicant and VLAN Associations" on page 655. ❒ MAC address-based authentication as an alternative to 802.1x username and password authentication. For background information, refer to "Authentication Modes" on page 646.	Chapter 29, "802.1x Port-based Network Access Control" on page 643 Modified procedure: ❒ "Configuring Authenticator Port Parameters" on page 665
Management Access Control List	
Simplified the menu interface for managing the access control entries in the Management ACL.	Chapter 35, "Management Access Control List" on page 759 Modified procedures: ❒ "Creating an ACE" on page 766 ❒ "Modifying an ACE" on page 768 ❒ "Deleting an ACE" on page 770

Chapter 1

Overview

This chapter reviews the functions of the AT-S62 management software, the types of management sessions supported by the switch, and the management access levels. This chapter contains the following sections:

Management Overview

The AT-S62 management software allows you to monitor and adjust the operating parameters of an AT-8500 Series switch and includes the following features:

- Basic operations such as configuring port and switch parameters, enhanced stacking, SNMPv1 and v2c, trunking, and mirroring
- Advanced operations including file uploads and downloads, event logging, traffic classifiers, access control lists, denial of service defense, Quality of Service (QoS), Class of Service (CoS), and IGMP snooping
- SNMPv3
- Spanning tree protocols including STP, RSTP, and MSTP
- Virtual LANs
- Port security options such as 802.1x Port-based Network Access Control and MAC address security levels
- Management security including encryption keys, PKI, SSL, Secure Shell, TACACS+, RADIUS, and management access control lists

The AT-S62 management software is preinstalled on the switch with default settings for all operating parameters. If the default settings are adequate for your network, you can use the device as an unmanaged switch by connecting it to your network, as explained in the hardware installation guide, and powering on the switch.

Note
The default settings for the management software can be found in Appendix A, “AT-S62 Default Settings” on page 773.

To actively manage a switch, you must connect to its management software. There are two general ways to connect to a switch:

- Locally using the RS232 Terminal Port on the switch
- Remotely using the Telnet protocol, the Secure Shell (SSH) protocol, or a web browser

The AT-S62 management software has three management interfaces. There is a menus interface, a command line interface, and a web browser interface. You can use the menus interface or the command line interface when managing the switch locally through the RS232 Terminal Port or remotely using the Telnet or SSH protocol. You use the web browser interface to manage the device with a web browser.

The following sections in this chapter briefly describe the different types of management sessions.

Local Management Session

To establish a local management session with an AT-8500 Series switch, you connect a terminal or a PC with a terminal emulator program to the RS232 Terminal Port on the switch, using the straight-through RS-232 management cable included with the unit. The RS232 Terminal Port is located on the front panel of the AT-8516F/SC, AT-8524M, and AT-8524POEswitches and the back panel of the AT-8550GB and AT-8550SP switches.

This type of management session is referred to as "local" because you must be physically close to the switch, such as in the wiring closet where the device is located.

Note
For instructions on starting a local management session, refer to "Starting a Local Management Session" on page 41.

A switch does not need an Internet Protocol (IP) address for you to manage it locally. You can start a local management session on a switch at any time and it will not affect the forwarding of frames by the device.

If you assign an AT-8500 Series switch an IP address and designate it as a master switch of an enhanced stack, you can manage all of the switches in the enhanced stack, all from the same local management session.

Note
For further information on enhanced stacking, refer to "Enhanced Stacking Overview" on page 80.

Telnet Management Session

You can remotely manage the switch from a workstation on your network using the Telnet application protocol. This type of management session is referred to in this guide as a remote management session because you do not have to be in the wiring closet where the switch is located.

To establish a Telnet management session with a switch, there must be at least one enhanced stacking switch in the subnet with an IP address. Only one switch in a subnet needs to have an IP address. Once you have established a Telnet management session with the switch that has an IP address, you can use the enhanced stacking feature of the management software to access all other enhanced stacking switches that reside in the same subnet.

Note
For further information on enhanced stacking, refer to “Enhanced Stacking Overview” on page 80.

Note
For instructions on how to start a Telnet management session, refer to “Starting a Telnet Management Session” on page 44.

A Telnet management session gives you access to nearly all of a switch’s operating parameters. You can perform nearly all the same functions from a Telnet management session as you can from a local management session.

Web Browser Management Session

You can also use a web browser from a management workstation on your network to manage a switch. This too is referred to as remote management because you can be anywhere on your network when managing the device.

This method of management, as with Telnet management, requires that the switch have an IP address or be part of an enhanced stack. Starting a web browser management session on a master switch of an enhanced stack allows you to manage all of the switches in the same enhanced stack, all from the same management session.

Note
For further information on the web browser interface, refer to the *AT-S62 Web Browser Interface User's Guide*.

SNMP Management Session

Another way to remotely manage the switch is with an SNMP management program. AT-S62 software supports SNMPv1, SNMPv2c, and SNMPv3. You need to be familiar with Management Information Base (MIB) objects to configure a switch using SNMP management.

The AT-S62 software supports the following MIBs:

- SNMP MIB-II (RFC 1213)
- Bridge MIB (RFC 1493)
- SNMPv3 (RFC 2571-6)
- User-based Security Model (USM) for SNMPv3 (RFC 2574)
- Interface Group MIB (RFC 2863)
- Ethernet MIB (RFC 1643)
- Remote Network MIB (RFC 1757)
- Allied Telesyn managed switch MIB

You must download the Allied Telesyn managed switch MIB files (atiChassisSwitch.mib and atiStackinginfo.mib) from the Allied Telesyn web site and compile the files with your SNMP program. For instructions on how to compile the MIB file with your SNMP program, refer to your SNMP management documentation.

For information about how to configure SNMP communities using a local or Telnet management session, see Chapter 5, “SNMPv1 and SNMPv2c Configuration” on page 89 and Chapter 21, “SNMPv3” on page 375.

Note
SNMP management can use the enhanced stacking feature through the private MIB (atiStackinginfo.mib). See Chapter 4, “Enhanced Stacking” on page 79.

Management Access Levels

There are two levels of management access in the AT-S62 management software: Manager and Operator. Manager access gives you the power to view and configure all of a switch's operating parameters. Operator access only allows you to view the operating parameters; you cannot change any values.

The switch has two default login accounts. For Manager access, the login name is "manager" and the default password is "friend". For Operator access, the login name is "operator" and the default password is also "operator". The usernames and passwords are case-sensitive.

You can create new Manager and Operator accounts with the RADIUS and TACACS+ authentication protocols, as explained in Chapter 34, "TACACS+ and RADIUS Authentication Protocols" on page 747.

Section I

Basic Operations

The chapters in this section cover a variety of basic switch features and functions. The chapters include:

Chapter 2

Starting a Local or Telnet Management Session

This chapter contains the procedures for starting a local or Telnet management session on an AT-8500 Series switch. The sections in the chapter are:

Local Management Session

To establish a local management session, you connect a terminal or PC with a terminal emulator program to the RS-232 terminal port on the switch. The RS232 Terminal Port is located on the front panel of the AT-8516F/SC, AT-8524M, and AT-8524POEswitches and the back panel of the AT-8550GB and AT-8550SP switches.

A local management session is so named because you must be close to the switch, usually within a few meters, to start this type of management session, meaning you must be in the wiring closet where the switch is located.

A switch does not need an IP address to be managed from a local management session, and a local management session will not interfere with the switch's forwarding of packets.

Starting a local management session on a switch configured as a Master switch allows you to manage all the switches in the same enhanced stack. This relieves you of having to start a separate local management session for each switch, simplifying network management.

Starting a local management session on a switch that is not part of an enhanced stack or that is a slave switch allows you to manage just that switch.

Note
For information on enhanced stacking, refer to "Enhanced Stacking Overview" on page 80.

Starting a Local Management Session

To start a local management session, perform the following procedure:

1. Connect one end of the straight-through RS232 management cable to the RS232 Terminal Port on the front panel of the switch.

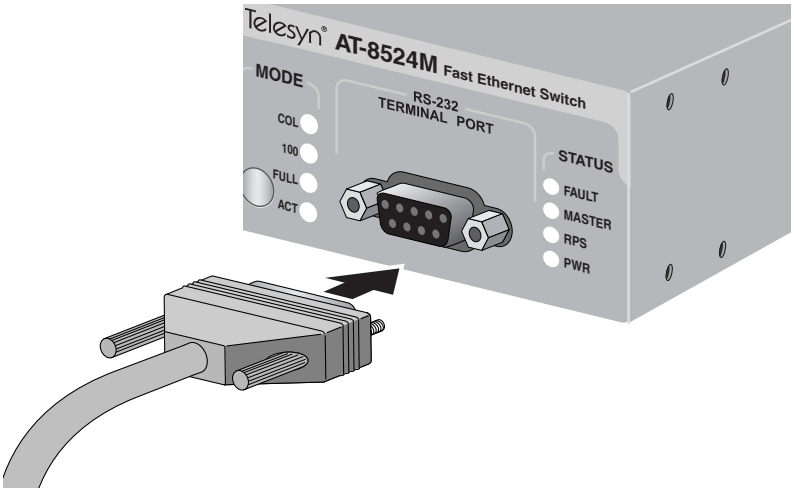

Figure 1. Connecting a Terminal or PC to the RS232 Terminal Port

2. Connect the other end of the cable to an RS-232 port on a terminal or PC with a terminal emulator program.

3. Configure the terminal or terminal emulator program as follows:

   - Baud rate: 9600 bps
   - Data bits: 8
   - Parity None
   - Stop bits: 1
   - Flow control: None

> **Note**
> The port settings are for a DEC VT100 or ANSI terminal, or an equivalent terminal emulator program.

> **Note**
> During boot up, the switch displays the following prompt: `Press <CTRL>B to go to Boot Prompt`. This message is intended for manufacturing purposes only. (If you inadvertently display the boot prompt (=>), type **boot** and press Return to start the switch.)

4. When prompted, enter a username and password.

   To configure the switch settings, enter "manager" as the user name. The default password for manager access is "friend". To just view the

settings, enter “operator” as the user name. The default password for operator access is “operator”. Usernames and passwords are case-sensitive. For information on the two access levels, refer to “Management Access Levels” on page 35. (For instructions on how to change a password, refer to “Configuring the Manager and Operator Passwords” on page 58.)

After logging on, you will see the window in Figure 2. This is the command prompt interface. You will see either a “#” symbol if you logged on as a manager or a “$” symbol if you logged on as an operator.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  <No System Name>
#
```

Figure 2. Command Prompt

For instructions on how to use the command line interface, refer to the *AT-S62 Command Line User’s Guide*, which is available from the Allied Telesyn web site.

5. To display the menu interface, type **menu** at the command prompt.

   The Main Menu is shown in Figure 3.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  <No System Name>
User: Manager                     11:20:02 02-Jan-2006
                        Main Menu
1 - Port Configuration
2 - VLAN Configuration
3 - Spanning Tree Configuration
4 - MAC Address Tables
5 - System Administration
6 - Advanced Configuration
7 - Security and Services
8 - Enhanced Stacking

C - Command Line Interface

Q - Quit

Enter your selection?
```

Figure 3. Main Menu

To select a menu item, type the corresponding letter or number.

Pressing the Esc key or typing the letter **R** in a submenu, returns you to the previous menu.

Enhanced Stacking

When you start a local management session on a switch configured as a Master switch, you can manage all the switches in the enhanced stack from the same management session. This saves you the time and trouble of having to start a separate local management session each time you want to manage a switch in your network. It also saves you from having to go to the different wiring closets where the switches are located.

For information on enhanced stacking and how to manage different switches from the same management session, refer to Chapter 4**,** "Enhanced Stacking" on page 79.

Quitting a Local Session

To quit a local session, return to the Main Menu and type **Q** for Quit.

You should always exit from a management session when you are finished managing a switch. This can prevent unauthorized individuals from making changes to a switch's configuration should you leave your management station unattended.

Note
You cannot run both a local management session and a Telnet management session on the same switch simultaneously. Failure to properly exit from a local or Telnet management session may block future management sessions.

Telnet Management Session

You can use the Telnet application protocol from a workstation on your network to manage an AT-8500 Series switch. This type of management is referred to as remote management because you do not have to be physically close to the switch to start the session, such as with a local management session. Any workstation on your network that has the application protocol can be used to manage the unit.

In terms of functionally, there are almost no differences between managing a switch locally through the RS232 Terminal Port and remotely with the Telnet application protocol. You see the same menu selections and have nearly the same management capabilities.

To manage a switch using Telnet, it must have an IP address or be part of an enhanced stack.

Note
For background information on enhanced stacking, refer to "Enhanced Stacking Overview" on page 80.

Starting a Telnet Management Session

To start a Telnet management session, specify the IP address of the Master switch of the enhanced stack in the Telnet application protocol and enter a user name and password when prompted.

To configure a switch's settings, enter "manager" as the user name. The default password for manager access is "friend". To just view the settings, enter "operator" as the username. The default password for operator access is "operator". User names and passwords are case-sensitive. For information on the two access levels, refer to "Management Access Levels" on page 35.

The management software displays the command line prompt shown in Figure 2 on page 42. For instructions on how to use the command line interface, refer to the *AT-S62 Command Line User's Guide*, available from the Allied Telesyn web site.

To use the menu interface instead, type **menu** and press Return. The Main Menu of a Telnet management session is the same menu for a local management session, shown in Figure 3 on page 42. You can perform nearly all the same functions from a Telnet management session as you can from a local management session.

The menus also function the same. To make a selection, type its corresponding number of letter. To return to a previous menu, type **R** or press ESC.

Note
You can run only one Telnet management session on a switch at a time. Additionally, you cannot run both a Telnet management session and a local management session on the same switch at the same time.

Quitting a Telnet Management Session

To end a Telnet management session, return to the Main Menu and type **Q** for Quit.

Saving Your Parameter Changes

When you make a change to a switch parameter, the change is, in most cases, immediately activated on the switch as soon as you enter it. However, most parameter changes are initially saved only to temporary memory in the switch and will be lost the next time you reset or power cycle the unit. To permanently save your changes, you must select the S - Save Configuration Changes option from the Main Menu. The switch saves your changes to its active configuration file. You should select that menu option whenever you have made a change to a switch parameter that you want the switch to retain even when it is reset or power cycled. If you do not see the option in the Main Menu, there are no parameter changes to be saved.

Ports 49R and 50R on the AT-8550GB and AT-8550SP Switches

This section applies to the 10/100/1000Base-T twisted pair ports 49R and 50R and the SFP and GBIC slots on the AT-8550GB and AT-8550SP switches. Note the following when configuring these ports:

- Twisted pair ports 49R and 50R change to the redundant status mode when an SFP or GBIC module is installed and establishes a link with its end node. An SFP or GBIC port is only active while it has a valid link. At all other times the corresponding twisted pair port 49R or 50R is the active port.
- A twisted pair port and its corresponding SFP or GBIC module share the same configuration settings, including port settings, VLAN assignments, access control lists, and spanning tree. When an SFP or GBIC module becomes active, it operates with the same settings as its corresponding twisted pair port.
- An exception is port speed: If you disable Auto-Negotiation on the twisted pair port and set the speed and duplex mode manually, the speed reverts to Auto-Negotiation when you install an SFP or GBIC module and the module establishes a link with an end node.

Chapter 3

Basic Switch Parameters

This chapter contains a variety of information and procedures. There is a discussion on when to assign an IP address to a switch and the different ways to do it. There are also procedures for resetting the switch, activating the switch default settings, and more.

Sections in the chapter include:

When Does a Switch Need an IP Address?

One of the tasks to building or expanding a network is deciding which managed switches need to be assigned a unique IP address. The rule used to be that a managed switch needed an IP address if you wanted to manage it remotely, such as with the Telnet application protocol. However, if a network contained a lot of managed switches, having to assign each one an IP address was often cumbersome and time consuming. It was also often difficult keeping track of all the IP addresses.

The enhanced stacking feature of the AT-8000 Series, AT-8400 Series, and AT-9400 Series switches simplifies all this. With enhanced stacking, you only need to assign an IP address to one switch in each subnet in your network. The switch with the IP address is referred to as the Master switch of the enhanced stack. All switches in the same subnet share the IP address.

Starting a local or remote management session on the Master switch automatically gives you complete management access to all the other enhanced stacking switches in the same enhanced stack.

This feature has two primary benefits. First, it helps reduce the number of IP addresses you have to assign to your network devices. Second, it allows you to configure multiple switches through the same local or remote management session.

If your network consists of multiple subnets, you must assign a unique IP address to at least one switch in each subnet. The switch with the IP address will be the Master switch of that subnet.

When you assign a switch an IP address, you must also assign it a subnet mask. The switch uses the subnet mask to determine which portion of an IP address represents the network address and which the node address.

You must also assign the switch a gateway address if there is a router between the switch and the remote management workstation. This gateway address is the IP address of the router through which the switch and management station will communicate.

Note
For further information on enhanced stacking, refer to “Enhanced Stacking Overview” on page 80.

How Do You Assign an IP Address?

After you have decided which, if any, switches on your network need an IP address, you must access the AT-S62 software on the switches and assign the addresses. There are two ways in which a switch can obtain an IP address.

The first method is for you to assign the IP configuration information manually. The procedure for this is explained in "Configuring an IP Address and Switch Name" on page 52. Initially assigning an IP address to a switch can only be done through a local management session.

The second method is for you to activate the BOOTP or DHCP client software on the switch and have the switch automatically download its IP configuration information from a BOOTP or DHCP server on your network. This procedure is explained in "Activating the BOOTP or DHCP Client Software" on page 55.

Configuring an IP Address and Switch Name

The procedure in this section explains how to manually assign an IP address, subnet mask, and gateway address to the switch from a local or Telnet management session. (If you want the switch to obtain its IP configuration from a DHCP or BOOTP server on your network, go to the procedure “Activating the BOOTP or DHCP Client Software” on page 55.)

This procedure also explains how to assign a name to the switch, along with the name of the administrator responsible for maintaining the unit and the location of the switch.

To manually set a switch’s IP address, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch

User: Manager                  11:20:02 02-Jan-2006

              System Administration

1 - System Information
2 - System Configuration
3 - Console (Serial/Telnet) Configuration
4 - Web Server Configuration
5 - SNMP Configuration
6 - Authentication Configuration
7 - Management ACL
8 - Event Log
9 - System Utilities

R - Return to Previous Menu

Enter your selection?
```

Figure 4. System Administration Menu

2. From the System Administration menu, type **2** to select System Configuration.

The System Configuration menu is shown in Figure 5.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch

User: Manager                 11:20:02 02-Jan-2006

               System Configuration

1 - BOOTP/DHCP ............... DISABLE
2 - IP Address ............... 0.0.0.0
3 - Subnet Mask .............. 0.0.0.0
4 - Default Gateway ......... 0.0.0.0
5 - System Name .............. Production Switch
6 - Location ................. Bldg. 12 Rm. 201
7 - Administrator ............ Jane Smith
8 - Configure System Time
9 - Fan Control Configuration

A - ARP Cache Timeout ....... 400 seconds

R - Return to Previous Menu

Enter your selection?
```

Figure 5. System Configuration Menu

3. Adjust the parameters as desired.

Note
A change to any parameter in this menu, including the IP address, subnet mask, or gateway address, is activated immediately on the switch.

The parameters in the System Configuration menu are described below:

1 - BOOTP/DHCP
This selection activates and deactivates the BOOTP and DHCP client software on the switch. For information on this selection, refer to "Activating the BOOTP or DHCP Client Software" on page 55.

2 - IP Address
This parameter specifies the IP address of the switch. You must specify an IP address if you want the switch to function as the Master switch of an enhanced stack or if the switch is not part of an enhanced stack and you want to remotely manage it using a web browser, a Telnet utility, SSH, or an SNMP management program. The IP address must be entered in the format: xxx.xxx.xxx.xxx. The default value is 0.0.0.0. Alternatively, you can activate the BOOTP or DHCP client software and have the switch obtain its IP configuration from a BOOTP or DHCP server on your network. For instructions, refer to "Activating the BOOTP or DHCP Client Software" on page 55.

3 - Subnet Mask
This parameter specifies the subnet mask for the switch. You must specify a subnet mask if you assigned an IP address to the switch. The subnet mask must be entered in the format: xxx.xxx.xxx.xxx. The default value is 255.255.0.0.

4 - Default Gateway
This parameter specifies the default router's IP address. This address is required if you intend to remotely manage the switch from a management station that is separated from the switch by a router. The address must be entered in the format: xxx.xxx.xxx.xxx. The default value is 0.0.0.0.

5 - System Name
This parameter specifies a name for the switch (for example, Sales Ethernet switch). The name is displayed at the top of the AT-S62 management menus and pages. The name can be from 1 to 39 characters. The name can include spaces and special characters, such as exclamation points and asterisks. The default is no name. This parameter is optional.

Note
Allied Telesyn recommends that you assign each switch a name. Names can help you identify the various switches in your network and help you avoid performing a configuration procedure on the wrong switch.

6 - Location
This parameter specifies the location of the switch, (for example, 4th Floor - rm 402B). The location can be from 1 to 39 characters. The location can include spaces and special characters, such as dashes and asterisks. The default is no location. This parameter is optional.

7 - Administrator
This parameter specifies the name of the network administrator responsible for managing the switch. The name can be from 1 to 39 characters. It can include spaces and special characters, such as dashes and asterisks. The default is no name. This parameter is optional.

Note
Option "8 - Configure System Time" is described in "Setting the System Time" on page 61. Option "9 - Fan Control Configuration" is described in "Setting Fan Control" on page 69.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Activating the BOOTP or DHCP Client Software

The BOOTP and DHCP application protocols can simplify network management by automatically assigning IP configuration information, such as IP addresses and subnet masks, to your network devices.

An AT-8500 Series switch contains the client software for these protocols and can obtain its IP configuration information from a BOOTP or DHCP server on your network. If you activate this feature, the switch seeks its IP address and other IP configuration information from a BOOTP or DHCP server on your network whenever you reset or power ON the device.

Review the following prior to activating the BOOTP or DHCP client:

- The switch can be running either BOOTP or DHCP, but not both simultaneously.
- There must be a BOOTP or DHCP server residing on your network.
- The BOOTP or DHCP server must be a member of the switch's management VLAN. The BOOTP or DHCP server must be communicating with the switch through a tagged or untagged port of the switch's management VLAN. For further information, refer to "Specifying a Management VLAN" on page 579.
- Any static IP address, subnet mask, or gateway address manually assigned to the switch is deleted from the System Configuration menu and replaced with the value the switch receives from the BOOTP or DHCP server. If you later disable BOOTP or DHCP, these values are returned to their default settings.

BOOTP and DHCP services allow you to specify how the IP address is to be assigned to the switch. The choices are static and dynamic. If you choose static, the server always assigns the same IP address to the switch when the switch is reset or powered ON. This is the preferred configuration. Since the switch is always assign the same IP address, you will always know which IP address to use when you need to remotely manage the device.

If you choose dynamic, the server assigns any unused IP address that it has not already assigned to another device. This means that a switch might have a different IP address each time you reset or power cycle the device, making it difficult for you to remotely manage the unit.

Note
The BOOTP and DHCP client software is disabled by default on the switch.

To activate or deactivate the BOOTP or DHCP client software, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **2** to select System Configuration.

   The System Configuration menu is shown in Figure 5 on page 53.

3. From the System Configuration menu, type **1** to select BOOTP/DHCP.

   The following prompt is displayed:

   ```
   DHCP/BOOTP/DISABLE (1-DHCP, 2-BOOTP, 3-DISABLE) :
   ```

4. Type **1** to activate DHCP, **2** to activate BOOTP, or **3** to disable both application protocols. The default is disabled.

   Note
   If you activate the BOOTP or DHCP client software, the switch immediately begins to query the network for the corresponding server. The switch continues to query the network for its IP configuration until it receives a response.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Rebooting a Switch

This procedure reboots the switch.

Note
Any configuration changes not save will be lost once the switch reboots. To save your configuration changes, return to the Main Menu and type **S** to select Save Configuration Changes.

To reboot the switch, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration. The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities. The System Utilities menu is shown in Figure 6.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                 11:20:02 02-Jan-2006
                 System Utilities

1 - File Operations
2 - Downloads and Uploads
3 - Ping a remote system
4 - Reset to Factory Defaults
5 - Reboot the switch
6 - Networking Stack

R - Return to Previous Menu

Enter your selection?
```

Figure 6. System Utilities Menu

3. From the System Utilities menu, type **5** to select Reboot the switch. The following prompt is displayed:

   ```
   The switch is about to reboot. Do you want to
   proceed? [Yes/No] ->
   ```

4. Type **Y** to reboot the switch or **N** to cancel the procedure.

Caution
The switch does not forward traffic while it initializes its management software and reloads the active boot configuration file. This process can take several minutes to complete. Some packet traffic may be lost. When the switch is finished rebooting, you can reestablish your management session if you want to continue managing the unit.

Configuring the Manager and Operator Passwords

There are two levels of management access on an AT-8500 Series switch: manager and operator. When you log in as manager, you can view and configure all of a switch's operating parameters. When you log in as an operator, you can only view the operating parameters; you cannot change any values.

You log in as a manager or an operator by entering the appropriate username and password when you start an AT-S62 management session. The default password for manager access is "friend". The default password for operator access is "operator". Passwords are case-sensitive.

This section contains these two procedures:

- "Changing the Manager or Operator Password" on page 58
- "Resetting the Manager Password" on page 59

The first procedure allows you to change a manager or operator password. The second allows you to bypass the manager password in the event you lose or forget it.

Changing the Manager or Operator Password

To change the manager or operator password, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **6** to select Authentication Configuration.

3. From the Authentication Configuration menu, type **5** to select Passwords Configuration.

   The Passwords Configuration menu is shown in Figure 7.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                 Production Switch
User: Manager                  11:20:02 02-Jan-2006
              Passwords Configuration
1 - Set Manager Password
2 - Set Operator Password

R - Return to Previous Menu

Enter your selection?
```

Figure 7. Passwords Configuration Menu

4. Type **1** to change the Manager password or type **2** to change the Operator password.

5. When prompted, enter the current manager password. (This step does not apply for the operator password.)

6. When prompted, enter the new manager or operator password. The new password will be case-sensitive.

7. When prompted, re-enter the new password.

Note
A password can be from 0 to 16 alphanumeric characters. Passwords are case-sensitive. You should not use spaces or special characters, such as asterisks (*) or exclamation points (!), in a password if you will be managing the switch from a web browser. Many web browsers cannot handle special characters in passwords.

Resetting the Manager Password

This procedure explains how to reset the manager password if you lost or forgot it. Note the following about this feature:

- You must perform this procedure from a local management session. You cannot perform it through enhanced stacking or from a Telnet or web browser management session.
- If the AT-S62 management software detects another active management session when you perform this procedure, a message is displayed for the other user stating that the user will be logged off. Thus, this type of session takes precedence over any other user's management session.

Caution
This procedure gives any person with physical access to the switch the ability to access its management software without having to provide a username and password. For this reason, all AT-8500 Series switches should be maintained in a locked wiring closet or other secure location to prevent unauthorized management access.

Note
This procedure requires resetting the switch. Some network traffic may be lost.

To reset the manager password on a switch, perform the following procedure:

1. Establish a local management session with the switch.

2. Reboot the switch. For instructions, refer to “Rebooting a Switch” on page 57.

3. When the switch displays “Press <Ctrl> B to go to Boot prompt,” type **S** or **s**.

   The switch continues its normal boot up and initialization process. Once complete, the management software automatically logs you in with manager access and displays the command line prompt. You are not prompted for a login username or password.

4. Type **menu** to display the Main Menu.

5. Follow the procedure in “Changing the Manager or Operator Password” on page 58 to reset the manager password.

   This completes the procedure for resetting the manager password. You can continue to manage the switch or you can quit from the management session. You must use the new password the next time you log on to the switch to start another management session.

Setting the System Time

This procedure explains how to set the switch's date and time. Setting the date and time is a good idea if you plan to monitor the switch by viewing the events in the event log or if the events are going to be sent to a syslog server. The correct date and time is also important if the management software will be sending traps to your management workstation. Events and traps contain the date and time of when they occurred so that you know when they transpired. The current date and time is also important if you intend to use the Secure Sockets Layer (SSL) certificate feature described in Chapter 32, "PKI Certificates and SSL" on page 705, because certificates must contain the date and time of when they were created.

There are two ways to set the switch's date and time. One method is to set it manually. The drawback to this approach is that the switch loses the information whenever it is reset or power cycled. This means that you must reset the values whenever you reset the device.

The second method uses the Simple Network Time Protocol (SNTP). The AT-S62 management software comes with the client version of this protocol. You can configure the AT-S62 software to obtain the current date and time from an SNTP or Network Time Protocol (NTP) server located on your network or the Internet.

SNTP is a reduced version of the NTP. However, the SNTP client software in the AT-S62 management software is interoperable with NTP servers.

Note
The SNTP or NTP server must be a member of the management VLAN. The server must be communicating with the switch through an untagged or tagged port of the management VLAN.

To set the system time manually or to configure SNTP, do the following:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **2** to select System Configuration.

   The System Configuration menu is shown in Figure 5 on page 53.

3. From the System Configuration menu, type **8** to select Configure System Time.

The Configure System Time menu is shown in Figure 8.

```
      Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                      Production Switch
User: Manager                            11:20:02 02-Jan-2006
                    Configure System Time
1 - System Time ..................... 00:04:22 on 01-Jan-1980
2 - SNTP Status ..................... Disabled
3 - SNTP Server ..................... 0.0.0.0
4 - UTC Offset ...................... +0
5 - Daylight Savings Time (DST) ... Enabled
6 - Poll Interval ................... 600 seconds
7 - Last Delta ...................... +0 seconds

U - Update System Time
R - Return to Previous Menu

Enter your selection?
```

Figure 8. Configure System Time Menu

4. To set the system time manually, do the following:

   a. Type **1** to select System Time

      The following prompt appears:

      ```
      Enter new system time [hh:mm:ss] ->
      ```

   b. Enter a new time for the system in the following format: hours, minutes, and seconds all separated by colons.

      The following prompt appears:

      ```
      Enter new system date [dd-mm-yyyy] ->
      ```

   c. Enter a new date for the system. Use two numbers to specify the day and month. Use four numbers to specify the year. Separate the values with hyphens. For example, December 5, 2003 is specified 05-12-2003.

      The new time and date are immediately activated on the switch.

5. To configure the switch to obtain its date and time from an SNTP or NTP server on your network or the Internet, do the following:

   a. Type **3** to select SNTP Server to enter the IP address of an SNTP server.

Note
If the switch is obtaining its IP address and subnet mask from a DHCP sever, you can configure the DHCP server to provide the switch with an IP address of an NTP or SNTP server. If you configured the DHCP server to provide this address, then you do not need to enter it here, and you can skip ahead to Step C.

The following prompt is displayed:

```
Enter SNTP server IP address ->
```

b. Enter an IP address of an SNTP or NTP server.

c. Type **4** to select UTC Offset to specify the difference between the UTC and local time.

Note
If the switch is using DHCP, it automatically attempts to determine this value. In this case, you do not need to configure a value for the UTC Offset parameter.

The following prompt is displayed:

```
Enter UTC Offset [-12 to 12] -> 0
```

d. Enter a UTC Offset time.

The default is 0 hours. The range is -12 to +12 hours.

e. Type **5** to select Daylight Savings Time (DST) to enable or disable the switch's ability to adjust its system time to daylight savings time. The following prompt is displayed:

```
Adjust for Daylight Savings Time (E - Enabled,
D - Disabled) ->
```

f. Select one of the following:

E - Enabled to allow the switch to adjust system time to daylight savings time. This is the default value.

D - Disabled to not allow the switch to adjust system time to daylight savings time.

Note
The switch does not set DST automatically. If the switch is in a locale that uses DST, you must remember to enable this in April when DST begins and disable it in October when DST ends. If the switch is in a locale that does not use DST, this option should be set to disabled all the time.

g. Type **6** - Poll Interval to specify the time interval between queries to the SNTP server.

   The following prompt is displayed:

   ```
   Enter interval to poll SNTP server [60 to 1200]
   -> 600
   ```

h. Enter the number of seconds the switch waits between polling the SNTP or NTP server. The default is 600 seconds. The range is from 60 to 1200 seconds.

i. Type **2** to select SNTP Status to enable or disable the SNTP client.

   The following prompt appears:

   ```
   SNTP Status (E-Enabled, D-Disabled) ->
   ```

j. Select one of the following:

   E - Enables the SNTP client software on the switch.

   D - Disables the SNTP client software

   Once enabled, the switch immediately polls the SNTP or NTP server for the current date and time. (The switch will also automatically poll the server whenever a change is made to any of the parameters in this menu, so long as SNTP is enabled.)

The Last Delta option in the menu displays the last adjustment that was applied to system time due to a drift in the system clock between two successive queries to the SNTP server. This is a read only field.

The U - Update System Time selection in the menu allows you to prompt the switch to poll the SNTP or NTP server for the current time and date. You can use this selection to update the time and date immediately rather than wait for the switch's next polling period. This selection has no effect if you set the date and time manually.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the Console Startup Mode

You can configure the AT-S62 software to initially display either the Main Menu or the command line interface prompt when you start a local, Telnet, or SSH management session. The default is the command line interface.

To change the console startup mode, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **3** to select Console (Serial/Telnet) Configuration.

   The Console (Serial/Telnet) Configuration menu is shown in Figure 9.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                  11:20:02 02-Jan-2006
       Console (Serial/Telnet) Configuration
1 - Console Startup Mode ............. CLI
2 - Console Disconnect Interval ..... 10 minute(s)
3 - Console Baud Rate ................ 9600
4 - Telnet Server .................... Enabled

R - Return to Previous Menu

Enter your selection?
```

Figure 9. Console (Serial/Telnet) Configuration Menu

3. Type **1** to toggle Console Startup Mode between Menu and CLI. When set to Menu, a management session starts by displaying the Main Menu. When set to CLI, a management session starts with the command line interface prompt. The default is CLI.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

   A change to the console startup mode takes effect the next time you start a management session.

Configuring the Console Timer

The AT-S62 management software uses the console timer, also referred to as the console disconnect interval, to automatically end inactive local and remote management sessions. The management software automatically ends a local or remote management session if a management session is inactive for the length of time specified by the console timer. For example, specifying two minutes for the console timer causes the AT-S62 management software to automatically end a management session if it does not detect any activity from the local or remote management station after two minutes.

This security feature prevents unauthorized individuals from using your management station should you step away from your system while configuring a switch. The default for the console timeout value is 10 minutes.

To adjust the console timer, do the following:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **3** to select Console (Serial/Telnet) Configuration.

   The Console (Serial/Telnet) Configuration menu is shown in Figure 9 on page 65.

3. From the Console (Serial/Telnet) Configuration menu, type **2** to select Console Disconnect Interval and, when prompted, enter a new console timer value. The range is 1 to 60 minutes. The default is 10 minutes.

   A change to the console timer is immediately activated on the switch.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Enabling or Disabling the Telnet Server

This procedure explains how to enable or disable the Telnet server on the switch. You might disable the server to prevent individuals from managing the switch with the Telnet application protocol or if you intend to use the Secure Shell (SSH) protocol.

> **Note**
> You cannot disable the Telnet server if there is an active Telnet management session on the switch.

To enable or disable the Telnet server, do the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **3** to select Console (Serial/Telnet) Configuration.

   The Console (Serial/Telnet) Configuration menu is shown in Figure 9 on page 65.

3. Type **4** to toggle Telnet Server between Enabled and Disabled. The default is enabled.

   A change to the Telnet server is immediately activated on the switch.

4. After making the change, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Setting the Baud Rate of the RS-232 Terminal Port

The default baud rate of the RS-232 Terminal Port on the switch is 9600 bps. To change the baud rate, do the following:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **3** to select Console (Serial/Telnet) Configuration.

   The Console (Serial/Telnet) Configuration menu is shown in Figure 9 on page 65.

3. From the Console (Serial/Telnet) Configuration menu, type **3** to select Console Baud Rate.

   The following message is displayed:

   ```
   Supported baud rates are:
   1200, 2400, 4800, 9600, 19200, 38400, 57600, or 115200
   Enter new baud rate value --> [1200 to 115200]
   ```

4. Type the desired baud rate value and press Return.

   The following message is displayed:

   ```
   Baud rate changed to [baud rate you typed] bps.
   Please change your terminal baud rate correspondingly.
   Press <Enter> to continue.
   ```

   Note
   If you are running a local management session, be sure to change your terminal's baud rate.

   A change to the baud rate is automatically saved to permanent memory in the switch. You do not need to use the Save Configuration Changes option in the Main Menu to permanently save this change.

Setting Fan Control

The AT-8524POE switch has a fan control feature that automatically adjusts the speed of four of its five cooling fans based on the ambient temperature of the room or wiring closet where the unit is installed and the load requirements of the PoE devices connected to the ports on the device. The lower the ambient temperature and load requirements of the powered devices, the lower the fan speed required by the system to maintain proper cooling.

The purpose of this feature is to decrease fan noise from the unit by taking advantage of building and networking environments where a reduction in fan speed will not compromise system cooling. A decrease in fan noise can lessen the chance of the switch being an annoyance to individuals when the device is installed in a public or work area.

When the fan control feature is deactivated, the default setting, the cooling fans operate at maximum speed at all times. When activated, fan speeds are continuously adjusted according to the ambient temperature as measured at the point where the air enters the cooling vents on the side of the switch, and the current load requirement of the PoE devices.

The four cooling fans controlled by this feature operate as a unit and have an operating range of approximately 5,000 to 11,000 RPM. The fans are operated at full speed when the ambient temperature reaches 40° C (104° F) or the PoE load exceeds 8.5 amps.

If a fan in a switch fails when the fan control feature is activated, the switch proportionally increases the speed of the remaining operational fans to compensate for the failed fan.

The fifth cooling fan is not controlled by this feature and operates are full speed at all times.

Enabling and Disabling Fan Control

To enable or disable fan control, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **2** to select System Configuration.
3. From the System Configuration menu, type **9** to select Fan Control Configuration.

The Fan Control Configuration menu is shown in Figure 10.

```
Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                Production Switch
User: Manager                  11:20:02 02-Jan-2006
            Fan Control Configuration
1 - Fan Control ..................... Off
2 - Show Fan Control Status

R - Return to Previous Menu

Enter your selection?
```

Figure 10. Fan Control Configuration Menu

4. Type **1** to toggle the fan control feature On or Off. The default setting is Off.

   A change to the status of the fan control feature is immediately implemented on the switch.

5. After making the change, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying Fan Control Status

To view the status of the fan control feature and the cooling fans in the AT-8524POE switch, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **2** to select System Configuration.
3. From the System Configuration menu, type **9** to select Fan Control Configuration.

   The Fan Control Configuration menu is shown in Figure 10.

4. Type **2** to select Show Fan Control Status.

Figure 11 illustrates the fan control information.

```
   Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                    Production Switch
 User: Manager                        11:20:02 02-Jan-2006
                  Show Fan Control Status

Fan Control Mode: Off
              Speed
Fan#      RPM       %           Status
-----------------------------------------
1         10700     100         Ok
2         10750     100         Ok
3         10700     100         Ok
4         10700     100         Ok
5         6200      100         Ok

Temperature = 24 C, PoE Current Load = 2.2 Amps (Max. 8.5)

U - Update System Time
R - Return to Previous Menu

Enter your selection?
```

Figure 11. Show Fan Control Status

The information is defined here:

- Fan Control Mode - The status of the fan control feature. If Off, the feature is disabled and all fans are operating at their maximum possible speed. If On, the feature is activated and the switch is adjusting the speed of the fans, as dictated by the ambient temperature and PoE load requirements.
- Fan# - The fan number. The system has five cooling fans. Fans 1 to 4 can be managed by the fan control feature. Fan 5 operates at its highest possible speed at all times.
- RPM - The current speed of the fan in revolutions per minute (RPM). The highest speed is approximately 11,000 RPM for fans 1 to 4 and 6,500 RPM for fan 5. RPM is displayed in increments of 100. The minimum operating speed for a fan is 4,000 RPM. A fan falling below or unable to attain that speed is considered as failed.
- % - The speed of the fan as a percentage of its highest possible operating speed.
- Status - A status message.
- Temperature - The ambient air temperature measured where the air enters the cooling vents of the switch.
- PoE Current Load - The total current load of the PoE devices.

Pinging a Remote System

You can instruct the switch to ping a remote device on your network. This procedure is useful in determining whether a valid link exists between the switch and another device. Note the following before performing the procedure:

- ❒ The switch must have an IP address.
- ❒ The device being pinged must be a member of the management VLAN. This means the device must be communicating with the switch through an untagged or tagged port of the management VLAN.

To instruct the switch to ping a network device, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. For the System Utilities menu, type **3** to select Ping a Remote System.

   The following prompt is displayed:

   ```
   Please enter an IP address ->
   ```

4. Enter the IP address of the end node you want the switch to ping.

   The results of the ping command are displayed on the screen.

5. To stop the ping, press any key.

Returning the AT-S62 Software to the Factory Default Values

There are two procedures for returning the settings on a switch to the factory default values. The first returns the switch's settings to the default values, but retains all files in the switch's file system (i.e., configuration files, SSL certificates, event logs, etc). The second method deletes all the files in the file system, including all configuration files. The AT-S62 software default values can be found in Appendix A**,** "AT-S62 Default Settings" on page 773.

Retaining the System Files

This procedure returns all operating parameters on the switch back to their default values, but retains the files in the file system. Review the following before performing this procedure:

- A switch's IP address and subnet mask, if assigned, are deleted.
- All port-based and tagged VLANs are deleted.
- All files in the AT-S62 file system are retained.
- All encryption keys stored in the key database are retained.
- The contents of the active boot configuration file is retained. To reset the file back to the default settings, you need to reestablish your management session after the switch reboots at the completion of this procedure and select Save Configuration Changes. Otherwise, the switch will revert back to the previous configuration the next time you reset the unit.

Caution

This procedure results in a switch reset. The switch will not forward traffic while it initializes its operating software, a process that can take approximately 20 seconds to complete. Some network traffic may be lost.

To return the AT-S62 software to the default settings while retaining the files in the file system, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. For the System Utilities menu, type **4** to select Reset to Factory Defaults.

The following prompt is displayed:

```
This operation requires a switch reboot. Continue?
[Yes/No] ->
```

4. Type **Y** for yes or **N** to cancel the procedure.

   If you respond with yes, the following prompt is displayed:

   ```
   Do you want to reset serial baud rate to 9600 bps?
   [Yes/No] ->
   ```

5. Typing **Y** for yes will change the baud rate of the RS232 Terminal Port to its default value of 9600 bps. Typing **N** leaves the baud rate at its current setting.

   The following prompt is displayed:

   ```
   NOTE: Please save configuration after reboot in
   order to make the configuration changes
   permanent!!!
   Waiting for background file operations to complete
   .....
   Rebooting the Switch .....
   ```

   Once the reset process is complete, the unit is again operating with its default settings.

6. Reestablish your management session.

7. From the Main Menu, type **S** to select Save Configuration Changes. This step returns the active boot configuration file back to the default settings. If you omit this step, the switch will revert back to the prior configuration the next time you reset or power cycle the unit.

Deleting the System Files

This procedure deletes all of the files in the switch's file system and resets the switch. This process returns the switch's operating parameters to their default settings.

> **Note**
> To return the switch to its default setting without deleting the files in the file system, perform the procedure "Retaining the System Files" on page 73.

Please note the following before performing this procedure:

- A switch's IP address and subnet mask, if assigned, are deleted.
- All port-based and tagged VLANs are deleted.
- All files in the AT-S62 file system are deleted.
- All encryption keys stored in the key database are deleted.

- The current speed setting of the RS232 console port on the switch is retained.

Caution

This procedure results in a switch reset. The switch will not forward traffic while it initializes its operating software, a process that takes approximately 20 seconds to complete. Some network traffic may be lost.

To delete all files from the file system and return the switch's operating parameters to the default settings, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. Form the System Administration menu, type **9** to select System Utilities.
3. For the System Utilities menu, type **1** to select File Operations.
4. From the File Operations menu, type **9** to select Format Flash Drive.

   The following prompt is displayed:

   ```
   This command will format the flash drive and
   requires a switch reboot.
   Do you want to continue ?  [Yes/No] ->
   ```

5. Type **Y** to proceed or **N** to cancel the procedure.

   If you type Y for yes, the switch deletes all of the files in the file system and then resets. After the system has reinitialized, all switch settings are returned to their default settings.

Viewing System Hardware and Software Information

The procedure in this section displays hardware and software information about the switch. The information includes the switch's serial number and MAC address, as well as the status of the power supply and fan.

To display this information, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **1** to select System Information.

   The System Information menu is shown in Figure 12.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                        11:20:02 02-Jan-2006
                             System Information

MAC Address ..... 00:30:84:01:00:00    IP Address ....... 167.11.11.11
Model Name ...... AT-8524M             Subnet Mask ...... 255.255.255.0
Serial Number ... S05525A023600000     Gateway .......... 0.0.0.0
                                       System Up Time ... 6D:11H:47M:34S

Bootloader ...... ATS62_LOADER v1.2.0  Build Date ....... Nov 14 2005 15:56:24
Application ..... ATS62 v1.4.0         Build Date ....... Jan 13 2006 17:57:17

System Name ..... Production Switch
Administrator ... John Doe
Location ........ Bldg. 5, Floor 4

H - System Hardware Status
U - Uplink Information

R - Return to Previous Menu

Enter your selection?
```

Figure 12. System Information Menu

   You cannot change the information in this menu.

3. To display system hardware information, type **H** to select System Hardware Status.

The System Hardware Information menu is shown in Figure 13.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                  11:20:02 02-Jan-2006

              System Hardware Status

System 1.8V Power ............... 1.79V
System 2.5V Power ............... 2.53V
System 3.3V Power ............... 3.30V
System 5V Power ................. 5.07V
System Temperature (Celsius) .... 30C
System Fan 1 Speed .............. 4720 RPM
System Fan 2 Speed .............. Off
Main Power Supply ............... AC - On
Redundant Power Supply .......... Not Present

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 13. System Hardware Information Menu

You cannot change the information in this menu. Note the following:

- The number of fans vary by model. The AT-8516F/MT, AT-8516F/SC, and AT-8524M switches have one fan, the AT-8524POE switch has five fans, and the AT-8550GB and AT-8550SP switches have two fans.
- The Redundant Power Supply status will be "Not Present" if the switch is not connected to an RPS unit. if the switch is connected to an RPS unit, the status will be "On," even when the RPS module itself is powered off.

Chapter 4

Enhanced Stacking

This chapter explains the enhanced stacking feature. The sections in this chapter include:

Enhanced Stacking Overview

The enhanced stacking feature can make it easier for you to manage the AT-8500 Series switches in your network. It offers the following benefits:

- You can manage up to 24 switches from one local or remote management session. This eliminates the need of having to initiate a separate management session with each switch in your network.
- The switches can share the same IP address. This reduces the number of IP addresses you have to assign to your network devices for remote management.
- Remotely managing a new switch in your network is simplified. You simply connect it to your network. Once connected to the network, you can begin to manage it immediately from any workstation in your network.

Guidelines

There are a few guidelines to keep in mind when implementing enhanced stacking for your network:

- An enhanced stack cannot span subnets.
- All of the switches in an enhanced stack must use the same Management VLAN. For information about Management VLANs, refer to “Specifying a Management VLAN” on page 579.
- You can create multiple enhanced stacks within a subnet by assigning the switches to different Management VLANs.
- An enhanced stack must have at least one master switch.
- The master switch can be any switch that supports enhanced stacking, such as an AT-8000 Series, AT-8400 Series, AT-8500 Series, or AT-9400 Series switch.
- You should assign the master switch an IP address and subnet mask.

Note
No IP address is required if you intend to manage an enhanced stack solely through the RS232 Terminal Port on a master switch. However, remote management of a stack using Telnet, a web browser, or an SNMP application does require assigning a master switch an IP address and subnet mask.

- You must set a master switch’s stacking status to Master. For instructions, refer to “Setting a Switch’s Enhanced Stacking Status” on page 83.
- The enhanced stacking feature uses the IP address 172.16.16.16. Do not assign this address to any device if you intend to use the enhanced stacking feature.

There are three basic steps to implementing this feature on your network:

1. You must select a switch to function as the master switch of the enhanced stack.

   The master switch can be any switch that supports enhanced stacking, such as an AT-8000 Series, AT-8400 Series, AT-8500 Series, or AT-9400 Series switch. For networks that consist of more than one subnet, there must be at least one master switch in each subnet.

   It is recommended that each enhanced stack have two master switches, each assigned a unique IP address. That way, should you remove one of the master switches from the network, such as for maintenance, you all still be able to remotely manage the switches in the stack using the other master switch.

2. You should assign each master switch a unique IP address and a subnet mask.

   A master switch should have a unique IP address and a subnet mask. The other switches in an enhanced stack, referred to as slave switches, do not need an IP address. If an enhanced stack will have more than one master switch, you should assign each master switch a unique IP address.

   You can set the IP address manually or activate the BOOTP or DHCP service on the master switch and have the switch obtain its IP information from a BOOTP or DHCP server on your network. Initially assigning an IP address or activating the BOOTP and DHCP services can only be performed through a local management session of the master switch.

   For instructions on how to set the IP address manually, refer to "Configuring an IP Address and Switch Name" on page 52. For instructions on activating the BOOTP or DHCP service, refer to "Activating the BOOTP or DHCP Client Software" on page 55.

   Note
   No IP address is required if you intend to manage an enhanced stack solely through the RS232 Terminal Port on a master switch. However, remote management using Telnet, a web browser, or an SNMP application does require assigning a master switch an IP address and subnet mask.

3. Change the enhanced stacking status of the master switch to Master.

   This is explained in "Setting a Switch's Enhanced Stacking Status" on page 83.

Figure 14 is an example of the enhanced stacking feature.

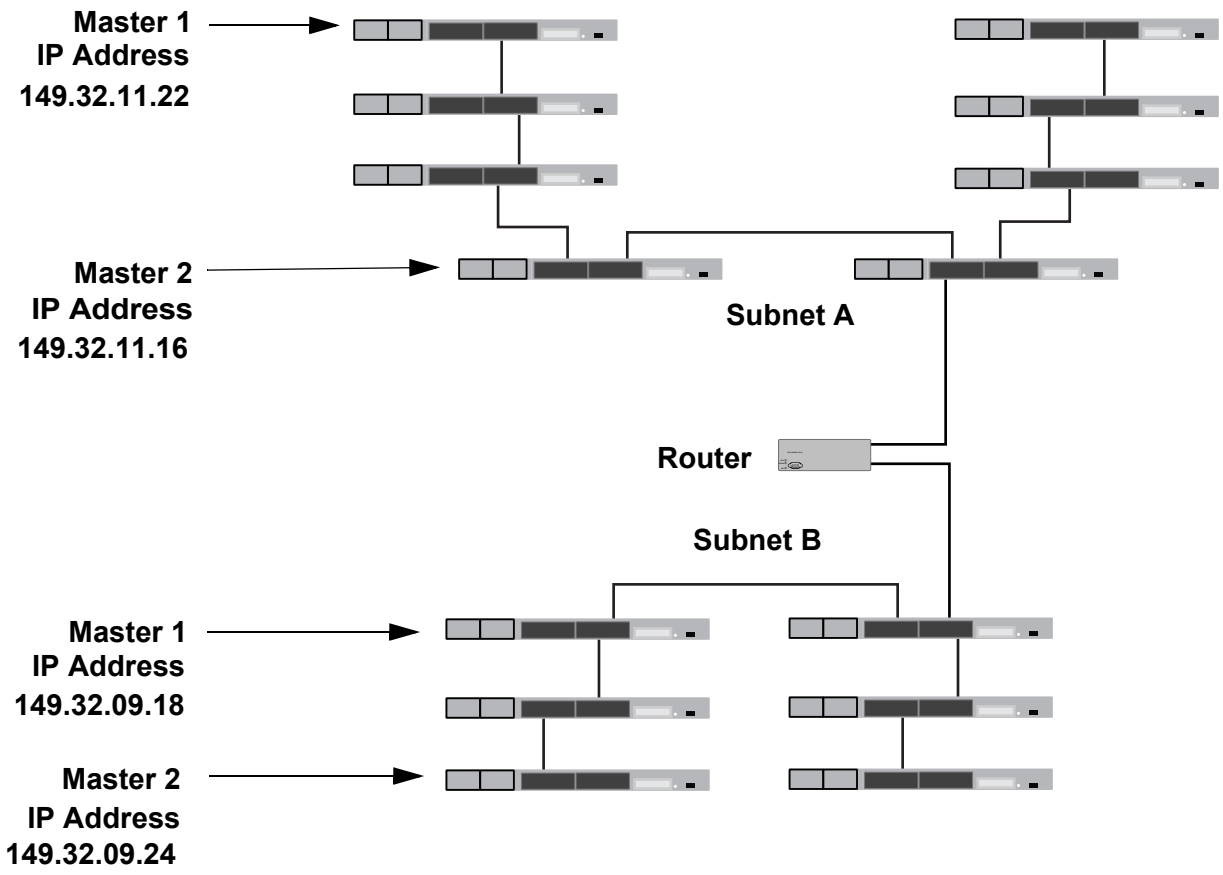

Figure 14. Enhanced Stacking Example

The example consists of a network of two subnets interconnected with a router. Two AT-8524M switches in each subnet have been selected as the master switches of their respective subnets, and each has been assigned a unique IP address.

To manage the switches of a subnet, you can start a local or remote management session on one of the master switches in the subnet. You would then have management access to all enhanced stacking switches in the same subnet.

Setting a Switch's Enhanced Stacking Status

The enhanced stacking status of the switch can be master switch, slave switch, or unavailable. Each status is described below:

- Master switch - A master switch of a stack can be used to manage all the other switches in a subnet. Once you establish a local or remote management session with the Master switch, you can access and manage all the switches in the stack.

  A master switch should have a unique IP address. You can manually assign a master switch an IP address or activate the BOOTP or DHCP client software on the switch.
- Slave switch - A slave switch can be remotely managed through a master switch. It does not need an IP address or subnet mask. This is the default setting.
- Unavailable - A switch with an unavailable stacking status cannot be remotely managed through a master switch. A switch with this designation can be managed locally. To be managed remotely, a switch with an unavailable stacking status must be assigned a unique IP address.

Note
You cannot change the stacking status of a switch through enhanced stacking. If a switch does not have an IP address or subnet mask, such as a slave switch, you must use a local management session to set its stacking status. If the switch has an IP address and subnet mask, such as a master switch, you can use either a local or remote management session.

To adjust a switch's enhanced stacking status, perform the following procedure:

1. From the Main Menu, type **8** to select Enhanced Stacking.

The Enhanced Stacking menu is shown in Figure 15.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                          11:20:02 02-Jan-2006
                     Enhanced Stacking
1 - Switch State-(M)aster/(S)lave/(U)navailable.... Master
2 - Stacking Services

R - Return to Previous Menu

Enter your selection?
```

Figure 15. Enhanced Stacking Menu

The menu displays the current status of the switch at the end of selection “1 - Switch State.” For example, the switch’s current status in the figure above is Master.

Note
The “2 - Stacking Services” selection in the menu is displayed only on master switches.

2. To change a switch’s stacking status, type **1** to select Switch State.

   The following prompt is displayed.

   ```
   Enter new setup (M/S/U) ->
   ```

3. Type **M** to change the switch to a master switch, **S** to make it a slave switch, or **U** to make the switch unavailable. Press Return.

   A change to the status is immediately activated on the switch.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Selecting a Switch in an Enhanced Stack

Before you perform a procedure on a switch in an enhanced stack, you should first check to be sure that you are performing it on the correct switch. If you assigned system names to your switches, this should be easy. The name of the switch being managed is always displayed at the top of every management menu.

When you start a local or remote management session on the Master switch of an enhanced stack, you are by default addressing that particular switch. The management tasks that you perform affect only the master switch.

To manage a slave switch or another Master switch in the stack, you need to select it from the management software.

To select a switch to manage in an enhanced stack, perform the following procedure:

1. From the Main Menu, type **8** to select Enhanced Stacking.
2. From the Enhanced Stacking menu, type **2** to select Stacking Services.

Note
The Stacking Services selection is only available on a Master switch.

The Stacking Services menu is shown in Figure 16.

```
       Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                              11:20:02 02-Jan-2006
                        Stacking Services
                                  Switch    Software   Switch
Num  MAC Address   Name           Mode      Version    Model
-------------------------------------------------------------


1 - Get/Refresh List of Switches
2 - Sort Switches in New Order
3 - Access Switch
4 - Load Image/Bootloader File
5 - Load Configuration File

R - Return to Previous Menu

Enter your selection?
```

Figure 16. Stacking Services Menu

3. Type **1** to select Get/Refresh List of Switches.

   The Master switch polls the subnet for all slave and Master switches that are a part of the enhanced stack and displays a list of the switches in the Stacking Services menu.

   The Master switch on which you started the management session is not included in the list, nor are any switches with an enhanced stacking status of Unavailable.

   By default, the switches are sorted in the menu by MAC address. You can sort the switches by name using the selection 2 - Sort Switches in New Order.

   Note
   Menu option “4 - Load Image/Bootloader File” uploads the AT-S62 image from the Master switch to another AT-8500 Series switch in the enhanced stack. The option is explained in “Uploading an AT-S62 Image File Switch to Switch” on page 196. Option “5 - Load Configuration File” allows you to upload a configuration file from a Master switch to another AT-8500 Series switch. This option is explained in “Uploading an AT-S62 Configuration File Switch to Switch” on page 199.

4. To manage a new switch, type **3** to select Access Switch.

   A prompt similar to the following is displayed:

   ```
   Enter the switch number -> [1 to 24}
   ```

5. Type the number of the switch in the list you want to manage.

6. Enter the appropriate username and password for the switch.

   The Main Menu of the selected switch is displayed. You now can manage the switch. Any management tasks you perform affect only the selected switch.

Returning to the Master Switch

When you have finished managing a slave switch, return to the Main Menu of the slave switch and type **Q** for Quit. This returns you to the Stacking Services menu. Once you see that menu, you are again addressing the Master switch from where you started the management session.

You can either select another switch in the list to manage or, if you want to manage the Master switch, return to the master switch's Main Menu by typing **R** twice.

Chapter 5

SNMPv1 and SNMPv2c Configuration

This chapter explains how to activate SNMP management on the switch and how to create, modify, and delete SNMPv1 and SNMPv2c community strings. Sections in the chapter include:

Note
For instructions on SNMPv3, refer to Chapter 21, “SNMPv3” on page 375.

SNMPv1 and SNMPv2c Overview

The Simple Network Management Program (SNMP) is another way for you to manage the switch. This type of management involves viewing and changing the management information base (MIB) objects on the device using an SNMP application program.

The AT-S62 management software supports SNMPv1, SNMPv2c, and SNMPv3. This chapter explains how to configure the switch's software for SNMPv1 and SNMPv2c. For instructions on how to configure the switch for SNMPv3, refer to Chapter 21, "SNMPv3" on page 375.

The procedures in this chapter show you how to create and manage SNMPv1 and SNMPv2c community strings through which your SNMP application program at your management workstation accesses the switch's MIB objects.

You can also configure SNMPv1 and SNMPv2c with the SNMPv3 Table menus described in Chapter 21, "SNMPv3" on page 375. However, because the SNMPv3 Table menus require a much more extensive configuration, Allied Telesyn recommends configuring SNMPv1 and SNMPv2c with the procedures in this chapter.

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. The procedure for this can be found in "Enabling or Disabling SNMP Management" on page 93.
- Load the Allied Telesyn MIBs for the switch onto your management workstation containing the SNMP application program. The MIBs are available from the Allied Telesyn web site at www.alliedtelesyn.com.

To manage a switch using SNMP, you need to know the IP address of the switch or of a master switch and at least one of the switch's community strings. A community string is a string of alphanumeric characters that gives you access to the switch.

A community string has several attributes that you can use to control 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 here.

Community String Name
You must give the community string a name. The name can be up to 32 alphanumeric characters. No spaces or special characters (such as /, #, or &) are allowed.

Access Mode
This defines what the community string will allow a network manager to do. 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 an unauthorized individual is using it to access the device. When a community string is enabled, it is available for use.

Open or Closed Access Status
You can use this feature to control which management stations on your network can use a community string. If you select the open access status, any network manager who knows the community string can use it. If you assign it a closed access status, then only those network managers working from particular workstations can use it. You specify the workstations by assigning their IP addresses to the community string. A closed community string can have up to eight IP addresses of management workstations assigned to it.

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 your 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 only Read.

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

Default SNMP Community Strings

The AT-S62 management software provides two default community strings: public and private. The public string has an access mode of just Read 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 its status from open to closed to prevent unauthorized changes to the switch.

Enabling or Disabling SNMP Management

To enable or disable SNMP management for the switch, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **5** to select SNMP Configuration.

   The SNMP Configuration menu is shown in Figure 17.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                   SNMP Configuration
1 - SNMP Status ......................... Disabled
2 - Authentication Failure Trap Status ..Disabled
3 - Configure SNMPv1 & SNMPv2c Community
4 - Display SNMPv1 & SNMPv2c Community
5 - Configure SNMPv3 Table
6 - Display SNMPv3 Table

R - Return to Previous Menu

Enter your selection?
```

Figure 17. SNMP Configuration Menu

3. Type **1** to toggle the SNMP Status option between its two settings of Enabled and Disabled. When set to Disabled, the default, you cannot manage the switch using SNMP. When set to Enabled, you can manage the switch using SNMP.

   A change to the SNMP status is immediately activated on the switch.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Setting the Authentication Failure Trap

As mentioned in the SNMP Overview section in this chapter, 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 to the switch using an SNMP application program. 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 will send this trap if either of the above events occur. If you disable it, the switch will 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.

To enable or disable the authentication trap, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **5** to select SNMP Configuration.

   The SNMP Configuration menu is shown in Figure 17 on page 93.

3. Type **2** to toggle Authentication Failure Trap Status between enabled and disabled. The default is disabled.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Creating an SNMP Community String

To create a new SNMP community string, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **5** to select SNMP Configuration.

   The SNMP Configuration menu is shown in Figure 17 on page 93.

3. From the SNMP Configuration menu, type **3** to select Configure SNMPv1 & SNMPv2c Community.

   The Configure SNMPv1 & SNMPv2c Community menu is shown in Figure 18.

```
              Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                            Production Switch
User: Manager                                          11:20:02 02-Jan-2006
                     Configure SNMPv1 & SNMPv2c Community
Community Name   AccessMode   Status   OpenAcc  Manager IP Addr   Trap Rec IP
-------------------------------------------------------------------------
Private          Read|Write   Enabled  Yes
Public           Read         \Enabled Yes

1 - Create SNMP Community
2 - Delete SNMP Community
3 - Modify SNMP Community

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 18. SNMPv1 & SNMPv2c Community Menu

   This menu lists the current community strings on the switch and their attributes. For attribute definitions, refer to "SNMPv1 and SNMPv2c Overview" on page 90.

4. Type **1** to select Create SNMP Community.

   This prompt is displayed:

   ```
   Enter SNMP Community Name:
   ```

5. Enter the new SNMP community string. The name can be up to 32 alphanumeric characters. No spaces or special characters (such as /, #, or &) are allowed.

   This prompt is displayed:

   ```
   Enter Access Mode [R-Read Only, W-Read/Write]:
   ```

6. Specify the access mode for the new SNMP community string. If you specify Read, the community string will only allow you to view the MIB objects on the switch. If you specify Read/Write, the community string will allow you to both view and change the SNMP MIB objects on the switch. This prompt is displayed:

   ```
   Enter Open Access Status [Y-Yes, N-No]:
   ```

7. Specify the open access status. If you enter Yes, any network manager who knows the community string can use it. If you respond with No, making it closed access, only those management workstations whose IP addresses you assign to the community string can use it. This prompt is displayed:

   ```
   Enter SNMP Manager IP Addr:
   ```

8. If in Step 7 you responded with No making this a closed community string, specify the IP address of the management workstation that can use the string. A community string can have up to eight IP addresses of management workstations. But you can assign only one to it initially with this procedure. To add additional IP addresses, refer to “Modifying a Community String” on page 98.

   If you assigned the community string an access status of open, leave this field blank by pressing Return.

   This prompt is displayed:

   ```
   Enter Trap Receiver IP Addr:
   ```

9. If you want the switch to send traps to a management workstation or server, enter the IP address of the node here. A community string can have up to eight IP addresses of trap receivers. But you can assign only one initially with this procedure. To add additional IP addresses, refer to “Modifying a Community String” on page 98.

   If you do not want to add a IP address of a trap receiver to the community string, leave this field blank by pressing Return.

   The AT-S62 software creates the new community string and adds it to the list in the SNMP Community menu. A new community string is immediately available for use to manage the switch.

10. If desired, repeat this procedure starting with Step 4 to create additional community strings.

11. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Community String

To modify a community string, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **5** to select SNMP Configuration.

   The SNMP Configuration menu is shown in Figure 17 on page 93.

3. From the SNMP Configuration menu, type **3** to select Configure SNMPv1 &SNMPv2c Community.

   The Configure SNMPv1 &SNMPv2c Community menu in shown in Figure 18 on page 95.

4. From the Configure SNMPv1 &SNMPv2c Community menu, type **3** to select Modify SNMP Community.

   The Modify SNMP Community menu is shown in Figure 19.

```
             Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                           Production Switch
User: Manager                                        11:20:02 02-Jan-2006
                     Modify SNMPv1 & SNMPv2c Community
Community Name   AccessMode   Status   OpenAcc  Manager IP Addr   Trap Rec IP
-------------------------------------------------------------------------
Private          Read|Write   Enabled  Yes
Public           Read         Enabled  Yes

1 - Add Attributes to Community
2 - Delete Attributes from Community
3 - Set Community Access Mode
4 - Set Community Status
5 - Set Community Open Access

U - Update Display
R - Return to Previous Menu

Enter your selection:
```

Figure 19. Modify SNMP Community Menu

This menu lists the current community strings on the switch and their attributes. For attribute definitions, refer to “SNMPv1 and SNMPv2c Overview” on page 90.

The menu options are described below:

1 - Add Attributes to Community
If a community string has a closed access mode, you can use this selection to add new IP addresses of management workstations that can use the string. You can also use this option to add IP addresses of new trap receivers. To use this option, do the following:

1. From the Modify SNMP Community menu, type **1** to select Add Attributes to Community. The following prompt is displayed:

   ```
   Enter SNMP Community Name:
   ```

2. Enter the community string you want to modify. Community strings are case sensitive. This prompt is displayed:

   ```
   Enter SNMP Manager IP Addr:
   ```

3. If you are modifying a community string with a closed access mode and you want to add an IP address of a management workstation to it, enter the workstation's IP address at the prompt. Otherwise, just press Return. A community string can have a maximum of eight IP addresses, but you can add only one at a time with this procedure. This prompt is displayed:

   ```
   Enter Trap Receiver IP Addr:
   ```

4. If you want the switch to send traps to a trap receiver, enter the IP address of the receiver at this prompt. Otherwise, just press Return.

   The community string is modified and the Modify SNMP Configuration menu is displayed again.

5. Repeat this procedure to modify other community strings.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

2 - Delete Attributes from Community
Use this option to delete an IP address of a management workstation or a trap receiver from a community string. To use this option, do the following:

1. From the Modify SNMP Community menu, type **2** to select Delete Attributes from Community. The following prompt is displayed:

   ```
   Enter SNMP Community Name:
   ```

2. Enter the community string you want to modify. Community strings are case sensitive. This prompt is displayed:

   ```
   Enter SNMP Manager IP Addr:
   ```

3. If you want to remove the IP address of a management workstation from the community string, enter the IP address at the prompt. Otherwise, just press Return. This prompt is displayed:

   ```
   Enter Trap Receiver IP Addr:
   ```

4. If you want to remove the IP address of a trap receiver from the community string, enter the IP address at the prompt. Otherwise, just press Return.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

3 - Set Community Access Mode
Use this option to change a community string's Read or Read/Write status. To use the selection, do the following:

1. From the Modify SNMP Community menu, type **3** to select Set Community Access Mode. The following prompt is displayed:

   ```
   Enter SNMP Community Name:
   ```

2. Enter the community string you want to modify. Community strings are case sensitive. This prompt is displayed:

   ```
   Enter Access Mode [R-Read Only, W-Read/Write]:
   ```

3. Type **R** to change the string's status to Read only, or **W** for Read/Write. This confirmation prompt is displayed:

   ```
   Do you want to change this Community Access Mode? (Y/N):
   [Yes/No] ->
   ```

4. Type **Y** to change the string's access mode or **N** to cancel the change.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

4 - Set Community Status
Use this option to enable or disable a community string. When disabled, no one can use the community string to access the switch. To use the selection, do the following:

1. From the Modify SNMP Community menu, type **4** to select Set Community Status. The following prompt is displayed:

   ```
   Enter SNMP Community Name:
   ```

2. Enter the community string you want to modify. Community strings are case sensitive. This prompt is displayed:

   ```
   Enter Community Status [E-Enable, D-Disable]:
   ```

3. Type **E** to enable the community string or **D** to disable it. This confirmation prompt is displayed:

   ```
   Do you want to change Community Status? (Y/N): [Yes/No] -
   >
   ```

4. Type **Y** to change the string's status or **N** to cancel the change.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

5 - Set Community Open Status
Use this selection to change a string's open status. A string with an open status can be used by any network administrator. A string with a closed status can only be used from management workstations whose IP addresses are assigned to the community string. To use the option, do the following:

1. From the Modify SNMP Community menu, type **5** to select Set Community Open Status. The following prompt is displayed:

   ```
   Enter SNMP Community Name:
   ```

2. Enter the community string you want to modify. Community strings are case sensitive. This prompt is displayed:

   ```
   Enter Open Access Status [Y-Yes, N-No]:
   ```

3. Type **Y** to assign the string an open status or **N** to assign it a closed status. This confirmation prompt is displayed:

   ```
   Do you want to change Open Access Status? (Y/N): [Yes/No]
   ->
   ```

4. Type **Y** to change the string's open status or **N** to cancel the change.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting a Community String

To delete an SNMPv1 or SNMPv2c community string, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **5** to select SNMP Configuration.

   The SNMP Configuration menu is shown in Figure 17 on page 93.

3. From the SNMP Configuration menu, type **3** to select Configure SNMPv1 &SNMPv2c Community.

   The Configure SNMPv1 &SNMPv2c Community menu in shown in Figure 18 on page 95.

4. From the Configure SNMPv1 &SNMPv2c Community menu, type **2** to select Delete SNMP Community. This prompt is displayed:

   ```
   Enter Trap Receiver IP Addr:
   ```

5. Enter the community string to be deleted. Community strings are case sensitive. A confirmation prompt is displayed.

6. Type Y for yes to delete the string or N for no to cancel the procedure.

   If you selected yes, the community string is immediately deleted from the switch.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying the SNMP Community Strings

To display the attributes of all the SNMP community strings on the switch, use the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **5** to select SNMP Configuration.

   The SNMP Configuration menu is shown in Figure 17 on page 93.

3. From the SNMP Configuration menu, type **4** to select Display SNMPv1 & SNMPv2c Community.

   The Display SNMPv1 & SNMPv2c Community menu is shown in Figure 20.

```
                  Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                                Production Switch
User: Manager                                                  11:20:02 02-Jan-2006
                          Display SNMPv1 & SNMPv2c Community
Community Name   Access Mode    Status    OpenAcc  Manager IP Addr      Trap Receiver IP
-------------------------------------------------------------------------------------
Private125       Read|Write    Enabled    No       147.41.11.30         147.45.16.70
                                                   147.45.16.80         147.45.16.80
PublicATI78      Read Only     Enabled    No       147.41.11.12         147.42.22.22
                                                   147.44.16.86         147.45.16.86
                                                   147.45.16.88         147.45.16.88
                                                   147.45.16.90         147.45.16.90
HighSchool2      Read|Write    Enabled    No       147.45.10.80         147.45.10.80

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 20. Display SNMP Community Menu

For attribute definitions, refer to "SNMPv1 and SNMPv2c Overview" on page 90.

Chapter 6

Port Parameters

The chapter contains the procedures for viewing and adjusting the parameter settings for the individual ports on a switch. It also describes how to display port statistics.

This chapter contains the following procedures:

Displaying Port Status

To display the current status and settings of the ports on the switch, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.

   The Port Configuration menu is shown in Figure 21.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                        11:20:02 02-Jan-2006
                    Port Configuration
1 - Port Configuration
2 - Port Status
3 - Port Statistics
4 - Port Trunking and LACP
5 - Port Security
6 - Port Mirroring

R - Return to Previous Menu

Enter your selection?
```

Figure 21. Port Configuration Menu

2. From the Port Configuration Menu, type **2** to select Port Status.

   An example of the Port Status menu is shown in Figure 22.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006
                             Port Status
Port  Link   Neg    MDIO   Speed   Duplex       PVID    Flow Ctl
-------------------------------------------------------------------
1     Up     Auto   Auto   0010    Half         0012    Disabled
2     Up     Auto   Auto   0100    Full         0012    Disabled
3     Up     Auto   Auto   0100    Full         0012    Disabled
4     Up     Auto   Auto   0100    Full         0023    Disabled
5     Up     Auto   Auto   0010    Half         0012    Disabled
6     Up     Auto   Auto   0100    Full         0011    Disabled
7     Up     Auto   Auto   0100    Full         0011    Disabled
8     Up     Auto   Auto   0010    Half         0011    Disabled

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 22. Port Status Menu

Note
The speed, duplex mode, and flow control settings will be blank for ports that have not established a link with their end node.

To view the settings of a GBIC or SFP module in Port 49 or 50 of an AT-8550GB or AT-8550SP switch, there must be a valid connection between the module's port and the end node. Otherwise, Ports 49 and 50 in the menu represent the twisted pair ports 49R and 50R.

The information in this menu is for viewing purposes only. The columns in the menu are described below:

Port

The port number.

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

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

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

Neg
The status of Auto-Negotiation on the port. Possible values are:

Auto - Indicates that the port is using Auto-Negotiation to set operating speed and duplex mode.

Manual - Indicates that the operating speed and duplex mode have been set manually.

MDIO
The operating configuration of the port. Possible values are Auto, MDI, MDI-X. The status Auto indicates that the port will automatically determine the appropriate MDI or MDI-X setting.

Speed
The operating speed of the port. Possible values are:

0010 - 10 Mbps

0100 - 100 Mbps

1000 - 1000 Mbps (Gigabit Ethernet ports only)

Duplex
The duplex mode of the port. Possible values are half-duplex and full-duplex.

PVID
The port's VLAN identifier (PVID). This number corresponds to the VID of the VLAN in which the port is an untagged member. This column will not include the VIDs of the VLANs where the port is a tagged member.

Flow Ctl
The flow control setting for the port. Possible values are:

Disabled - No flow control on the port.

Enabled - Flow control is activated.

Configuring Port Parameters

To configure the parameter settings of a port, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.

   The Port Configuration menu is shown in Figure 21 on page 106.

2. From the Port Configuration menu, type **1** to select Port Configuration.

   The following prompt is displayed:

   ```
   Enter port-list ->
   ```

3. Enter the number of the port you want to configure. You can specify more than one port at a time. You can specify the ports individually (for example, 5,7,22), as a range (for example, 18-23), or both (for example, 1,5,14-22).

   To configure a GBIC or SFP module in Port 49 or 50 of an AT-8550GB or AT-8550SP switch, there must be a valid connection between the port and the end node. Otherwise, specifying Port 49 or 50 configures the twisted pair port 49R or 50R, respectively.

   The Port Configuration menu is shown in Figure 23.

```
        Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                              11:20:02 02-Jan-2006
                       Port Configuration
Configuring Port 11
0 - Port Description ..................... Port-1
1 - Status ............................... Enabled
2 - Broadcast Filter ..................... Disabled
3 - MDI/MDIX Crossover ................... Auto
4 - Negotiation .......................... Manual
5 - Speed ................................ 0100
6 - Duplex ............................... Full
7 - HOL Blocking Prevention Threshold .... 61440 cells
8 - Flow Control
9 - Back Pressure
L - Rate Limiting
U - Unknown Multicast Filtering .......... Disabled
D - Set Default Port Configuration
F - Force Renegotiation
X - Reset Port
R - Return to Previous Menu

Enter your selection?
```

Figure 23. Port Configuration (Port) Menu

Selections 3, 5, and 6 appear in the menu only when selection 4 - Negotiation is set to Manual. When selection 4 is set to Auto, these options are hidden.

Note
The Port Configuration menu in the figure above is for a 10/100 Mbps twisted pair port. The menu for a fiber optic port will contain a subset of the parameters.

If you are configuring multiple ports and the ports have different settings, the Port Configuration menu displays the settings of the lowest numbered port. Once you have configured the settings of the port, all of its settings are copied to the other selected ports.

4. Adjust the port parameters as necessary. You adjust a parameter by typing its number. The parameters are described below.

Note
A change to a parameter is immediately activated on the port.

0 - Port Description
You use this selection to assign a name to a port. The name can be from one to fifteen alphanumeric characters. Spaces are allowed, but you should not use special characters, such as asterisks or exclamation points. (You cannot set a port name if you are configuring more than one port.)

1 - Status
You use this selection to enable or disable a port. When disabled, a port will not forward frames to or from the node connected to the port.

You might want to disable a port and prevent packets from being forwarded if a problem occurs with the node or cable connected to the port. Once the problem has been fixed, you can enable the port again to resume normal operation.

You might also want to disable a port that is not being used to secure it from unauthorized connections.

Possible settings for this parameter are:

Enabled The port will forward packets. This is the default setting.

Disabled The port will not forward packets.

2 - Broadcast Filter
Most frames on an Ethernet network are usually unicast frames. A unicast frame is a frame that is sent to a single destination. A node sending a unicast frame intends the frame for a particular node on the network.

Broadcast frames are different. Broadcast frames are directed to all nodes on the network or all nodes within a particular virtual LAN. Broadcast packets can perform a variety of functions. For example, some network operating systems use broadcast frames to announce the presence of devices on a network.

The problem with broadcast frames is that too many of them traversing a network can impact network performance. The more bandwidth consumed by broadcast frames, the less available for unicast frames.

Should the performance of your network be impacted by heavy broadcast traffic, you can use this parameter to limit the number of broadcast frames forwarded by the switch ports and so limit the number of broadcast frames on your network.

Activating this feature on a port discards all egress broadcast packets on the port.

Filtering only occurs on egress broadcast packets — packets that a port is transmitting. This filter is not applied to ingress broadcast packets.

Possible settings for this parameter are:

Enabled — The port discards all egress broadcast frames.

Disabled — The port transmits egress broadcast frames. This is the default setting.

3 - MDI/MDIX Crossover

You use this selection to set the wiring configuration of a twisted pair port. This option only appears when option 4 - Negotiation, which is used to activate and deactivate Auto-Negotiation, is set to Manual.

When selection 4 - Negotiation is set to Auto, which activates Auto-Negotiation on a port, this option is hidden in the menu and a twisted pair port uses auto-MDI/MDI-X to automatically set its wiring configuration. This feature enables a port to configure itself automatically as MDI or MDI-X when connected to an end node. This allows you to use a straight-through twisted pair cable when connecting any type of network device to a port on the switch.

The auto-MDI/MDI-X feature is only available when a port is using Auto-Negotiation to set its speed and duplex mode. It is also the only setting available when a port's speed and duplex are set through Auto-Negotiation.

If you set option 4 - Negotiation to Manual, which disables Auto-Negotiation on a port, the auto-MDI/MDI-X feature is disabled as well and this menu option appears with the two possible settings of MDI and MDI-X. The default is MDI-X.

4 - Negotiation
You use this selection to activate or deactivate Auto-Negotiation on a twisted pair port. This parameter has the two settings Auto and Manual. If you select Auto, a twisted pair port uses Auto-Negotiation to set its speed, duplex mode, and MDI/MDI-X settings. This is the default setting. If you select Manual, additional options appear in the menu for manually configuring these port settings. If you are configuring a fiber optic port, the only setting available is Manual.

You should note the following concerning the operation of Auto-Negotiation:

- In order for a twisted pair port to successfully Auto-Negotiate its duplex mode with an end node, the end node should also be using Auto-Negotiation. Otherwise, a duplex mode mismatch can occur. 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, you should disable Auto-Negotiation on the port and set the port's speed and duplex mode manually.

- When the port is set to Auto-Negotiate, the MDI/MDI-X setting is locked at auto-MDI/MDI-X. The switch automatically determines the correct MDI/MDI-X setting. You cannot manually set MDI/MDI-X manually.
- When Auto-Negotiation is disabled on a twisted pair port, the auto-MDI/MDI-X feature on a port is also disabled, and the port defaults to the MDI-X configuration. If you disable Auto-Negotiation and set a port's speed and duplex mode manually, you might also need to set the port's MDI/MDI-X setting as well.

5 - Speed
This selection is used to set the speed of a twisted pair port. It only appears when option 4 - Negotiation is set to Manual. The possible settings are:

0010 10 Mbps

0100 100 Mbps

You cannot change the speed of a fiber optic port.

Note
Ports 49R and 50R on an AT-8550GB Series switch must be set to Auto-Negotiation to operate at 1000Mbps. You cannot manually configure these ports to 1000Mbps.

6 - Duplex
This selection is used to set the duplex mode of a port. The option only appears when option 4 - Negotiation is set to Manual. The possible settings are:

Full Full-duplex

Half Half-duplex.

7 - HOL Blocking Prevention Threshold
Head of line (HOL) blocking is a problem that occurs when a port on a switch becomes oversubscribed. An oversubscribed port is receiving more packets from other switch ports than it can transmit in a timely manner.

The problem an oversubscribed port can create is that it can prevent other ports from forwarding packets to each other. This is because ingress packets on a port are buffered in a First In, First Out (FIFO) manner. If the head of an ingress queue consists of a packet destined for an oversubscribed port, the ingress queue will not be able to forward any of its other packets to the egress queues of other ports.

A simplified version of the problem is illustrated in Figure 24. It shows four ports on a switch. Port D is receiving packets from two ports—50% of the ingress traffic on Port A and 100% of the ingress traffic on Port B. The result is that not only is Port A unable to forward packets to Port D because the latter's egress queues are filled with packets from Port B, but it is also unable to forward traffic to Port C because its ingress queue has frames destined to Port D that it is unable to forward.

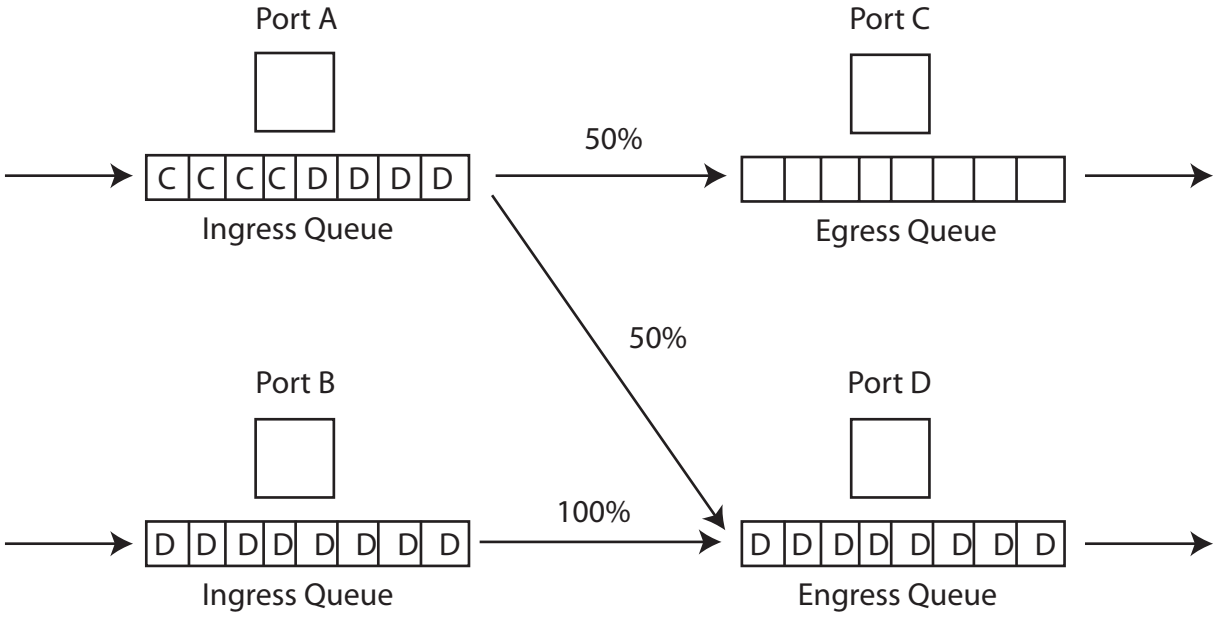

Figure 24. Head of Line Blocking

The HOL Limit parameter can help prevent this problem from occurring. This parameter sets a threshold on the utilization of a port's egress queue. When the threshold for a port is exceeded, the switch signals other ports to discard packets to the oversubscribed port.

For example, referring to the figure above, when the utilization of the storage capacity of Port D exceeds the threshold, the switch signals the other ports to discard packets destined for Port D. Port A drops the D packets, enabling it to once again forward packets to Port C.

The number for this value represents cells. A cell is 64 bytes. The range is 1 to 61,440 cells. The default is 61,440.

8 - Flow Control
Sets flow control on the port. This option applies only to ports operating in full-duplex mode.

A switch port uses flow control to control the flow of ingress packets from its end node.

A port using flow control issues a special frame, referred to as a PAUSE frame, as specified in the IEEE 802.3x standard, to stop the transmission of data from an end node. When a port needs to stop an end node from transmitting data, it issues this frame. The frame instructs the end node to cease transmission. The port continues to issue PAUSE frames until it is ready again to receive data from the end node.

The default setting for flow control on a switch port is disabled.

Selecting this option displays the Flow Control menu, shown in Figure 25.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                       11:20:02 02-Jan-2006
                        Flow Control

Configuring Port 11

1 - Flow Control ................... Disabled
2 - Flow Control (Cell Limit) .... 57344

R - Return to Previous Menu

Enter your selection?
```

Figure 25. Flow Control Menu

The options in the Flow Control menu are described below:

1 - Flow Control
Disabled - No flow control on the port. This is the default setting.

Enabled - Flow control is activated. This setting is appropriate only when the end node connected to the port is also using flow control.

Auto - The port uses flow control only if it detects that the end node is using it.

2 - Flow Control (Cell Limit)
Specifies the number of cells. A cell represents 64 bytes. The range is 1 to 57,344 cells. The default is 57,344.

B -Back Pressure
Sets backpressure on a port. This option only applies for ports operating in half-duplex mode.

Backpressure performs much the same function as flow control. Both are used by a port to control the flow of ingress packets from the end node.

Where they differ is that while flow control applies to ports operating in full-duplex, backpressure applies to ports operating in half-duplex mode.

When a twisted pair port on the switch operating in half-duplex mode needs to stop an end node from transmitting data, it forces a collision. A collision on an Ethernet network occurs when two end nodes attempt to transmit data using the same data link at the same time. A collision causes the end nodes to stop sending data. This is called backpressure.

When a switch port needs to stop a half-duplex end node from transmitting data, it forces a collision on the data link, which stops the end node. Once the port is ready to receive data again, it stops forcing collisions.

The default setting for backpressure on a switch port is disabled.

Selecting this option displays the Back Pressure menu shown in Figure 26.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                     Back Pressure
Configuring Port 11
1 - Back Pressure .................. Disabled
2 - Back Pressure Cell Limit ...... 8192

R - Return to Previous Menu

Enter your selection?
```

Figure 26. Back Pressure Menu

The options on the Back Pressure menu are described below:

1 - Back Pressure
Enables and disables backpressure on a port. Possible values are:

Disabled The port will not use backpressure. This is the default setting.

Enabled The port will use backpressure.

2 - Back Pressure Cell Limit
Specifies the number of cells. A cell represents 64 bytes. The range is 1 to 57,344 cells. The default is 8192.

L - Rate Limit
For instructions, refer to “Setting the Rate Limit” on page 118.

U - Unknown Multicast Filtering
Discards all unknown ingress multicast packets on a port when activated. This feature can help improve switch performance in instances where a multicast stream arrives on a port of a switch where there are no host nodes. Without the feature, the unknown multicast packets would be flooded out all ports of the same VLAN as the ingress port. With the feature, the unknown multicast packets are discarded at the ingress port, reducing the number of packets that the switch must forward. The default setting is disabled.

A dynamic multicast address is added to the address table only after a host node responds with a request to join the group. Consequently, it is possible to activate this feature on a port and have it filter the packets even after the port has started receiving a multicast stream.

This feature does not filter multicast queries.

Possible settings are:

Enabled The port discards unknown ingress multicast packets.

Disabled The port forwards unknown ingress multicast packets. This is the default setting.

D - Set Default Port Configuration
Resets all port settings to the default values.

F - Force Renegotiation
If the port is already operating in Auto-Negotiation, this options prompts the port to Auto-Negotiate again with the end node. This can be helpful if you believe that a port and end node are not operating at the same speed and duplex mode. If the port’s speed and duplex mode have been set manually, this option returns the port to Auto-Negotiation.

X - Reset Port
Resets the speed and duplex mode of the selected port to the default value of Auto-Negotiation. Also returns the MDI/MDIX setting to the default value of Auto-Detect.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Setting the Rate Limit

This feature sets the maximum number of ingress packets the switch ports accept each second. Packets exceeding the threshold are discarded. You can enable the rate limiting threshold independently for multicast, broadcast, and unknown unicast packets. However, the same threshold applies to all packet types.

To configure this feature, you must enter a rate limit. This establishes the maximum number of packets the individual ports will accept per second. This limit applies to all ports and to all three packet types. There can be only one packet limit value for the switch.

Here is an example. Assume that you set a rate limit of 5,000 packets and you activate multicast and broadcast rate limiting. Each switch port will accept up to 5,000 ingress multicast packets and 5,000 ingress broadcast packets each second. If a port receives more of either type, it discards the extra packets. Since the feature was not activated for unknown unicast packets, ports do not restrict their number. (An unknown unicast packet is a packet with a MAC address not stored in the switch's MAC address table.)

To set rate limiting, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.

   The Port Configuration menu is shown in Figure 21 on page 106.

2. From the Port Configuration menu, type **1** to select Port Configuration.

   The following prompt is displayed:

   ```
   Enter port-list ->
   ```

3. Enter any port on the switch.

   This feature cannot be set on a per-port basis. You can enter any port or range of ports and the change will apply to all switch ports.

   The Port Configuration menu is shown in Figure 23 on page 109.

4. Type **L** to select Rate Limit.

The Rate Limiting menu is shown in Figure 27.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                      11:20:02 02-Jan-2006
                              Rate Limiting
Configuring Port 1
1 - Broadcast Rate Limiting Status ........... Disabled
2 - Multicast Rate Limiting Status ........... Disabled
3 - Unknown Unicast Rate Limiting Status ..... Disabled
4 - Rate Limit ............................... 262143 packets/second

R - Return to Previous Menu

Enter your selection?
```

Figure 27. Rate Limiting Menu

5. Type **4** to select Rate Limit and, when prompted, enter the maximum number of broadcast, multicast, and unknown unicast ingress packets you want all switch ports to accept each second. This threshold is applied independently to each packet type.

6. Type **1**, **2**, or **3** to activate the threshold for broadcast packets, multicast packets, and unknown unicast packets, respectively. You can enable this feature on one, two, or all three packet types.

   Rate limiting changes are immediately implemented on all switch ports.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying Port Statistics

To display Ethernet port statistics, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **3** to select Port Statistics.

   The Port Statistics menu is shown in Figure 28.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                     Port Statistics
1 - Display Port Statistics
3 - Clear Port Statistics

R - Return to Previous Menu

Enter your selection?
```

Figure 28. Port Statistics Menu

3. From the Port Statistics menu, type **1** to select Display Port Statistics.

   This prompt is displayed:

   ```
   Enter port-list:
   ```

4. Enter the port whose statistics you want to view. You can specify more than one port at a time.

   A menu is displayed containing the statistics for each port. The information in this menu is for viewing purposes only. The statistics are defined below:

   Bytes Received
   Number of bytes received on the port.

   Bytes Sent
   Number of bytes transmitted from the port.

   Frames Received
   Number of frames received on the port.

   Frames Sent
   Number of frames transmitted from the port.

   Broadcast Frames Received
   Number of broadcast frames received on the port.

Broadcast Frames Sent
Number of broadcast frames transmitted from the port.

Multicast Frames Received
Number of multicast frames received on the port.

Multicast Frames Sent
Number of multicast frames transmitted from the port.

Frames 64 Bytes
Frames 65 - 127 Bytes
Frames 128 - 255 Bytes
Frames 256 - 511 Bytes
Frames 512 - 1023 Bytes
Frames 1024 - 1518 Bytes
Number of frames transmitted from the port, grouped by size.

CRC Error
Number of frames with a cyclic redundancy check (CRC) error but with the proper length (64-1518 bytes) received on the port.

Jabber
Number of occurrences of corrupted data or useless signals appearing on the port.

No. of Rx Errors
Total number of frames received on the port containing errors.

No. of Tx Errors
Total number of frames transmitted on the port containing errors.

Undersize Frames
Number of frames that were less than the minimum length specified by IEEE 802.3 (64 bytes including the CRC) received on the port.

Oversize Frames
Number of frames exceeding the maximum specified by IEEE 802.3 (1518 bytes including the CRC) received on the port.

Fragments
Number of undersized frames, frames with alignment errors, and frames with frame check sequence (FCS) errors (CRC errors) received on the port.

Tx Collisions
Number of collisions that have occurred on the port. This applies only to ports operating in half duplex.

Clearing Port Counters

To return the statistics counters of a port to zero, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **3** to select Port Statistics.

   The Port Statistics menu is shown in Figure 28 on page 120.

1. From the Port Statistics menu, type **2** to select Clear Port Statistics.

   This prompt is displayed:

   ```
   Enter port-list:
   ```

2. Enter the port whose statistics counters you want to return to zero. You can specify more than one port at a time.

   The port counters are returned to zero.

Chapter 7

MAC Address Table

The chapter contains the procedures for viewing the static and dynamic MAC address table.

This chapter contains the following sections:

MAC Address Overview

The AT-8500 Series switch has a MAC address table with a storage capacity of 8,000 entries. The switch uses the table to store the MAC addresses of the network nodes connected to its ports, along with the port number on which each address was learned.

The switch learns the MAC addresses of the end nodes by examining the source address of each packet received on a port. It adds the address and port on which the packet was received to the MAC table if the address is not already in the table. The result is a table that contains the MAC addresses of all the devices connected to the switch's ports, and the port number where each address was learned.

When the switch receives a packet, it also examines the destination address and, by referring to its MAC address table, determines the port where the destination node is connected. It then forwards the packet to the appropriate port and on to the end node. This increases network bandwidth by limiting each frame to the appropriate port when the intended end node is located, freeing the other switch ports for receiving and transmitting other network packets.

If the switch receives a packet with a destination address that is not in the MAC address table, it floods the packet to all the ports on the switch, excluding the port where the packet was received. If the ports have been grouped into virtual LANs, the switch floods the packet only to those ports that belong to the same VLAN from where the packet originated. This prevents packets from being forwarded onto inappropriate LAN segments and increases network security. When the destination node responds, the switch adds its MAC address and port number to the table.

If the switch receives a packet with a destination address that is on the same port where the packet was received, it discards the packet without forwarding it on to any port. Since both the source node and the destination node for the packet are located on the same port on the switch, there is no reason for the switch to forward the packet. This too increases network performance by preventing frames from being forwarded unnecessarily to other network devices.

The type of MAC address described above is referred to as a *dynamic MAC address*. Dynamic MAC addresses are addresses that the switch learns by examining the source MAC addresses of the frames received on the ports.

Dynamic MAC addresses are 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 with that MAC address 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 period of time that the switch waits before purging an inactive dynamic MAC address is called the *aging time*. This value is adjustable on the AT-8500 Series switch. The default value is 300 seconds (5 minutes). For instructions on changing the aging timer, refer to "Changing the Aging Time" on page 134.

The MAC address table can also store *static MAC addresses*. A static MAC address 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.

You might need to enter static MAC addresses of end nodes the switch might not learn in its normal dynamic learning process, or if you want a MAC address to remain permanently in the table, even when the end node is inactive.

Displaying MAC Addresses

The management software has two menu selections for displaying the MAC addresses of a switch. One selection displays the static and dynamic unicast MAC addresses while the other displays the static and dynamic multicast addresses.

To display the MAC address tables, perform the following procedure:

1. From the Main Menu, type **4** to select MAC Address Tables. The MAC Address Tables menu is shown in Figure 29.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                   MAC Address Tables
1 - MAC Address Aging Time ......... 300 second(s)
2 - MAC Addresses Configuration
3 - Display Unicast MAC Addresses
4 - Display Multicast MAC Addresses

R - Return to Previous Menu

Enter your selection?
```

Figure 29. MAC Address Tables Menu

2. From the MAC Address Tables menu, type **3** to select Display Unicast MAC Addresses or **4** to select Display Multicast MAC Addresses. The Display Unicast MAC Addresses menu is shown in Figure 30. The Display Multicast MAC Addresses menu has the same selections.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
             Display Unicast MAC Addresses
1 - Display All
2 - Display Static
3 - Display Dynamic
4 - Display by Port
5 - Display Specified MAC
6 - Display by VLAN ID
7 - Display on Base Ports

R - Return to Previous Menu

Enter your selection?
```

Figure 30. Display Unicast MAC Addresses Menu

3. Select the desired option. The options are explained below:

   1 - Display All
   This selection displays all dynamic addresses learned on the ports of the switch and all static addresses that have been assigned to the ports. An example of a unicast MAC address table is shown in Figure 31.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                        11:20:02 02-Jan-2006
                                 Display All
                                    Page 1
Total Number of MAC Addresses: 121
MAC Address           Port      VlanID    Type
-------------------------------------------------------------------------
01:80:C1:00:02:01     0         0         Static (fixed, non-aging)
00:a0:d2:18:1a:c8     1         1         Dynamic
00:a0:c4:16:3b:80     2         1         Dynamic
00:a0:12:c2:10:c6     3         1         Dynamic
00:a0:c2:09:10:d8     4         1         Dynamic
00:a0:33:43:a1:87     5         1         Dynamic
00:a0:12:a7:14:68     6         1         Dynamic
00:a0:d2:22:15:10     7         1         Dynamic
00:a0:d4:18:a6:89     8         1         Dynamic

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 31. Display All Menu - Unicast MAC Addresses

Note
The first address in the unicast MAC address table is the address of the switch.

The information in this menu is for viewing purposes only. The columns in a unicast MAC address menu are defined below.

MAC - The static or dynamic unicast MAC address.

Port - The port where the address was learned or assigned. The MAC address with Port 0 is the address of the switch.

VlanID - The ID number of the VLAN where the port is an untagged member.

Type - The type of the address: static or dynamic.

An example of a multicast MAC address table is shown in Figure 32.

```
              Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                            Production Switch
User: Manager                                          11:20:02 02-Jan-2006
                                  Display All
                                    Page 1
Total Number of MCAST MAC Addresses: 1
MAC Address        VLAN ID Type         Port Maps (U:Untagged T:Tagged)
-------------------------------------------------------------------------
01:00:51:00:00:01  1       Static       U:1-4
                                        T:

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 32. Display All Menu - Multicast MAC Addresses

The information in this menu is for viewing purposes only. The columns in a multicast MAC address menu are defined below.

MAC Address- The static or dynamic multicast MAC address.

VlanID - The ID number of the VLAN where the port is an untagged member.

Type - The type of address: static or dynamic.

Port Maps - The tagged and untagged ports on the switch that are members of a multicast group. This column is useful in determining which ports belong to different groups.

The other options in the Display Unicast MAC Addresses menu or Display Multicast MAC Addresses menu are:

2 - Display Static
This selection displays just the static addresses assigned to the ports on the switch.

3 - Display Dynamic
This selection displays just the dynamic addresses learned on the ports on the switch.

4 - Display by Port
Displays the dynamic and static MAC addresses of a particular port. When you select this option, you are prompted for a port number. You can specify more than one port at a time.

5 - Display Specified MAC
Displays the port number on which a MAC address was assigned or learned.

In some situations, you might want to know on which port a particular MAC address was learned. You could display the MAC address table and scroll through the list looking for the MAC address. But if the switch is part of a large network, finding the address could prove difficult.

This menu option offers an easier way. You can specify the MAC address and let the management software automatically locate the port on the switch where the device is connected.

6 - Display by VLAN ID
Displays all the static and dynamic addresses learned on the tagged and untagged ports of a specific VLAN. When you select this option, you are prompted for the VLAN ID number of the VLAN. You can specify only one VLAN at a time

7 - Display on Base Ports
This displays the static and dynamic MAC addresses learned on the base ports. Base ports are the standard ports on the switch, excluding optional expansion modules, GBIC modules, or SFP modules.

Adding Static Unicast and Multicast MAC Addresses

This section contains the procedure for adding static unicast and multicast MAC addresses to the switch. You can assign up to 255 static addresses per port on an AT-8500 Series switch.

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

1. From the Main Menu, type **4** to select MAC Address Tables.

   The MAC Address Tables menu is shown in Figure 29 on page 126.

2. From the MAC Address Tables menu, type **2** to select MAC Addresses Configuration.

   The MAC Addresses Configuration menu is shown in Figure 33.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
             MAC Addresses Configuration
1 - Add Static MAC Address
2 - Delete MAC Address
3 - Delete All Dynamic MAC Addresses

R - Return to Previous Menu

Enter your selection?
```

Figure 33. Configure MAC Addresses Menu

3. From the Configure MAC Addresses menu, type **1** to select Add static MAC address.

   The following prompt is displayed:

   ```
   Please enter MAC address ->
   ```

4. Enter the static unicast or multicast MAC address in either of the following formats:

   ```
   XXXXXXXXXXXX or XXXXXX XXXXXX
   ```

5. Once you have specified the MAC address, the following prompt is displayed:

   ```
   Enter port-list: ->
   ```

6. Enter the number of the port on the switch where you want to assign the static address. If you are adding a static unicast address, you can specify only one port.

If you are entering a static multicast address, you must specify the port when the multicast application is located as well as the ports where the host nodes are connected. Assigning the address only to the port where the multicast application is located will prevent the forwarding of the multicast packets to the host nodes. You can specify the ports individually (e.g., 1,4,5), as a range (e.g., 11-14) or both (e.g., 15-17,22,24).

The following prompt is displayed:

```
Please enter VLAN ID: [1 to 4094] -> 1
```

7. Enter the VLAN ID where the port is a member.
8. Repeat this procedure starting with Step 3 to enter additional static unicast or multicast MAC addresses.
9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting Unicast and Multicast MAC Addresses

To delete a dynamic or static unicast or multicast address from the MAC address table, perform the following procedure:

1. From the Main Menu, type **4** to select MAC Address Tables.

   The MAC Address Tables menu is shown in Figure 29 on page 126.

2. From the MAC Address Tables menu, type **2** to select Configure MAC Addresses.

   The Configure MAC Addresses menu is shown in Figure 33 on page 130.

3. From the Configure MAC Addresses menu, type **2** to select Delete MAC Address.

   The following prompt is displayed:

   ```
   Please enter a MAC address ->
   ```

4. Enter the unicast or multicast MAC address to be deleted in either of the following formats:

   ```
   XXXXXXXXXXXX or XXXXXX XXXXXX
   ```

   After you have entered the MAC address, the following prompt is displayed:

   ```
   Please enter VLAN ID -> [1 to 4094] -> 1
   ```

5. Enter the VLAN ID of the port where the address was assigned or learned.

   The MAC address is deleted from the switch's MAC address table.

   Note
   You cannot delete a switch's MAC address, an STP BPDU MAC address, or a broadcast address.

6. Repeat the procedure to delete additional MAC addresses.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting All Dynamic MAC Addresses

To delete all dynamic unicast and multicast MAC address from the MAC address table, do the following:

1. From the Main Menu, type **4** to select MAC Address Tables.

   The MAC Address Tables menu is shown in Figure 29 on page 126.

2. From the MAC Address Tables menu, type **2** to select MAC Addresses Configuration.

   The MAC Addresses Configuration menu is shown in Figure 33 on page 130.

3. From the MAC Addresses Configuration menu, type **3** to select Delete All Dynamic MAC Addresses.

   The following prompt is displayed:

   ```
   All learned MAC (non-static) addresses will be deleted
   Do you want to continue? [Yes/No] ->
   ```

4. Enter **Y** to delete the addresses or **N** to cancel the procedure.

   If you respond with yes, all dynamic unicast and multicast addresses are deleted from the table, and the switch begins to learn new addresses.

Changing the Aging Time

The switch uses the aging time to delete inactive dynamic MAC addresses from the MAC address table to prevent the table from becoming full of addresses of inactive nodes. The switch deletes an address from the table when it detects no packets sent to or received from the address after the expiration of the time specified by the aging time.

The default setting for the aging time is 300 seconds (5 minutes).

To adjust the aging time, perform the following procedure:

1. From the Main Menu, type **4** to select MAC Address Tables.

   The MAC Address Tables menu is shown in Figure 29 on page 126.

2. From the MAC Address Tables menu, type **1** to select MAC Address Aging Time.

   The following prompt is displayed:

   ```
   Enter your new value -> [0 to 1048575]
   ```

3. Enter a new value in seconds.

   The range is 0 to 1048575 seconds. The default is 300 seconds (5 minutes). The value 0 (zero) disables the aging timer. When disabled, no dynamic addresses are deleted from the table, even addresses of inactive nodes, and the table stops adding new entries after it reaches maximum capacity.

   The new value is immediately activated on the switch.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Chapter 8

Static and LACP Port Trunks

This chapter contains the procedures for creating, modifying, and deleting static and LACP port trunks. Sections in the chapter include:

Port Trunk 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 the other network device and is useful in situations where a single physical link between the devices is insufficient to handle the traffic load.

The AT-8500 Series switch supports two types of port trunks:

- Static trunks
- Link Aggregate Control Protocol (LACP) IEEE 802.3ad trunks

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 simply designate the ports on the switch that are to be in the trunk and the management software on the switch automatically groups them together. The management software also gives you control over how the traffic is to be distributed over the trunk ports, as described in “Load Distribution Methods” on page 144.

The example in Figure 34 illustrates a static port trunk of four links between two AT-8524M switches.

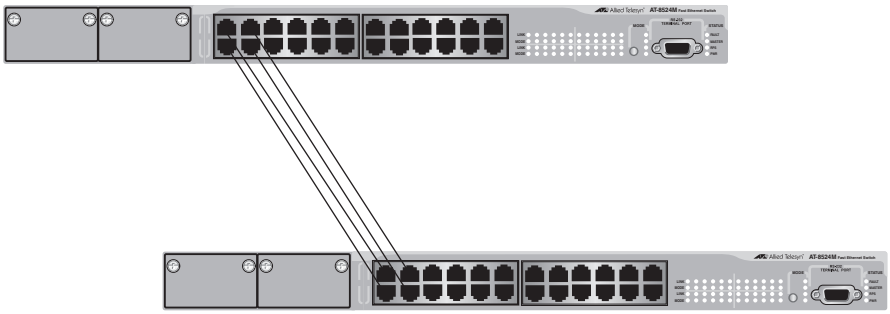

Figure 34. Static Port Trunk Example

Network equipment vendors tend to employ different techniques to implement static trunks. Consequently, a static trunk on one device might not be compatible 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 Telesyn layer 2 managed switch cannot form a static trunk with a device

from another manufacturer; but there is the possibility that the implementations of static trunking on the two devices might not be compatible.

It should also be noted that this type of 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 reduced. Though 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

Here are the guidelines to creating a static trunk:

- Allied Telesyn recommends using static port trunks between Allied Telesyn networking devices to ensure compatibility. While an Allied Telesyn device might be able to form a static trunk with a device from another equipment vendor, the implementations might not be fully compatible, resulting in undesirable 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 5-9) or nonconsecutive (for example, Ports 4, 8, 11, 20).
- Before creating a port trunk, examine the speed, duplex mode, flow control, and back pressure settings of the lowest number port that will be in the trunk. Verify that its settings are correct for the device to which the trunk will be connected. When you create a static port trunk, the management software copies the current settings of the lowest numbered port in the trunk to the other ports, because all ports in a static trunk must have the same settings. For example, if you create a port trunk consisting of ports 5 to 8, the parameter settings for port 5 are copied to ports 6, 7, and 8 so that all the ports of the trunk have the same settings.
- Once you have created a port trunk, do not change the speed, duplex mode, flow control or back pressure of any port in the trunk without making the same change to the other ports.
- A port can belong to only one static trunk at a time.
- A port cannot be a member of a static trunk and a LACP trunk at the same time.
- The switch can support up to six static and LACP 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 a static trunk must be untagged members of the same VLAN. A trunk cannot consist of untagged ports from different VLANs.
- ❒ The switch selects the lowest numbered port in the trunk to handle broadcast packets and packets of unknown destination. For example, a trunk of ports 11 to 15 would use port 11 for broadcast packets.
- ❒ You can create a port trunk of the ports in two expansion modules in an AT-8500 Series switch, providing that the ports are of the same medium type and have the same operating specifications.

LACP Trunk Overview

An LACP (Link Aggregation Control Protocol) trunk is another type of port trunk. Like a static trunk, it can increase the bandwidth between two 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 trunks tend to be vendor specific, the AT-8500 Series implementation of LACP is compliant with the IEEE 802.3ad standard. This makes it interoperable with equipment from other vendors that also comply with the standard, allowing you to create trunks between Allied Telesyn devices and networking 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. Should a link in a static trunk go down, the overall bandwidth of the trunk is reduced and restoring it requires reestablishing the link or manually modifying the trunk by adding another port to it. In contrast, an LACP trunk can activate ports in a stand-by mode when an active link fails to maintain the maximum possible bandwidth of the trunk.

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 two reserve ports to maintain maximum bandwidth of the trunk.

The main component of an LACP trunk is an *aggregator*. An aggregator 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 one time, the extra ports are placed in a stand-by 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 aggregator. 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 a switch is to support more than one aggregate trunk, it may be necessary to place each trunk in a separate aggregator, while in other cases you may be able to create just one aggregator and let the switch discern the individual aggregate trunks for you, automatically. The determining factor is whether the trunks are going to the same or different devices. If the trunks are going to the same device, you need to create a different aggregator for each trunk. If they are going to different devices, you can create just one aggregator and the switch can form the aggregate trunks itself.

Here are a two examples. Figure 35 illustrates an AT-8524M switch with two LACP trunks, each containing three links. Since both aggregate trunks go to the same 802.3ad-compliant device, in this case another Fast Ethernet switch, each trunk requires a separate aggregator.

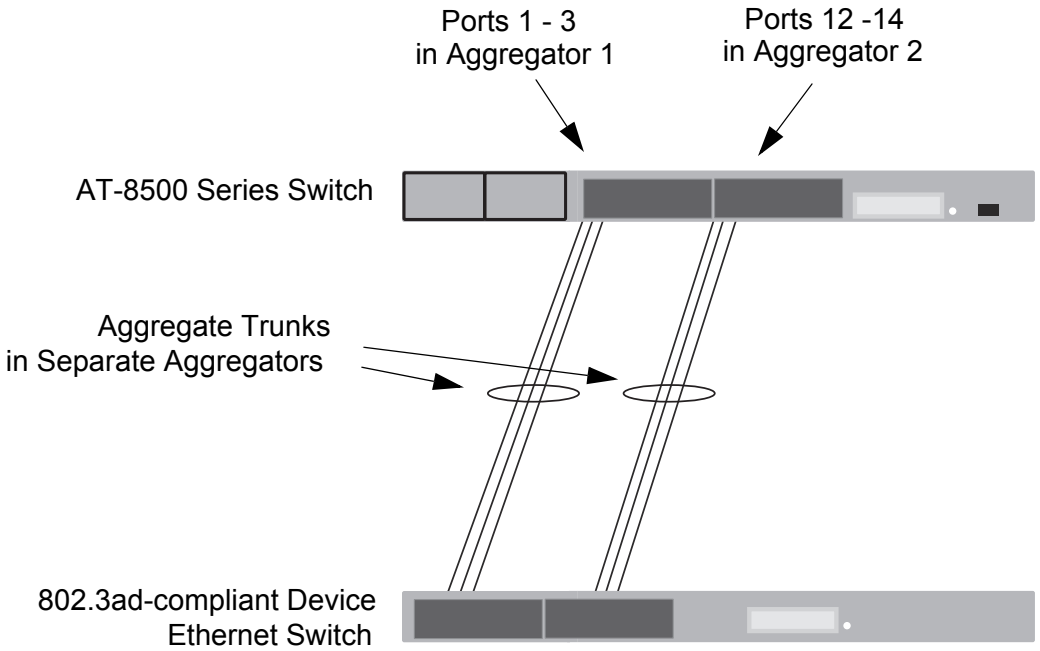

Figure 35. Example of Multiple Aggregators for Multiple Aggregate Trunks

Here is how the example might look in table format for the ports on the AT-8500 Series switch.

Aggregator Description	Aggregator Ports	Aggregate Trunk Ports
Aggregator 1	1-3	1-3
Aggregator 2	12-14	12-14

Caution
The example cited here illustrates a loop in a network. Network loops should be avoided to prevent broadcast storms.

If the aggregate trunks go to different devices, you can create one aggregator and let the AT-8500 Series switch form the trunks for you automatically. This is illustrated in Figure 36. The ports of the two aggregate trunks on the AT-8500 Series switch are members of the same aggregator. It is the switch that determines that there are actually two separate aggregate trunks.

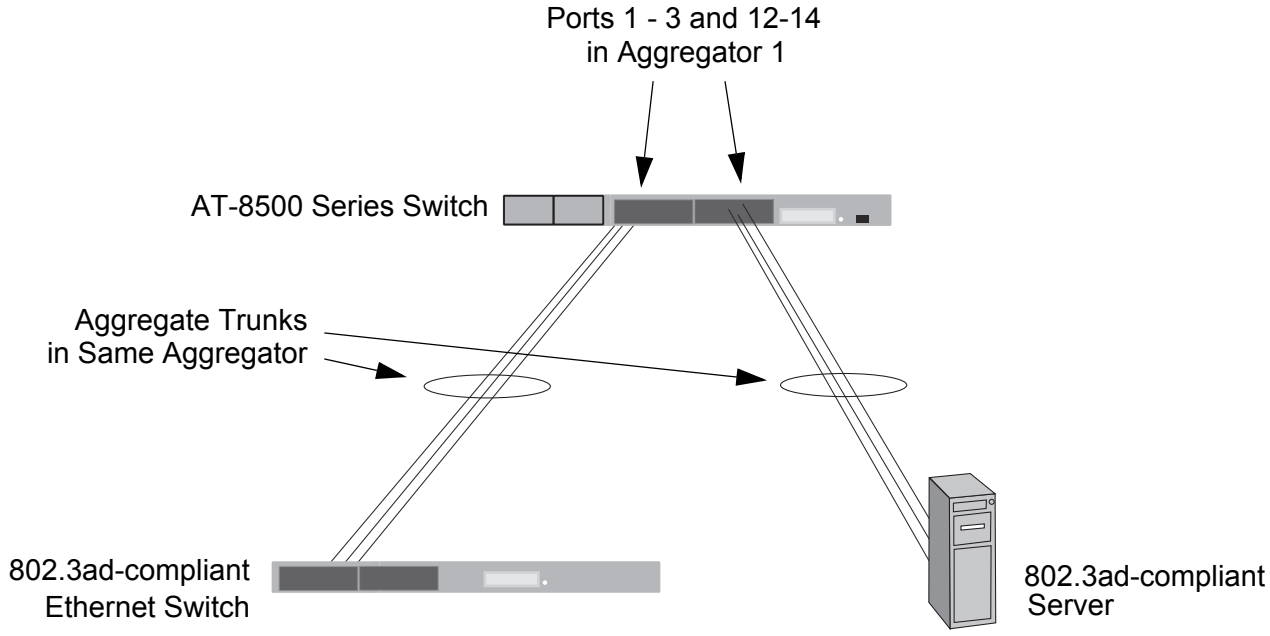

Figure 36. Example of an Aggregator with Multiple Trunks

Here is how this example looks in table format for the ports on the AT-8500 Series switch.

Aggregator Description	Aggregator Ports	Aggregate Trunk Ports
Aggregator 1	1-3, 12-14	1-3
		12-14

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 can save you 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 forming a 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 to be active and which are to be in standby.

If a conflict occurs, the devices need a mechanism for resolving the problem, a means by which they can decide whose LACP settings are to take precedence. That is the function of the system LACP priority value. A hexadecimal value of from 1 to FFFF, this parameter is used whenever the devices encounter a conflict creating a trunk. The lower the number, the higher the priority. The settings on the device with the higher priority takes 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.

This parameter can prove useful when connecting an aggregate trunk between an AT-8500 Series switch and another 802.3ad-compliant device that does not have the same LACP trunking capabilities. If the other device's capability is less than that of the AT-8500 Series', you should give that device the higher priority so that its settings are used by both devices when forming the trunk.

For example, an aggregate trunk of six links between an AT-8500 Series switch and an 802.3ad-compliant device that supported up to four active links at one time could possibly result in a conflict. The AT-8500 Series switch would try to use all six links as active, since it can handle up to eight active links in a trunk at one time, while the other device would want to use only four ports as active. By giving the other 802.3ad device the higher priority, the conflict would be avoided. The AT-8500 Series switch would use only four active links because the settings on the other device would take precedence. The other ports would be in standby mode.

Adminkey Parameter

The adminkey is a hexadecimal value from 1 to FFFF that identifies an aggregator. Each aggregator on a switch must have a unique adminkey. The adminkey is limited to a switch. Two aggregators on different switches can have the same adminkey without creating a conflict.

LACP Port Priority Parameter

The switch uses a port's LACP priority to determine which ports are to be 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 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.

A port's priority value is not adjustable.

The selection of the active links in an aggregate trunk is dynamic. It 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 it to return to the active state by virtue of its having a higher priority, while the port that replaced it is returned to the standby mode.

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

Load Distribution Methods

The load distribution method determines the manner in which the switch distributes the traffic across the active ports of an aggregate trunk. The method is assigned to an aggregator and applies to all aggregate trunks within it. If you want to assign different load distribution methods to different aggregate trunks, you must create a separate aggregator for

each trunk. For further information, refer to "Load Distribution Methods" on page 144.

LACP Trunk Guidelines

Here are the guidelines to follow when 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-8500 Series switch supports up to eight active ports in an aggregate trunk at a time.
- The switch supports a maximum of three aggregate trunks.
- The ports of an aggregate 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 consecutive (for example Ports 5-9) or nonconsecutive (for example, Ports 4, 8, 11, 20).
- 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. (The switch's management software does not display an error message if you create an aggregator with ports from different untagged VLANs. However, the ports are not added to the aggregate trunk when the trunk is established.)
- The switch can support up to six static and LACP 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.
- 10/100Base-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 expansion modules or GBIC modules with 1000Base-X fiber optic ports.
- Only those ports that are members of an aggregator transmit LACPDU packets.
- The load distribution method is applied at the aggregator level. If you want aggregate trunks to have different load distribution methods, you must create a separate aggregator for each trunk. For further information, refer to "Load Distribution Methods" on page 144.
- A port that is a member 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 AT-8500 Series switch 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 eight, the maximum number for the AT-8500 Series switch, you should probably assign it a higher system LACP priority than the AT-8500 Series switch. If it is more than eight, assign the AT-8500 Series switch the higher priority. This can help avoid a possible conflict between the devices if some ports are placed in the standby mode when the trunk is created by the devices. For background information, refer to "LACP System Priority" on page 141.
- LACPDU packets are transmitted as untagged packets.

Load Distribution Methods

This section discusses the load distribution methods and applies to both static and LACP port trunks.

One of the steps to creating a static or LACP port trunk is the selection of a load distribution method. This step determines how the switch distributes the traffic load across the ports in the trunk. The AT-S62 management software offers the following load distribution methods:

- Source MAC Address (Layer 2)
- Destination MAC Address (Layer 2)
- Source MAC Address / Destination MAC Address (Layer 2)
- Source IP Address (Layer 3)
- Destination IP Address (Layer 3)
- Source IP Address / Destination IP Address (Layer 3)

The load distribution methods examine the last three bits of a packet's MAC or IP address and compare the bits against mappings assigned to the ports in the trunk. The port mapped to the matching bits is selected as the transmission port for the packet.

In cases where you select a load distribution that employs either a source or destination address but not both, the last three bits of only the designated address are used in the selection of a transmission port in a trunk. If you select one of the two load distribution methods that employs both source and destination addresses, port selection is achieved through an XOR operation of the last three bits of both addresses.

As an example, assume you created a static or LACP aggregate trunk of Ports 7 to 14 on a switch. The table below shows the mappings of the switch ports to the possible values of the last three bits of a MAC or IP address.

Last 3 Bits	000 (0)	001 (1)	010 (2)	011 (3)	100 (4)	101 (5)	110 (6)	111 (7)
Trunk Ports	7	8	9	10	11	12	13	14

Now assume you selected source MAC address as the load distribution method and that the switch needed to transmit over the trunk a packet with a source MAC address that ended in 9. The binary equivalent of 9 is 1001, making the last three bits of the address 001. An examination of the table above indicates that the switch would use Port 8 to transmit the frame because that port is mapped to the matching bits.

The same method is used for the two load distribution methods that employ both the source and destination addresses. Only here the last three bits of both addresses are combined by an XOR process to derive a single value which is then compared against the mappings of bits to ports. The XOR rules are as follows:

0 XOR 0 = 0

0 XOR 1 = 1

1 XOR 0 = 1

1 XOR 1 = 0

As an example, assume that you had selected source and destination MAC addresses for the load distribution method in our previous example, and that a packet for transmission over the trunk had a source MAC address that ended in 9 and a destination address that ended in 3.

The binary values would be:

9 = 1001

3 = 0011

Applying the XOR rules above on the last three bits would result in 010. A examination of the table above shows that the packet would be transmitted from port 9.

Port trunk mappings on an AT-8500 Series switch can consist of up to eight ports. This corresponds to the maximum number of ports allowed in a static trunk and the maximum number of active ports in an LACP trunk. (Inactive ports in an LACP trunk are not applied to the mappings until they transition to the active status.)

You can assign different load distribution methods to different static trunks on the same switch. The same is true for LACP aggregators. However, it should be noted that all aggregate trunks within an LACP aggregator must use the same load distribution method.

The load distribution methods assume that the final three bits of the source and/or destination addresses of the packets from the network nodes are varied enough to support adequate distribution of the packets over the trunk ports. A lack of variation can result in one or more ports in a trunk being used more than others, with the potential loss of a trunk's efficiency and performance.

Managing Static Port Trunks

The following procedures explain how to create, modify, and delete static port trunks:

- "Creating a Static Port Trunk" on page 147
- "Modifying a Static Port Trunk" on page 150
- "Deleting a Static Port Trunk" on page 152

For background information, refer to "Static Port Trunk Overview" on page 136.

Creating a Static Port Trunk

This section contains the procedure for creating a static port trunk on a switch. Be sure to review the guidelines in "Port Trunk Overview" on page 136 before performing the procedure.

Caution

Do not connect the cables to the trunk ports on the switches until after you have configured the trunk with the management software. Connecting the cables before configuring the software will create a loop in your network topology. Data loops can result in broadcast storms and poor network performance.

Note

Before creating a port trunk, examine the speed, duplex mode, and flow control settings of the lowest numbered port that will be a part of the trunk. Check to be sure that the settings are correct for the end node to which the trunk will be connected. When you create the trunk, the AT-S62 management software copies the settings of the lowest numbered port in the trunk to the other ports so that all the settings are the same.

You should also check to be sure that the ports are untagged members of the same VLAN. You cannot create a trunk of ports that are untagged members of different VLANs.

To create a port trunk, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **4** to select Port Trunking and LACP.

The Port Trunking and LACP menu is shown in Figure 37.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                Port Trunking and LACP
1 - Static Port Trunking
2 - LACP Configuration

R - Return to Previous Menu

Enter your selection?
```

Figure 37. Port Trunking and LACP Menu

3. From the Port Trunking and LACP menu, type **1** to select Static Port Trunking.

   The Static Port Trunking menu is shown in Figure 38.

```
       Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                      Production Switch
User: Manager                                11:20:02 02-Jan-2006
                     Static Port Trunking
ID   Name        Ports          Method        Status
-------------------------------------------------------

C - Create Trunk
D - Delete Trunk
M - Modify Trunk

R - Return to Previous Menu

Enter your selection?
```

Figure 38. Static Port Trunking Menu

This menu lists the trunks that already exist on the switch. The information includes the following:

- ID - The ID number of the static port trunk.
- Name - The name of the static port trunk.
- Ports - The ports of the static port trunk.
- Method - One of the following load distribution methods:

  SRC MAC — Source MAC address.

  DST MAC — Destination MAC address.

  SRC/DST MAC — Source address/destination MAC address.

  SRC IP — Source IP address.

DST IP Destination IP address.

SRC/DST IP Source address/destination IP address.

- Status - The operational status of the trunk. Up means at least one port in the trunk has established a link with a port on the other device. Down means no ports in the trunk have established a link with the other device.

4. Type **C** to select Create Trunk.

   The Create Trunk menu is shown in Figure 39.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                      Create Trunk

1 - Trunk ID ......... 1
2 - Trunk Name .......
3 - Trunk Method ..... SRC/DST MAC
4 - Trunk Ports ......

C - Create Trunk
R - Return to Previous Menu

Enter your selection?
```

Figure 39. Create Trunk Menu

5. Type **1** to select Trunk ID and, when prompted, enter an ID number for the trunk of from 1 to 6. A trunk must be assigned a unique ID number. The default value is the next unused ID number.

6. Type **2** to select Trunk Name and, when prompted, enter a name for the trunk. The name can be up to sixteen alphanumeric characters. No spaces or special characters, such as asterisks and exclamation points, are allowed. Each trunk must have a unique name.

7. To set the load distribution method, type **3** to toggle the selection through the following possible settings:

   - SRC MAC - Source MAC address
   - DST MAC - Destination MAC address
   - SRC/DST MAC - Source address /destination MAC address
   - SRC IP - Source IP address trunking
   - DST IP - Destination IP address trunking
   - SRC/DST IP - Source address /destination IP address

   The default is SRC/DST MAC. For background information, refer to "Load Distribution Methods" on page 144.

8. Type **4** to select Trunk Ports and, when prompted, enter the ports of the trunk. A trunk can contain up to eight ports. You can identify the ports individually (for example, 3,7,10), as a range (for example, 5-11), or both (for example, 2,4,11-14).

9. Type **C** to select Create Trunk.

   The port trunk is now active on the switch.

10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

11. Configure the ports on the remote switch for port trunking.

12. Connect the cables to the ports of the trunk on the switch.

    The port trunk is ready for network operations.

Modifying a Static Port Trunk

This section contains the procedure for modifying a static port trunk on the switch. Be sure to review the guidelines in “Static Port Trunk Guidelines” on page 137 before performing the procedure.

Caution
If you will be adding or removing ports from the trunk, you should disconnect all network cables from the ports of the trunk on the switch before performing the procedure. Adding or removing ports from a static port trunk without first disconnecting the cables may result in loops in your network topology, which can result in broadcast storms and poor network performance.

Note the following before performing this procedure:

- If you are adding a port and the port will be the lowest numbered port in the trunk, its parameter settings will overwrite the settings of the existing ports in the trunk. Consequently, you should check to see if its settings are appropriate prior to adding it.
- If you are adding a port and the port will not be the lowest numbered port in the trunk, its settings will be changed to match the settings of the existing ports in the trunk.
- If you are adding a port to a static trunk, you should check to be sure that the new port is an untagged member of the same VLAN as the other trunk ports. A trunk cannot contain ports that are untagged members of different VLANs.

To modify a port trunk, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.

2. From the Port Configuration menu, type **4** to select Port Trunking and LACP.

   The Port Trunking and LACP menu is shown in Figure 37 on page 148.

3. From the Port Trunking and LACP menu, type **1** to select Static Port Trunking.

   The Static Port Trunking menu is shown in Figure 38 on page 148.

4. Type **M** to select Modify Trunk.

   The following prompt is displayed:

   ```
   Enter Trunk ID: [1 to 6] ->
   ```

5. Enter the ID number of the trunk you want to modify.

   The Modify Trunk menu is displayed. The menu displays the operating specifications of the selected trunk. An example is shown in Figure 40.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                       Modify Trunk

1 - Trunk ID ......... 2
2 - Trunk Name ....... Server11
3 - Trunk Method ..... SRC/DST MAC
4 - Trunk Ports ...... 12-16

M - Modify Trunk
R - Return to Previous Menu

Enter your selection?
```

Figure 40. Modify Trunk Menu

Note
You cannot change a trunk's ID number.

6. To modify a port trunk's name, type **2** to select Trunk Name and, when prompted, enter the new name for the trunk. The name can be up to sixteen alphanumeric characters. No spaces or special characters, such as asterisks and exclamation points, are allowed. Each trunk must have a unique name.

7. To change the trunk's load distribution method, type **3** to toggle the selection through the following possible settings.

   - SRC MAC - Source MAC address
   - DST MAC - Destination MAC address

- ❒ SRC/DST MAC - Source address /destination MAC address
- ❒ SRC IP - Source IP address trunking
- ❒ DST IP - Destination IP address trunking
- ❒ SRC/DST IP - Source address /destination IP address

For background information on these selections, refer to “Load Distribution Methods” on page 144.

8. To change the ports of a trunk, type **4** to select Trunk Ports and, when prompted, enter the new ports of the trunk. A trunk can contain up to eight ports. You can identify the ports individually (for example, 3,7,10), as a range (for example, 5-11), or both (for example, 2,4,11-14). The new list of ports replaces the existing ports of the trunk.
9. Type **M** to select Modify Trunk.

   The modifications to the port trunk are activated on the switch.
10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.
11. Reconnect the cables to the ports of the trunk on the switch.

    The modified port trunk is ready for network operations.

Deleting a Static Port Trunk

To delete a static port trunk from the switch, perform the following procedure:

> ⚠ **Caution**
> Disconnect the cables from the port trunk on the switch before performing the following procedure. Deleting a port trunk without first disconnecting the cables can create loops in your network topology. Data loops can result in broadcast storms and poor network performance.

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Menu, type **4** to select Port Trunking and LACP.

   The Port Trunking and LACP menu is shown in Figure 38 on page 148.
3. From the Port Trunking and LACP menu, type **1** to select Static Port Trunking.

   The Static Port Trunking menu is shown in Figure 38 on page 148.
4. Type **D** to select Delete Trunk.

The following prompt is displayed:

```
Enter Trunk ID: [1 to 6] ->
```

5. Enter the ID number of the trunk to be deleted.

   A confirmation prompt is displayed.

6. Type **Y** for yes to delete the port trunk or **N** for no to cancel this procedure.

   The port trunk is deleted from the switch.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Managing LACP Trunks

The following procedures explain how to create and manage LACP trunks:

For background information, refer to “LACP Trunk Overview” on page 138.

Enabling or Disabling LACP

This procedure explains how to enable or disable LACP on the switch. When you enable LACP, the switch begins to transmit LACPDU packets from ports assigned to aggregators. If ports in an aggregator receive LACPDU packets from a remote device, the switch creates aggregate trunks. If no aggregators are defined, no LACPDU packets are transmitted. When you disable LACP, any ports in existing aggregators stop sending LACPDU packets and function as regular Fast Ethernet ports.

Caution
Do not disable LACP if there are defined aggregators. without first disconnecting all cables connected to the aggregate trunk ports. Otherwise, a network loop might occur, resulting in a broadcast storm and poor network performance.

To enable or disable LACP, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **4** to select Port Trunking and LACP.

   The Port Trunking and LACP menu is shown in Figure 38 on page 148.
3. Type **2** to select LACP Configuration.

The LACP (IEEE 8023ad) Configuration menu is shown in Figure 41.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                        11:20:02 02-Jan-2006
            LACP (IEEE 802.3ad) Configuration
1 - LACP Status ....................... Disabled
2 - Priority .......................... 0x0080
3 - Create Aggregator
4 - Modify Aggregator
5 - Configure Port
6 - Delete Aggregator
7 - Show LACP Port Status
8 - Show LACP Aggregator Status

R - Return to Previous Menu

Enter your selection?
```

Figure 41. LACP (IEEE 8023ad) Configuration Menu

4. Type **1** to toggle LACP Status between Disabled and Enabled. The default is disabled.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Setting a LACP System Priority

This procedure explains how to set the LACP system priority value on a switch. The switch uses this parameter if a conflict occurs when establishing an aggregate trunk with the other device. The LACP settings on the device with the higher priority take precedence over the settings on the other device. The lower the value, the higher the priority. A switch can have only one LACP system priority. For more information, refer to "LACP System Priority" on page 141.

To set the LACP system priority for the switch, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.

2. From the Port Configuration menu, type **4** to select Port Trunking and LACP.

   The Port Trunking and LACP menu is shown in Figure 38 on page 148.

3. Type **2** to select LACP Configuration.

   The LACP (IEEE 8023ad) Configuration menu is shown in Figure 41 on page 155.

4. Type **2** to s elect Priority.

The following prompt is displayed:

```
Enter Priority [0x1 - 0xFFFF]: [0x1 to 0xffff] -> 0x
```

5. Enter the new value is hexadecimal. The range is 1 to FFFF. The lower the value, the higher the priority. The prefix “0x” indicates that the number is hexadecimal.

   The new priority value takes effect immediately on the switch.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Creating an Aggregator

To create an aggregator, perform the following procedure:

Caution
Do not connect the cables to the ports of the aggregator on the switch until after you have configured the aggregator with the management software and enabled LACP. Connecting the cables before configuring the software and activating the protocol will create a loop in your network topology. Data loops can result in broadcast storms and poor network performance.

Note
Before creating an aggregator, verify that the ports that will be members of the aggregator are set to Auto-Negotiation or 100 Mbps, full-duplex. Aggregate trunks do not support half-duplex mode.

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **4** to select Port Trunking and LACP.

   The Port Trunking and LACP menu is shown in Figure 38 on page 148.

3. Type **2** to select LACP Configuration.

   The LACP (IEEE 8023ad) Configuration menu is shown in Figure 41 on page 155.

4. Type **3** to select Create Aggregator.

The Create LACP (IEEE 8023ad) Aggregator menu is shown in Figure 41 on page 155.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
         Create LACP (IEEE 802.3ad) Aggregator
1 - Aggregator ..................
2 - Adminkey ..................... 0x0000
3 - Distribution Mode ............ SRC/DST MAC
4 - Port Range ..................
C - Create Aggregator

R - Return to Previous Menu

Enter your selection?
```

Figure 42. Create LACP (IEEE 8023ad) Aggregator Menu

5. Configure the parameters as necessary. The parameters are defined here:

   1 - Aggregator
   Specifies a name for the aggregator. The name can be up to twenty alphanumeric characters. Spaces are allowed, but special characters, such as asterisks and exclamation points, are not. Each aggregator must have a unique name.

   2 - Adminkey
   Specifies a unique adminkey value for the aggregator. The value is entered in hexadecimal. The range is 1 to FFFF. For background information, refer to "Adminkey Parameter" on page 142.

   3 - Distribution Mode
   Sets the load distribution method. Possible settings are:

   - SRC MAC - Source MAC address
   - DST MAC - Destination MAC address
   - SRC/DST MAC - Source address /destination MAC address
   - SRC IP - Source IP address trunking
   - DST IP - Destination IP address trunking
   - SRC/DST IP - Source address /destination IP address

   The default is SRC/DST MAC. For background information, refer to "Load Distribution Methods" on page 144.

   4 - Port Range
   Specifies the aggregator ports. An aggregator can contain any number of ports on the switch. You can identify the ports individually (for example, 3,7,10), as a range (for example, 5-11), or both (for example, 2,4,11-14).

6. After you configure the parameters, type **C** to select Create Aggregator.

   The aggregator is created on the switch.

7. If LACP is not enabled on the switch, perform the procedure “Enabling or Disabling LACP” on page 154 and activate the protocol.

8. Configure LACP on the other network device.

9. Connect the cables to the ports of the aggregator on both the switch and the other network device.

   The aggregator and its aggregate trunk(s) are now ready for network operations.

> **Caution**
> Do not connect the cables to the ports of the aggregator on the switch until after you have enabled LACP. Connecting the cables before activating the protocol will create a loop in your network topology. Data loops can result in broadcast storms and poor network performance.

10. Repeat this procedure to create additional aggregators, if needed.

11. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an Aggregator

This procedure explains how to modify an aggregator. You can change an aggregator’s name, adminkey, or load distribution method. You can also use this procedure to add or remove ports. To modify an aggregator, you need to know its name or adminkey key. It is recommended that you review the section “LACP Trunk Guidelines” on page 143 before modifying an aggregator.

> **Caution**
> If you will be adding or removing ports from the aggregator, you should disconnect all network cables from the ports of the aggregator on the switch before performing the procedure. Adding or removing ports without first disconnecting the cables can result in loops in your network topology, which can result in broadcast storms and poor network performance.

To modify an aggregator, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.

2. From the Port Configuration menu, type **4** to select Port Trunking and LACP.

The Port Trunking and LACP menu is shown in Figure 38 on page 148.

3. Type **2** to select LACP Configuration.

   The LACP (IEEE 8023ad) Configuration menu is shown in Figure 41 on page 155.

4. Type **4** to select Modify Aggregator.

   The Modify LACP (IEEE 8023ad) Aggregator menu is shown in Figure 43.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
         Modify LACP (IEEE 802.3ad) Aggregator
1 - Aggregator ..................
2 - Adminkey ..................... 0x0000
3 - Distribution Mode ........... SRC/DST MAC
4 - Port Range ..................
M - Modify Aggregator

R - Return to Previous Menu

Enter your selection?
```

Figure 43. Modify LACP (IEEE 8023ad) Aggregator Menu

5. Type **1** to select Aggregator or **2** for Adminkey and, when prompted, enter the name or adminkey of the aggregator you want to modify. You can specify the aggregator by its name or adminkey number. The name is case-sensitive.

   After you enter the aggregator's name or adminkey, the specifications of the aggregator are displayed in the menu.

6. Adjust the settings as necessary. The parameters are defined here:

   1 - Aggregator
   Specifies a name for the aggregator. The name can be up to twenty alphanumeric characters. Spaces are allowed, but special characters, such as asterisks and exclamation points, are not. Each aggregator must have a unique name.

   2 - Adminkey
   Specifies a unique adminkey value for the aggregator. The value is entered in hexadecimal. The range is 1 to FFFF. For background information, refer to "Adminkey Parameter" on page 142.

3 - Distribution Mode
Sets the load distribution method. Possible settings are:

- ❒ SRC MAC - Source MAC address
- ❒ DST MAC - Destination MAC address
- ❒ SRC/DST MAC - Source address /destination MAC address
- ❒ SRC IP - Source IP address trunking
- ❒ DST IP - Destination IP address trunking
- ❒ SRC/DST IP - Source address /destination IP address

The default is SRC/DST MAC. For background information, refer to “Load Distribution Methods” on page 144.

4 - Port Range
Specifies the aggregator ports. An aggregator can contain any number of ports on the switch. You can identify the ports individually (for example, 3,7,10), as a range (for example, 5-11), or both (for example, 2,4,11-14).

7. After configuring the parameters, type **M** to select Modify Aggregator.

   The aggregator is modified on the switch.

8. Reconnect the cables to the ports of the aggregator.

   The modified aggregator is now ready for network operations.

Deleting an Aggregator

This procedure deletes an aggregator from the switch. The ports that are members of the aggregator stop transmitting LACPDU packets after the aggregator is deleted.

Caution
Disconnect the cables from the ports of the aggregator before performing the following procedure. Deleting an aggregator without first disconnecting the cables can create loops in your network topology. Data loops can result in broadcast storms and poor network performance.

To delete an aggregator, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **4** to select Port Trunking and LACP.

   The Port Trunking and LACP menu is shown in Figure 38 on page 148.

3. Type **2** to select LACP Configuration.

The LACP (IEEE 8023ad) Configuration menu is shown in Figure 41 on page 155.

4. Type **6** to select Delete Aggregator.

   The following prompt is displayed:

   ```
   Enter Aggregator Name [Max up to 20 alphanumeric
   characters]:
   ```

5. Enter the name of the aggregator you want to delete. The name is case-sensitive. You can delete only one aggregator at a time.

   A confirmation prompt is displayed.

6. Type **Y** to delete the aggregator or **N** to cancel the procedure.

   If you entered Yes, the aggregator is deleted.

Displaying LACP Port or Aggregator Status

To display LACP port or aggregator status, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **4** to select Port Trunking and LACP.

   The Port Trunking and LACP menu is shown in Figure 38 on page 148.

3. Type **2** to select LACP Configuration.

   The LACP (IEEE 8023ad) Configuration menu is shown in Figure 41 on page 155.

4. To view port status, type **7** to select Show LACP Port Status. To view aggregator status, type **8** to select Show LACP Aggregator Status.

Figure 44 is an example of the LACP (IEEE 802.3ad Port Status menu. The information in this window is for viewing purposes only. For definitions, refer to the IEEE 802.3ad standard.

```
                    LACP (IEEE 802.3ad) Port Status

Port ............. 01
Aggregator ....... Sales server

ACTOR                                PARTNER
=====================================++++++++============================
Actor Port ............. 06          Partner Port ......... 00
Selected ............... SELECTED    Partner System ....... 00-30-84-00-00-02
Oper Key ............... 0x0050      Oper Key  ............ 0x0004
Oper Port Priority  .... 0x0006      Oper Port Priority ... 0x0007
Individual ............. NO          Individual ........... NO
Synchronized............ YES         Synchronized.......... YES
Collecting  ............ YES         Collecting ........... YES
Distributing ........... YES         Distributing ......... NO
Defaulted .............. NO          Defaulted ............ NO
Expired ................ NO          Expired .............. NO
Actor Churn   .......... YES         Partner Churn ........ YES

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 44. LACP (IEEE 802.3ad Port Status Menu

Figure 45 is an example of the LACP (IEEE 802.3ad) Aggregator Status menu. The information is for viewing purposes only. An aggregator appears in the menu only if there is at least one active aggregate trunk between the switch and another network device.

```
             Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                        11:20:02 02-Jan-2006
                    LACP (IEEE 802.3ad) Aggregator Status

Aggregator #1 ................. Sales server
Adminkey ...................... 0x0050
Oper Key....................... 0x1405
Speed ......................... 100 Mbps
Ports in LAGID ................ 1-4
Aggregated Port ............... 1-4
Mode .......................... SRC/DST MAC

LAG ID:
[(0080,00-30-84-00-00-00,0041,00,0000),(0080,00-30-84-00-00-
02,0004,00,0000)]

R - Return to Previous Menu
Enter your selection?
```

Figure 45. LACP (IEEE 802.3ad) Aggregator Status Menu

If there are no active aggregate trunks on the switch, the following message is displayed:

```
No Aggregator with aggregatable Ports
```

Chapter 9

Port Mirroring

This chapter contains the procedures for creating and deleting a port mirror. Sections in the chapter include:

Port Mirroring Overview

The port mirroring feature is used to unobtrusively monitor the ingress and egress traffic on one or more ports by copying the traffic to another port. By connecting a network analyzer to the port where the traffic is being copied, you can monitor the traffic on the other ports without impacting network performance or speed.

The port(s) whose traffic is to be mirrored is called the *source port(s)*. The port where the traffic is copied to is called the *destination port*.

Observe the following guidelines when creating a port mirror:

- You can select only one destination port.
- You can select more than one source port. However, the more ports mirrored, the greater the likelihood the destination port will be 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 it will not provide an accurate mirror of the traffic of the six source ports.
- You can mirror either the ingress or egress traffic of the source ports, or both.
- The source and destination ports must be located on the same switch.
- The source ports and the destination port on an AT-8550GB Series switch must be located within the same port group. The port groups are:
  - Ports 1 to 24 and 49
  - Ports 25 to 48 and 50
- To create a mirror port for the Denial of Service defenses, specify only the destination port for the mirrored traffic. The management software automatically determines the source ports.

Creating a Port Mirror

To create a port mirror, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.

   The Port Configuration menu is shown in Figure 21 on page 106.

2. From the Port Configuration menu, type **6** to select Port Mirroring.

   The Port Mirroring menu is shown in Figure 46.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                     Port Mirroring
1 - Enable/Disable ..................... Disabled

R - Return to Previous Menu

Enter your selection?
```

Figure 46. Port Mirroring Menu #1

3. Type **1** to select Enable/Disable.

   The following prompt is displayed.

   ```
   Enter Enable(E)/Disable(D):
   ```

4. Type **E** to enable the feature.

   New options are added to the Port Mirroring menu, as shown in Figure 47.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                     Port Mirroring
1 - Enable/Disable ....................... Enabled
2 - Mirror-To (Destination) Port ........ None
3 - Ingress (Rx) Mirror (Source) Ports .. None
4 - Egress (Tx) Mirror (Source) Ports ... None

R - Return to Previous Menu

Enter your selection?
```

Figure 47. Port Mirroring Menu #2

5. Type **2** to select Mirror-To Port and, when prompted, enter the number of the port to function as the destination port. This is the port where the traffic from the source ports will be copied to and where the network analyzer will be located. You can specify only one destination port.

6. If you want to mirror the ingress (received) traffic on one or more ports, type **3** to select Ingress Mirror Port and, when prompted, enter the ports. You can identify the ports individually (for example, 3,7,10), as a range (for example, 5-11), or both (for example, 2,4,11-14). Entering “none” removes all ingress source ports.

7. If you want to mirror the egress (transmitted) traffic from one or more ports, type **4** to select Egress Mirror Port and, when prompted, enter the ports. Entering “none” removes all egress source ports.

   To monitor both the ingress and egress traffic of the source ports, you must specify the ports in both menu options 3 and 4.

   The port mirror is now functional. Attach a network analyzer to the destination port to monitor the traffic on the source ports.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Disabling a Port Mirror

To disable a port mirror, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.

   The Port Configuration menu is shown in Figure 21 on page 106.

2. From the Port Configuration menu, type **6** to select Port Mirroring.

   The Port Mirroring menu is shown in Figure 47 on page 167.

3. Type **1** to select Enable/Disable.

   The following prompt is displayed.

   ```
   Enter Enable(E)/Disable(D):
   ```

4. Type **D** to disable the feature.

   Port mirroring on the switch is now disabled. You can disconnect the network analyzer from the destination port and use the port for normal network operations.

Section II

Advanced Operations

The chapters in this section explain some of the more advanced features of an AT-8500 Series switch. The chapters include:

Chapter 10
File System

This chapter describes the AT-S62 file system, and how you can use the file system to copy, rename, and delete system files. This chapter also explains how you can use the file system to select which boot configuration file you want the switch to use the next time the device is reset or power cycled. This chapter contains the following sections:

File System Overview

The AT-S62 management software has a file system of 2 megabytes for storing system files. You can view the file system, as well as copy, rename, and delete files. The following file types are supported by the AT-S62 file system:

- Boot configuration files
- Encryption keys
- CA and self-signed certificates
- Certificate enrollment requests
- Event logs

For an explanation of a boot configuration file, refer to “Working with Boot Configuration Files” on page 176.

Public encryption keys, public certificates, and certificate enrollment request files are related to the Secure Sockets Layer (SSL) certificates feature described in Chapter 31, “Encryption Keys” on page 687, and Chapter 32, “PKI Certificates and SSL” on page 705. Refer to those chapters for background information on those files.

Note
The certificate file, certificate enrollment request file, and key file are only supported on the version of AT-S62 management software that features SSL and PKI security.

This chapter does not explain how to download or upload a file from the AT-S62 file system to a management workstation or an TFTP server. For those instructions, refer to Chapter 11, “File Downloads and Uploads” on page 187.

Note
The file system may contain one or more ENC.UKF files. These are encryption key pairs. These files cannot be deleted or copied in the file system. For instructions on deleting an encryption key pair, refer to “Deleting an Encryption Key” on page 699.

The file system should not be used to store the switch’s AT-S62 image file.

File Naming Conventions

The file system is a flat file system which means directories are not supported. Files are uniquely identified by a file name in the following format:

```
filename.ext
```

where:

- *filename* is a descriptive name for the file, and may be one to sixteen characters in length. Valid characters are lowercase letters (a–z), uppercase letters (A–Z), digits (0–9), and the following characters: ~ ' @ # $ % ^ & ( ) _ - { }+. Invalid characters are: ! * = "| \ [ ] ; : ? / , < >.
- *ext* is a file name extension of three characters in length, preceded by a period (.). The extension is used by the switch to determine the file type.

Table 2. File Extensions and File Types

Extension	File Type
.cfg	Configuration file
.cer	Certificate file
.csr	Certificate enrollment request
.key	Public encryption key

The following is an example of a valid file name for a configuration file:

```
standardconfig.cfg
```

The following is an example of an invalid file name:

```
sys/head_o.cfg
```

The backslash character (/) is not a valid character because subdirectories are not supported.

The file system displays filenames and directories in DOS 28.3 format. Filenames and directories longer than 32 bytes are represented in DOS 8.3 format.

Using Wildcards to Specify Groups of Files

You can use the asterisk character (*) as a wildcard character in some fields to identify groups of files. In addition, a wildcard can be combined with other characters. The following are examples of valid wildcard expressions:

```
*.cfg
*.key
28*.cfg
```

Working with Boot Configuration Files

A boot configuration file contains the commands for configuring the switch's parameter settings whenever you power cycle or reset the device. The commands in the file recreate the VLANs, port settings, spanning tree settings, port trunks, port mirrors, and so on.

You can store multiple boot configuration files on a switch, but only one can be active at a time. The switch uses the active boot file to configure itself whenever reset or power cycled. It also updates the active boot file with the latest parameter settings whenever you select the Save Configuration Changes option from the Main Menu or use the Save Configuration command from the command line interface.

You can create different configuration files and store them in the switch's file system. For instance, you might create a backup of a configuration file to protect against the loss of the file, or you might create different configuration files to see which works best on the switch and for your network. You can also copy configuration files onto different switches to save yourself the trouble of manually configuring AT-8500 Series switches that are to have similar configurations.

The procedures in this section explain how to create a boot configuration file, set the active boot configuration file, view the contents of a configuration file, and edit a file. The procedures are:

To display a list of the configuration files that exist on the switch, see "Displaying System Files" on page 185.

Creating a Boot Configuration File

This procedure explains how to create a new boot configuration file on the switch. You might want to create a boot configuration file to download it onto another switch. Or, you might want to create a backup of your current configuration. This procedure consists of three phases:

- ❒ Phase 1: Creating a Configuration File
- ❒ Phase 2: Configuring the Switch's Parameter Settings
- ❒ Phase 3: Selecting the Active Configuration File for the Switch

Phase 1: Creating a Configuration File

Before you begin to configure the switch with the parameter settings that you want to save in a new configuration file, you should first create the file. Configuring the parameters first and then creating the new configuration file might cause you to inadvertently change a configuration file you might not want to change.

To perform this phase, do the following:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **9** to select System Utilities.
3. From the System Utilities menu, type **1** to select File Operations.

   The File Operations menu is shown in Figure 48.

```
            Allied Telesyn AT-8524M Series - ATS62
                       Production Switch
User: Manager                                11:20:02 02-Jan-2006
                         File Operations

1 - Boot Configuration File ............. boot.cfg (Exists)
2 - Current Configuration ............... boot.cfg
3 - Create Configuration File

4 - Copy File
5 - Rename File
6 - Delete File

7 - View File
8 - List Files

9 - Format Flash Drive
R - Return to Previous Menu

Enter your selection?
```

Figure 48. File Operations Menu

Option 1 - Boot Configuration File specifies the file that is updated whenever you save a configuration change using the Save Configuration Changes option in the Main Menu or the Save Configuration command in the command line interface. It is also the boot file that the switch will use the next time you reset or power cycle the unit. Option 2 - Current Configuration specifies the boot configuration file the switch used the last time it was reset or power cycled.

Caution

Option 9 - Format Flash Drive should be used with care. It deletes all files in the file system, including configuration files, encryption keys, event logs, etc. For instructions, refer to “Deleting the System Files” on page 74.

4. Type **3** to select Create Configuration File.

   The following prompt is displayed:

   ```
   Enter the file name (or None):
   ```

5. Enter a file name for the new configuration file.

   The file name can be up to 16 alphanumeric characters. Spaces are allowed. The filename must include the extension “.cfg”. See “File Naming Conventions” on page 175.

   Note

   If the filename already exists, the system displays a message asking if you want to overwrite the existing file.

   Note

   You cannot name a configuration file “default.cfg.” This file name is reserved by the switch.

   The management software creates the new configuration file with the switch’s current settings and stores it in the file system.

6. Type **1** to select Boot Configuration File.

   The following prompt is displayed:

   ```
   Enter the file name:
   ```

7. Enter the same file name that you entered in Step 5.

   This makes your new configuration file the active file on the switch. Any changes you now make to the switch’s parameter settings are saved to this file.

   The file name will now appear following selection 1 in the File Operations menu. The file name should be followed by “Exist”, meaning that the file exists in the switch’s file system. If “Not Found” appears instead, you probably enter the name incorrectly, in which case you need to repeat Steps 6 and 7.

Phase 2: Configuring the Switch's Parameter Settings

Now that you have created a configuration file and designated it as the active boot configuration file on the switch, you can now configure the switch's parameter settings by making those changes that you want the new configuration file to contain. Once you have done that, be sure to save your changes to the configuration file by returning to the Main Menu and typing **S** to select Save Configuration Changes. Failure to save your changes will mean that the configuration file will not contain the new parameter settings.

Note
Only the active boot configuration file is changed when you select the Save Configuration Changes option in the Main Menu. No other boot configuration files stored on the switch are altered.

Phase 3: Selecting the Active Configuration File for the Switch

You have now created the configuration file, made the necessary changes to the switch's parameter settings, and saved the changes. If you want the switch to use this new configuration file the next time you reset or power cycle the switch, no further steps are necessary. The new configuration file is already the active boot file on the device.

If you want the switch to use a different file as the active configuration file, then perform the procedure in "Setting the Active Boot Configuration File" on page 179.

If you want to create another new configuration file, repeat this procedure starting with Phase 1.

Setting the Active Boot Configuration File

This procedure selects the active boot configuration file on the switch. The switch uses the active configuration file the next time the unit is reset or power cycled to set its parameter settings. You can select a configuration file that you created on the switch or that you downloaded onto the switch from another switch.

The switch comes with one default configuration file, called "default.cfg." This is the default active configuration file.

Note
The active boot configuration file is updated whenever you select the Save Configuration Changes from the Main Menu or the Save Configuration command from the command line interface.

To select the active boot configuration file for the switch, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **9** to select System Utilities.
3. From the System Utilities menu, type **1** to select File Operations.

   The File Operations menu is shown in Figure 48 on page 177.
4. Type **1** to select Boot Configuration File.

   The following prompt is displayed:

   ```
   Enter the file name:
   ```
5. Enter the file name of the configuration file you want the switch to use the next time it is reset or power cycled.

   The file name will now appear following selection 1 in the File Operations menu. The file name should be followed by “Exist”, which means that the file exists in the switch’s file system. In the future, the switch will use the newly selected configuration file whenever you reset the unit, unless you designate another boot configuration file as the active boot file.

   Note
   If “Not Found” appears, the file does not exist. If you reboot the switch using a nonexistent configuration file the switch is reset to its factory default settings.
6. Do one of the following:
   - If you want to configure the switch using the parameter settings in this boot configuration file, do **not** select Save Config. Instead, reset or power cycle the switch.
   - If you want to overwrite the settings in the configuration file with the switch’s current operating settings, select **Save Config**.

Viewing a Boot Configuration File

Use the following procedure to view the contents of a configuration file. (To display the names of the configuration files on the switch, see “Displaying System Files” on page 185.)

This procedure starts from the File Operations menu. If you are unsure how to display the menu, perform steps 1 to 3 in “Setting the Active Boot Configuration File” on page 179.

To view the contents of a configuration file, perform the following procedure:

1. From the File Operations menu, type **7** to select View File.

   The following prompt is displayed:

   ```
   Enter file name:
   ```

2. Enter the name of the configuration file you want to view.

   The contents of the configuration file are displayed in the View File menu. An example is shown in Figure 49.

```
            Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                           Production Switch
User: Manager                                    11:20:02 02-Jan-2006
                                  View File
Configuration File: mydefault.cfg
------------------------------------------------------------------------
#
# System Configuration
#
set system name="Production Switch"
set system contact="Jane Smith"
set system location="Building 5"

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 49. View File Menu

A configuration file contains those switch settings that differ from the AT-S62 default values. The parameter settings are shown in their command line equivalents. The switch executes the commands in the boot configuration file to configure its settings when it is reset or power cycled. For information on command line commands, refer to the **AT-S62 Command Line User's Guide**. The information in this menu is for viewing purposes only.

3. Type **N** for Next Page and **P** for Previous Page to scroll through the file.

Editing a Boot Configuration File

You can edit a boot configuration file using a text editor on your management workstation. To edit a file, you must upload it from the switch to your management workstation. You cannot edit a boot configuration file directly on the switch. Once you have edited the file, you can download it back to the switch and make it the active boot configuration file.

For instructions on how to upload a configuration file from a switch to your management workstation, refer to “Uploading a System File” on page 209. For instructions on how to download a configuration file from your workstation back to the switch, refer to “Downloading a System File” on page 202. For instructions on how to designate an active boot configuration file, refer to “Setting the Active Boot Configuration File” on page 179.

Here are several guidelines to editing a boot configuration file:

- The text editor must be able to store the file as ASCII text. Do not insert special formatting codes, such as boldface or italics, into a boot configuration file.
- The configuration file must contain AT-S62 command line commands. You enter the commands you want the switch to perform when reset or power cycled. For a description of the commands, refer to the **AT-S62 Command Line User’s Guide**.
- A boot configuration file is divided into sections with each section devoted to the commands of a particular function. For example, the VLAN Configuration section should contain commands for creating VLANs or for setting the VLAN mode. When entering new commands, be sure to place them in the appropriate sections.
- Each command must start flush left against the margin.
- To comment out a command so that the switch does not perform it, precede the command with the symbol “#”.
- You should test the commands manually by entering them at the command line before inserting them into a boot configuration file. This helps to insure that you understand the syntaxes and parameters of the commands and that the commands produce the desired results.

Troubleshooting a Boot Configuration File

If a boot configuration file contains an invalid or incorrect command, the switch, when reset or power cycled, will stop processing the configuration file at the point of the invalid command. The invalid command and any commands following it in the file will not be performed. To troubleshoot a configuration file, start a local management session with the switch and reset the device. Messages on the screen during the boot up and configuration process will indicate the line in the configuration file with the error. You can download the file to your management workstation and edit it to correct the error.

Copying, Renaming, and Deleting System Files

Use this procedure to copy, rename, and delete system files. To view a list of system file names, see "Displaying System Files" on page 185.

Note
Files with the extension UKF are encryption key pairs. These files cannot be copied, renamed, or deleted from the file system. To delete a key pair from the switch, refer to "Deleting an Encryption Key" on page 699.

To copy, rename, or delete a file in the file system, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **9** to select System Utilities.
3. From the System Utilities menu, type **1** to select File Operations.

   The File Operations menu is shown in Figure 48 on page 177.
4. To copy a file, do the following:

   a. From the File Menu, type **4** to select Copy File.

   Note
   Selecting Copy File does not allow you to overwrite files.

   The following prompt is displayed:

   ```
   Enter the source file name:
   ```

   b. Enter the name of the file you want to copy.

   The following prompt is displayed:

   ```
   Enter the destination file name:
   ```

   c. Enter the new file name.

   You can enter a file name of up to 16 alphanumeric characters, followed by a 3 letter extension. You should keep the same extension as the original filename.

   The following message is displayed:

   ```
   Please wait...
   Press any key ...
   ```

d. Press any key to return to the File Operations menu.

5. To rename a system file, do the following:

   a. From the File Operations menu, type **5** to select Rename File.

      The following prompt is displayed:

      ```
      Enter the source file name:
      ```

   b. Enter the name of the file you want to rename.

      The following prompt is displayed:

      ```
      Enter the destination file name:
      ```

   c. Enter the new name for the file.

      You can enter a file name of up to 16 alphanumeric characters, followed by a 3 letter extension. You must keep the same extension.

      The following message is displayed:

      ```
      Please wait...
      Press any key ...
      ```

   d. Press any key to return to the File Operations menu.

6. To delete a system file, do the following:

   a. From the File Operations menu, type **6** to select Delete File.

      The following prompt is displayed:

      ```
      Enter file name to be deleted:
      ```

   b. Enter the name of the file you want to delete.

      The following prompt is displayed:

      ```
      Please wait...
      Press any key ...
      ```

   c. Press any key to return to the File Operations menu.

Note
Deleting the active boot configuration file causes the switch to use its default settings the next time you reboot or power cycle the switch, unless you select another active boot configuration file. For instructions on how to change the active boot configuration file, see “Setting the Active Boot Configuration File” on page 179.

Displaying System Files

Use this procedure to display a list of the system files currently stored on the switch. For information about shortcuts for specifying file names, see "File Naming Conventions" on page 175.

To display a list of current system file names, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **9** to select System Utilities.
3. From the System Utilities menu, type **1** to select File Operations.

   The File Operations menu is shown in Figure 48 on page 177.
4. From the File Operations menu, type **8** to select List Files.

   The following prompt is displayed:

   ```
   Enter file name pattern to list:
   ```
5. Enter a configuration file name or pattern using the wildcard "*". Below are examples of how to use the wildcard to display different files.

   To display a list of all the files, enter:

   ```
   *.*
   ```

   To display a list of the certificate files, enter:

   ```
   *.cer
   ```

   To display a list of the configuration files, enter:

   ```
   *.cfg
   ```

   To display a list of the key files, enter:

   ```
   *.key
   ```

   To display a list of the files that begin with the letter t, enter:

   ```
   t*.*
   ```

The List Files menu is displayed. An example of the menu is shown in Figure 50.

```
            Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                          Production Switch
User: Manager                                          11:20:02 02-Jan-2006
                                List Files
File Name           Device      Size (Bytes) Last Modified
-------------------------------------------------------------------
default.cfg         flash       805          01/10/2002 12:01:16
boot.cfg            flash       1249         10/24/2003 16:50:40
newcfg.cg           flash       1082         07/12/2003 16:59:06
serverkey150.key    flash       768          11/30/2003 19:17:35
ProdSw.cer          flash       1024         11/30/2003 20:38:20
ProdSw2.cer         flash       560          12/11/2003 20:56:13

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 50. List Files Menu

The columns in the List Files menu are described below:

- The File Name column contains the name of the system file.
- The Device column indicates the location of the file. For an AT-8500 Series switch, this is always Flash.
- The Size column indicates the size of the file, in bytes.
- The Last Modified column lists the time the file was created or last modified, in the following date and time format: month/day/year hours:minutes:seconds.

The information in this menu is for viewing purposes only.

Chapter 11

File Downloads and Uploads

This chapter contains procedures for downloading a new AT-S62 image file onto the switch. This chapter also contains procedures for uploading and downloading system files, such as boot configuration files, from the file system in a switch. The procedures in this chapter are:

Note
For instructions on how to obtain the latest version of the AT-S62 management software, refer to “Management Software Updates” on page 25.

Downloading a New AT-S62 Image File onto a Switch

The procedures in this section explain how to download a new AT-S62 image file onto the switch. These procedures are used to update the AT-S62 image file on a switch with a new version of the file. If you have an enhanced stack of AT-8500 Series switches, the easiest way to update the switches is to first update the master switch by performing one of the procedures in this section and then instructing the master switch to upload its image file to the other switches, as explained in "Uploading an AT-S62 Image File Switch to Switch" on page 196.

There are two ways that you can download a new image file onto a switch. You can do it from a local management session using either Xmodem or TFTP, or from a remote Telnet management session using TFTP, exclusively. Each method is described in a separate section in this section. The procedures are:

- "Downloading an AT-S62 Image from a Local Management Session" on page 190
- "Downloading an AT-S62 Image from a Telnet Management Session" on page 194

Caution
Installing a new AT-S62 image file will reset the switch. Some network traffic may be lost.

Guidelines

The following guidelines apply to both Xmodem and TFTP downloads:

- The following procedures download a new AT-S62 image file into the application block portion of the switch's flash memory. The application block is the area of memory reserved for the active AT-S62 image file on a switch and is separate from the file system.
- Alternatively, you can download the image file into the switch's file system and then later copy it into the application block. The drawback to this approach is that the image file will require nearly all 2 megabytes of the file system, leaving almost no room for other files, such as configuration files and SSL certificates. To download an image file into the file system rather than the application block, refer to "Downloading a System File" on page 202.
- All models of the AT-8500 Series switches use the same AT-S62 management software image.
- The current configuration of a switch is retained when a new AT-S62 software image is installed. If you want to return a switch to its default configuration values, refer to "Returning the AT-S62 Software to the Factory Default Values" on page 73.

- The AT-S62 image file contains the bootloader for the switch. You cannot load the image file and bootloader separately.

The following guidelines apply to an Xmodem download:

- Xmodem can only download the image file onto the switch where you started the local management session. You cannot use Xmodem to download a new image file to a switch accessed through enhanced stacking.
- The new AT-S62 image file must be stored on the computer or terminal connected to the RS232 Terminal Port on the switch.

The following guidelines apply to a TFTP download:

- Your network must have a node with TFTP server software.
- You must store the new AT-S62 image file on the server.
- You should start the TFTP server software before you begin the download procedure.
- The switch where you are downloading the new image file must have an IP address and subnet mask, such as a master switch of an enhanced stack. If the switch does not have an IP address, such as a slave switch, you can perform the download from a local management session of the switch using Xmodem or, alternatively switch to switch, as explained in "Uploading an AT-S62 Image File Switch to Switch" on page 196.

The following procedures assume that you have already obtained the new software from Allied Telesyn and stored it on the management workstation from which you will be performing the procedure, or on the TFTP server.

Downloading an AT-S62 Image from a Local Management Session

Review the "Guidelines" on page 188 before performing the following download procedure.

To download a new software image onto a switch from a local management session using Xmodem or TFTP, perform the following procedure:

1. Establish a local management session on the switch where you want to download the new management software.

2. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

3. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

4. From the System Utilities menu, type **2** to select Downloads and Uploads.

   The Downloads and Uploads menu is shown in Figure 51.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                       11:20:02 02-Jan-2006
                  Downloads and Uploads
1 - Download Application Image/BootLoader
2 - Upload Application Image/BootLoader

3 - Download a file
4 - Upload a file

R - Return to Previous Menu

Enter your selection?
```

Figure 51. Downloads and Uploads Menu

Note
The Downloads and Uploads menu for an AT-8524POE switch includes the selection 5 - Download PoE Firmware. This selection is intended for Allied Telesyn Technical Support only.

5. From the Downloads and Uploads menu, type **1** to select Download Application Image/Bootloader.

   The following prompt is displayed:

   ```
   Download Method/Protocol [X-Xmodem, T-TFTP]:
   ```

6. To download the AT-S62 image file using Xmodem, go to Step 7. To download the file using TFTP, do the following:

   a. Type **T**.

      The following prompt is displayed:

      ```
      TFTP Server IP address:
      ```

   b. Enter the IP address of the TFTP server.

      The following prompt is displayed:

      ```
      Remote File Name:
      ```

   c. Enter the file name of the AT-S62 image file stored on the TFTP server. (Be sure to include the ".img" extension.)

      The following message is displayed:

      ```
      Getting the file from Remote TFTP Server - Please wait
      ...
      ```

   d. If you have not already done so, start the TFTP server software.

      After the switch has downloaded the image file, the following message is displayed:

      ```
      File received successfully!
      ```

      After receiving the file, the switch compares the version numbers of the new and existing image files. If the new image file has an earlier or the same version number as the file in the switch's application block, the update process is cancelled. If the new image file has a newer version number, the switch writes the file to the application block portion of flash memory and then resets.

Caution

The switch will not forward network traffic while writing the image to flash and during the reset process. This can take several minutes to complete.

      This completes the procedure for downloading a new AT-S62 image file from an Xmodem management session using TFTP.

7. To download a file using Xmodem, type **X** at the prompt displayed in Step 5.

The following prompt is displayed:

```
You are going to invoke the Xmodem download utility.
Do you wish to continue? [Yes/No]

Note: Please select 1K Xmodem protocol for faster
download.
```

8. Type **Y** for Yes.

   The prompt “Downloading” is displayed.

9. Begin the file transfer.

Note
The transfer protocol must be Xmodem or 1K Xmodem.

As an example, steps 10 through 13 illustrate how to download a file using the Hilgraeve HyperTerminal program.

10. From the HyperTerminal main window, select the **Transfer** menu. Then select **Send File** from the pull-down menu, as shown in Figure 52.

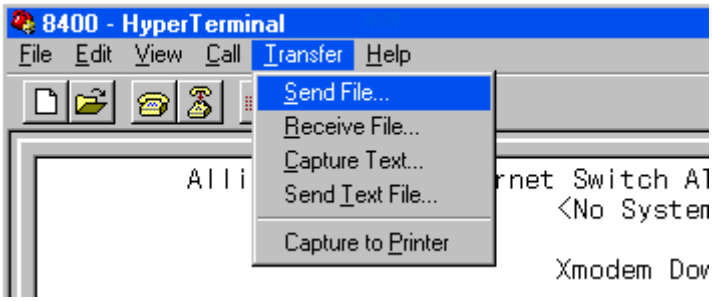

Figure 52. Local Management Window

The Send File window is shown in Figure 53.

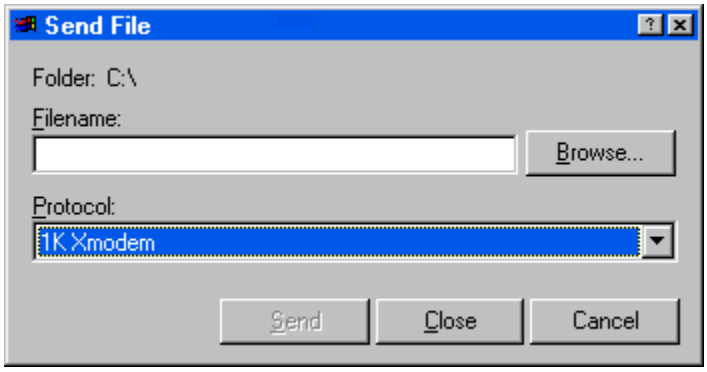

Figure 53. Send File Window

11. Click **Browse** and specify the location and file to be downloaded onto the switch.

12. Click in the Protocol field and select as the transfer protocol either Xmodem or, for a faster download, 1K XModem.

13. Click **Send**.

    The software immediately begins downloading the file onto the switch. The Xmodem File Send window in Figure 54 displays the status of the software download. The download process takes several minutes to complete.

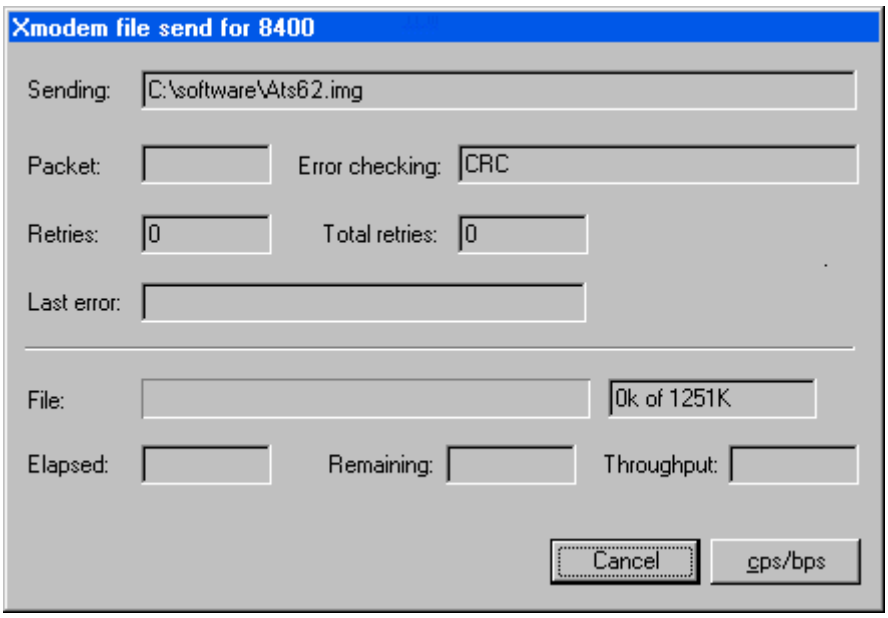

Figure 54. XModem File Send Window

After receiving the file, the switch compares the version numbers of the new and existing image files. If the new image file has the same or an earlier version number as the file in the application block, it cancels the update process. If the new image file has a newer version number, the switch writes the file to the application block portion of flash memory and then resets.

Caution

The switch will not forward any network traffic while writing the image to flash and during the reset process. This can take several minutes to complete.

This completes the procedure for downloading a new AT-S62 image file onto a switch from an Xmodem management session.

Downloading an AT-S62 Image from a Telnet Management Session

Review the "Guidelines" on page 188 before performing the following download procedure.

To download a new AT-S62 image onto the application block portion of the switch's flash memory, making it the active image file on the switch, from a Telnet management session using TFTP, perform the following procedure:

1. Establish a Telnet management session on the switch where you want to download the new management software.
2. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.
3. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.
4. For the System Utilities menu, type **2** to select Downloads and Uploads.

   The Downloads and Uploads menu is shown in Figure 51 on page 190.
5. Type **1** to select Download Application Image/Bootloader.

   The following prompt is displayed:

   ```
   Only TFTP downloads are available for a Telnet access

   TFTP Server IP address:
   ```
6. Enter the IP address of the TFTP server.

   The following prompt is displayed:

   ```
   Remote File Name:
   ```
7. Enter the file name of the AT-S62 image file on the TFTP server to download onto the switch. (Be sure to include the ".img" extension.)

   The following message is displayed:

   ```
   Getting the file from Remote TFTP Server - Please wait
   ...
   ```
8. If you have not already done so, start the TFTP server software.

   After downloading the image file, the switch displays the following message:

   ```
   File received successfully!
   ```

After receiving the file, the switch compares the version numbers of the new and existing image files. If the new image file has the same or an earlier version number as the existing file in the application block, the switch cancels the update process. If the new image file has a newer version number, the switch writes the file to the application block portion of flash memory and then resets.

Caution

The switch will not forward network traffic while writing the image to flash and during the reset process. This can take several minutes to complete.

This completes the procedure for downloading a new AT-S62 image file from a Telnet management session.

Uploading an AT-S62 Image File Switch to Switch

This procedure uploads an AT-S62 software image from a master AT-8500 Series switch to other AT-8500 Series switches in an enhanced stack. Commonly referred to as a switch to switch transfer, this transfer method can simplify the task of updating the AT-S62 image file in the AT-8500 Series switches in an enhanced stack. Rather than manually updating each switch, you can manually update the master switch's image file and then instruct it to upload its image file to the other switches, automatically. (For instructions on how to update the AT-S62 image file on a master switch, refer to "Downloading a New AT-S62 Image File onto a Switch" on page 188.)

You can perform this procedure from a local or Telnet management session.

Guidelines

Review the following guidelines before performing the procedure:

- This procedure downloads a new AT-S62 image file into the application block portion of the switch's flash memory. The application block is the area of memory reserved for the active AT-S62 image file on a switch and is separate from the file system.
- Alternatively, you can download the image file into the switch's file system and then later copy it into the application block using the LOAD command in the command line interface. To download an image file into the file system rather than the application block, refer to "Downloading a System File" on page 202.
- All models of the AT-8500 Series switches use the same AT-S62 management software image.
- The current configuration of a switch is retained when a new AT-S62 software image is installed. If you want to return a switch to its default configuration values, refer to "Returning the AT-S62 Software to the Factory Default Values" on page 73.
- The AT-S62 image file contains the bootloader for the switch. You cannot load the image file and bootloader separately.

To upload a management software image from a master switch to other AT-8500 Series switches in the same enhanced stack, perform the following procedure:

1. From the Main Menu, type **8** to select Enhanced Stacking.

   The Enhanced Stacking menu is shown in Figure 15 on page 84.

2. From the Enhanced Stacking menu, type **2** to select Stacking Services.

Note
The "2 - Stacking Services" selection is only available on a master switch.

The Stacking Services menu is shown in Figure 16 on page 85.

3. Type **1** to select Get/Refresh List of Switches. The master switch polls the subnet for the switches in its enhanced stack and displays the switches in the Stacking Services menu.

4. Type **4** to select Load Image/Bootloader File.

   The following prompt is displayed:

   ```
   Enter the list of switches ->
   ```

5. Enter the number (Num column in the menu) of the AT-8500 Series switch whose software you want to update. You can specify more than one switch at a time (for example, 2,4,5).

   Note
   The AT-S62 image file is only supported on AT-8500 Series switches.

   The following prompt is displayed:

   ```
   Do you want to show remote switch burning flash -> [Yes/
   No]
   ```

6. You can respond with Yes or No to this prompt. It does not affect the upload.

   The following prompt is displayed:

   ```
   Do you want confirmation before downloading each switch
   -> [Yes/No]
   ```

7. If you answer Yes to this prompt, the management software displays a confirmation message before uploading the image file to a switch. If you answer No, the management software does not display a confirmation prompt before uploading the file.

   The management software begins the upload. The management software notifies you when the upload is complete.

   After receiving the file, a switch compares the version numbers of the new and existing image files. If the new image file has the same or an earlier version number as the file in the application block, the switch cancels the update process. If the new image file has a newer version number, the switch writes the file to the application block portion of flash memory and then resets.

Caution

The switch will not forward network traffic while writing the image to flash and during the reset process. This can take several minutes to complete.

Uploading an AT-S62 Configuration File Switch to Switch

This procedure uploads a boot configuration file from a master AT-8500 Series switch to another AT-8500 Series switch in an enhanced stack. This procedure provides you with an easy way of distributing a configuration file to different switches that are to share a similar configuration. For background information on configuration files, refer to "Working with Boot Configuration Files" on page 176.

Note
You can perform this procedure from a local or Telnet management session.

Guidelines

Please review the following guidelines before performing the procedure:

- The procedure gives you the choice of uploading the master switch's active boot configuration file or another configuration file in the switch's file system. If you choose the switch's current boot configuration file, the following information in the file is not included in the upload: IP address, subnet mask, gateway address, switch name, contact, location, and the master mode setting. This is to prevent two switches from having the same IP configuration. However, the switch receiving the configuration file does not retain its current settings to these parameters. Instead, they are returned to their default values.
- If you choose to upload another configuration file from the master switch's file system, the entire file without modification is uploaded. This type of upload should be performed with care. If you upload onto more than one switch a configuration file that assigns the units a specific IP address, multiple switches will end up using the same IP address.
- This procedure can be performed from a local or Telnet management session.
- Once the upload is complete, the switch that received the configuration file marks it as its active boot configuration file and resets. Some network traffic may be lost while the switch reloads its operating software. After the reset is complete, the switch operates with the parameter settings contained in the uploaded configuration file.
- A configuration file should only be uploaded onto the same model of switch as the original switch (for example, AT-8524M to AT-8524M). Allied Telesyn does not recommend uploading a configuration file onto a switch of a different model (for example, AT-8524M to AT-8516F/SC). Undesirable switch behavior may result.

Caution
This procedure causes the switch to reset. Some network traffic may be lost.

To upload a boot configuration file from the master switch to another switch in an enhanced stack, perform the following procedure:

1. From the Main Menu, type **8** to select Enhanced Stacking.

   The Enhanced Stacking menu is shown in Figure 15 on page 84.

2. From the Enhanced Stacking menu, type **2** to select Stacking Services.

   Note
   The "2 - Stacking Services" selection is only available on master switches.

   The Stacking Services menu is shown in Figure 16 on page 85.

3. Type **1** to select Get/Refresh List of Switches. The master switch polls the network for all enhanced stacking switches in the subnet and displays the switches in the Stacking Services menu.

4. Type **5** to select Load Configuration File.

   The following prompt is displayed:

   ```
   Remote switches will reboot after load is complete
   Do you want to load the last saved master configuration?
   [Yes/No] ->
   ```

5. If you want the master switch to upload its active boot configuration file onto the other switch, type **Y** for yes and go to step 7. If you want the master switch to upload another configuration file from its file system, type **N** for no.

   The following prompt is displayed:

   ```
   Enter the configuration file name ->
   ```

6. Enter the name of the configuration file on the master switch you want to upload. The name must include the suffix ".cfg". (To view the names of the configuration files, refer to "Displaying System Files" on page 185.)

   The following prompt is displayed:

   ```
   Enter the list of switches ->
   ```

7. Enter the number (Num column in the menu) of the AT-8500 Series switch where you want to upload the configuration file. You can specify more than one switch at a time (for example, 2,4,5).

   Note
   An AT-8500 Series configuration file is only compatible with other AT-8500 Series switches. Do not upload the file onto any other type of enhanced stacking switch.

   The following prompt is displayed:

   ```
   Do you want confirmation before downloading each switch
   -> [Yes/No]
   ```

8. If you answer Yes to this prompt, the management software prompts you with a confirmation message before uploading the file to a switch. If you answer No, the management software does not display a confirmation prompt before uploading the file.

   The management software begins the upload. A switch, after receiving the configuration file, automatically designates it as its new active boot configuration file and resets. After the reset is complete, the switch operates with the parameter settings in its new configuration file.

   

   Caution
   The switch will not forward network traffic during the reset process. This can take several minutes to complete.

Downloading a System File

This section contains the following procedures:

- “Downloading a File from a Local Management Session” on page 203
- “Downloading a File from a Telnet Management Session” on page 207

Both procedures are used to download files into a switch’s file system. One procedure downloads files from a local management using either Xmodem or TFTP, and the other explains how to do it from a Telnet management session, which only supports TFTP.

There are only two files that you are ever likely to download into a switch’s file system:

- Boot configuration file
- Certificate enrollment request

You might have edited a boot configuration file at your management workstation and want to download it onto a switch prior to designating it as the active boot configuration file. Or, you might have obtained a CA certificate for the switch to add encryption to your web browser management sessions.

You can also use this procedure to store an AT-S62 image file in the switch’s file system. However, downloading an image file into the file system should be performed with care. First, for an image file to be the active image file on a switch it has to be stored in the switch’s application block, which is a separate part of flash memory from the file system. Second, the image file will take up almost all 2 megabytes of the file system, leaving little room for other files. If you want to download an AT-S62 image file so that it is the active image file on the unit, see “Downloading a New AT-S62 Image File onto a Switch” on page 188 or “Uploading an AT-S62 Image File Switch to Switch” on page 196.

Guidelines

Review the following guidelines before downloading a file onto a switch.

- You can use either Xmodem or TFTP to download files from a local management session.
- You must use TFTP to download files from a Telnet management session.
- You cannot download a private encryption key onto a switch, but you can a public key. However, since the switch can only use those encryption keys it has generated itself, Allied Telesyn recommends against downloading any keys onto the switch.

These guidelines apply to an Xmodem download:

- Xmodem can only download a file onto the switch where you started the local management session. You cannot use Xmodem to download a file onto a switch accessed through enhanced stacking.
- The file to be downloaded must be stored on the computer or terminal connected to the RS232 Terminal Port on the switch.

These guidelines apply to a TFTP download:

- Your network must have a node with TFTP server software.
- The file to be downloaded must be stored on the TFTP server.
- You should start the TFTP server software before you begin the download procedure.
- The switch where you are downloading the file must have an IP address and subnet mask, such as a master switch of an enhanced stack. For switches without an IP address, such as slave switches, you can download the file from a local management session of the switch using Xmodem.

Downloading a File from a Local Management Session

Review "Guidelines" on page 202 before performing this procedure.

To download a file onto a switch from a local management session using Xmodem or TFTP, perform the following procedure:

1. Establish a local management session on the switch where you want to download the system file.
2. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.
3. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.
4. For the System Utilities menu, type **2** to select Downloads and Uploads.

   The Downloads and Uploads menu is shown in Figure 51 on page 190.
5. Type **3** to select Download a File.

   The following prompt is displayed:

   ```
   Download Method/Protocol [X-Xmodem, T-TFTP]:
   ```

6. To download a system file using Xmodem, go to Step 7. To download a file using TFTP, do the following:

   a. Type **T**.

      The following prompt is displayed:

      ```
      TFTP Server IP address:
      ```

   b. Enter the IP address of the TFTP server.

      The following prompt is displayed:

      ```
      Remote File Name:
      ```

   c. Enter the name of the file on the TFTP server to download onto the switch. You can specify only one file.

      The following prompt is displayed:

      ```
      Local File Name:
      ```

   d. Enter a name for the file for when it is stored in the switch's file system. The name can be up to 32 alphanumeric characters. Spaces are allowed, but special characters (e.g., /, \, #, and %) should not be used. The extension should match its file type. Supported extensions and file types are listed in Table 3.

      Table 3. File Name Extensions

Extension	File Type
.cfg	Configuration file
.cer	Certificate file
.img	AT-S62 image file

      The following message is displayed:

      ```
      Getting the file from Remote TFTP Server - Please wait
      ...
      ```

   e. If you have not already done so, start the TFTP server software.

      After downloading the system file, the switch displays the following message:

      ```
      File received successfully!
      ```

      This completes the process for downloading a file using TFTP.

f. If you downloaded a new configuration file and you want to make it the switch's active boot file, go to "Setting the Active Boot Configuration File" on page 179. If you downloaded a CA certificate and want to add it to the certificate database, refer to "Adding a Certificate to the Database" on page 722.

7. To download a file using Xmodem, type **X** at the prompt displayed in Step 5.

   The following prompt is displayed:

   ```
   Local File Name:
   ```

8. Enter a name for the file. The file is given this name when stored in the switch's file system. When naming a file, be sure to give it an extension that corresponds to its file type. The extensions and file types are listed in Table 3 on page 204.

   The following prompt is displayed:

   ```
   You are going to invoke the Xmodem download utility.
   Do you wish to continue? [Yes/No]

   Note: Please select 1K Xmodem protocol for faster
   download.
   ```

9. Type **Y** for Yes.

   The prompt "Downloading" is displayed.

10. Begin the file transfer of the system file.

> **Note**
> The transfer protocol must be Xmodem or 1K Xmodem.

Steps 11 through 14 illustrate how to download a file with the Hilgraeve HyperTerminal program.

11. From the HyperTerminal main window, select **Send File** from the **Transfer** pull-down menu, as shown in Figure 52.

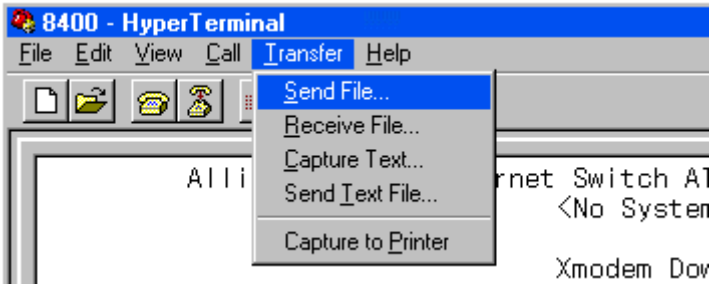

Figure 55. Local Management Window

The Send File window is shown in Figure 53.

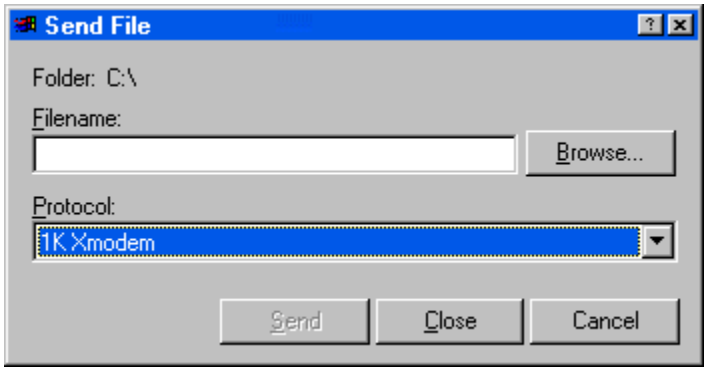

Figure 56. Send File Window

12. Click **Browse** and specify the location and system file to be downloaded onto the switch.

13. Click in the Protocol field and select as the transfer protocol either Xmodem or, for a faster download, 1K XModem.

14. Click **Send**.

    The file immediately begins downloading onto the switch. The Xmodem File Send window in Figure 54 displays the status of the download.

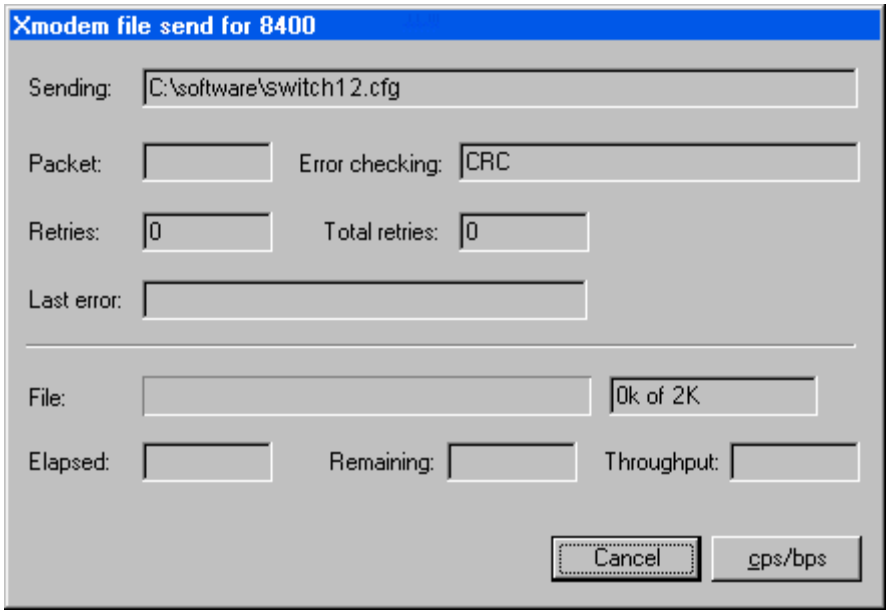

Figure 57. XModem File Send Window

The download is complete when the Downloads and Uploads menu is displayed.

15. If you downloaded a new configuration file and you want to make it the switch's active boot file, go to "Setting the Active Boot Configuration File" on page 179. If you downloaded a CA certificate and need to add it to the certificate database, refer to "Adding a Certificate to the Database" on page 722.

Downloading a File from a Telnet Management Session

Review "Guidelines" on page 202 before performing this procedure.

To download a file onto a switch from a Telnet management session using TFTP, perform the following procedure:

1. Establish a Telnet management session on the switch where you want to download the file.

2. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

3. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

4. For the System Utilities menu, type **2** to select Downloads and Uploads.

   The Downloads and Uploads menu is shown in Figure 51 on page 190.

5. Type **3** to select Download a File.

   The following prompt is displayed:

   ```
   Only TFTP downloads are available for a Telnet access

   TFTP Server IP address:
   ```

6. Enter the IP address of the TFTP server.

   The following prompt is displayed:

   ```
   Remote File Name:
   ```

7. Enter the name of the file on the TFTP server to download onto the switch.

   The following prompt is displayed:

   ```
   Local File Name:
   ```

8. Enter a name for the file for when it is stored in the switch's file system. The name can be up to 32 alphanumeric characters. Spaces are allowed, but special characters (e.g., /, \, #, and %) should not be used. The extension should correspond to its file type. Supported extensions and file types are listed in Table 3 on page 204.

The following message is displayed:

```
Getting the file from Remote TFTP Server - Please wait
...
```

9. If you have not already done so, start the TFTP server software.

   After downloading the system file, the switch displays the following message:

   ```
   File received successfully!
   ```

10. If you downloaded a new configuration file and want to make it the switch's active boot file, go to "Setting the Active Boot Configuration File" on page 179. If you downloaded a CA certificate and want to add it to the certificate database, refer to "Adding a Certificate to the Database" on page 722.

Uploading a System File

You use the procedures in this section to upload a system file from a switch's file system to a computer or TFTP server. Here are the system files you are most likely to upload from a switch:

- Boot configuration file
- Certificate enrollment request
- Public encryption key

You might, for instance, upload a switch's configuration file so that you can modify it with a text editor at your management workstation. Or, you might have created a CA certificate enrollment request on the switch and need to upload it prior to submitting it to a CA.

This section contains the following procedures:

- "Uploading a File from a Local Management Session" on page 210
- "Uploading a File from a Telnet Management Session" on page 213

Guidelines

This section contains the guidelines to uploading a file from the switch's file system.

These guidelines apply to both Xmodem and TFTP downloads.

- You can use either Xmodem or TFTP when uploading files from a local management session.
- You must use TFTP when uploading files from a Telnet management session.
- You cannot upload a private encryption key or encryption key pair. Key pairs have the file name extension ".ukf." (The prohibition against uploading an encryption key pair is to prevent an unauthorized individual from obtaining the private key.)
- To upload a public key, you must first export it from the key database into the switch's file system. For instructions, refer to "Exporting an Encryption Key" on page 701. Public keys have the file name extension ".key."

This guideline applies only to an Xmodem upload:

- Xmodem can upload a file only from the switch where you started the local management session. You cannot use Xmodem to upload a file from a switch accessed through enhanced stacking.

These guidelines apply to a TFTP upload:

- ❒ Your network must have a node with the TFTP server software.
- ❒ You should start the TFTP server software before beginning the download procedure.
- ❒ The switch must have an IP address and subnet mask, such as a master switch of an enhanced stack. For switches that do not have an IP address, such as slave switches, you can perform the upload from a local management session of the switch using Xmodem.

Uploading a File from a Local Management Session

Review “Guidelines” on page 209 before performing this procedure.

To upload a system file from a switch to a workstation or TFTP server from a local management session using Xmodem or TFTP, perform the following procedure:

1. Establish a local management session on the switch where you want to upload the system file.
2. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.
3. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.
4. For the System Utilities menu, type **2** to select Downloads and Uploads.

   The Downloads and Uploads menu is shown in Figure 51 on page 190.
5. Type **4** to select Upload a File.

   The following prompt is displayed:

   ```
   Upload Method/Protocol [X-Xmodem, T-TFTP]:
   ```
6. To upload a system file using Xmodem, go to Step 7. To upload a file using TFTP, do the following:

   a. Type **T**.

      The following prompt is displayed:

      ```
      TFTP Server IP address:
      ```

   b. Enter the IP address of the TFTP server.

The following prompt is displayed:

```
Remote File Name:
```

c. Enter a name for the file for when it is stored on the TFTP server. The extension should be the same as in the original file name (for example, ".cfg" for a configuration file and ".csr" for a CA certificate enrollment request).

The following message is displayed:

```
Local File Name:
```

d. Enter the name of the file in the switch's file system to upload to the TFTP server. You can specify only one file. You may not use wildcards in the filename.

The following message is displayed:

```
Sending the file to Remote TFTP Server - Please wait
...
```

The switch displays the following message when the upload process is finished:

```
File sent successfully!
```

The file is now stored on the TFTP server. This completes the procedure for uploading a file from the switch's file system from a local management session using TFTP.

7. To upload a file using Xmodem, type **X** at the prompt displayed in Step 5.

The following message is displayed:

```
Local File Name:
```

8. Enter the name of the file on the switch to upload to your computer. You can specify only one file. You may not use wildcards in the filename.

The following prompt is displayed:

```
You are going to invoke the Xmodem download utility.
Do you wish to continue? [Yes/No]

Note: Please select 1K Xmodem protocol for faster
download.
```

9. Type **Y** for Yes.

The following message is displayed:

```
Use Hyper Terminal's 'Transfer/Receive File' option to
select Protocol
```

```
Note: Please select '1K Xmodem' protocol for faster
upload...
```

10. Begin the file transfer.

> **Note**
> The transfer protocol must be Xmodem or 1K Xmodem.

Steps 11 through 14 illustrate how to upload a file with the Hilgraeve HyperTerminal program.

11. From the HyperTerminal main window, select **Receive File** from the **Transfer** pull-down menu, as shown in Figure 58.

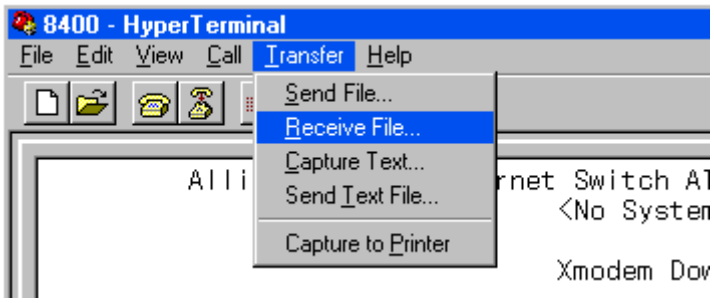

Figure 58. Local Management Window

The Receive File window is shown in Figure 59.

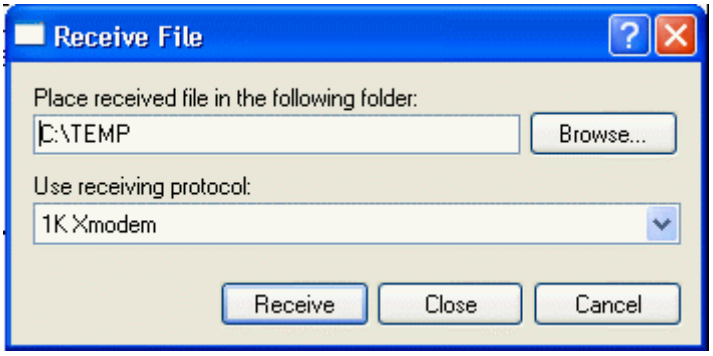

Figure 59. Receive File Window

12. Click **Browse** and specify a directory on your computer to store the file.

13. Click the **Use receiving protocol** field and select either Xmodem or, for a faster download, 1K XModem as the transfer protocol.

14. Click **Receive**.

15. When prompted, enter a name for the file for when it is stored on your workstation. The filename extension should be the same as in the original name (for example, ".cfg" for a configuration file and ".csr" for a CA certificate enrollment request).

    The switch uploads the file from the switch to your computer. This completes the procedure for uploading a file from the switch from a local management session using Xmodem.

Uploading a File from a Telnet Management Session

Review the "Guidelines" on page 209 before performing this procedure.

To upload a system file from the switch using a Telnet management session and TFTP, perform the following procedure:

1. Establish a Telnet management session on the switch containing the system file you want to upload to the TFTP server.

2. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

3. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

4. For the System Utilities menu, type **2** to select Downloads and Uploads.

   The Downloads and Uploads menu is shown in Figure 51 on page 190.

5. Type **4** to select Upload a File.

   The following prompt is displayed:

   ```
   Only TFTP uploads are available for a Telnet access

   TFTP Server IP address:
   ```

6. Enter the IP address of the TFTP server.

   The following prompt is displayed:

   ```
   Remote File Name:
   ```

7. Enter a name for the file for when it is stored on the TFTP server.

   The following message is displayed:

   ```
   Local File Name:
   ```

8. Enter the name of the file on the switch to upload to the TFTP server. You can specify only one file. You cannot use wildcards in the filename.

The following message is displayed:

```
Sending the file to Remote TFTP Server - Please wait ...
```

The switch displays the following message at the completion of the upload process:

```
File sent successfully!
```

The file is now stored on the TFTP server. This completes the procedure for uploading a file from a Telnet management session using TFTP.

Chapter 12

Event Log and Syslog Servers

This chapter describes how to view the event messages in the event log and how to configure the switch to send its event messages to a syslog server. Sections in the chapter include:

Event Log and Syslog Server Overview

A managed switch is a complex piece of computer equipment that includes both hardware and software components. Multiple software features operate simultaneously, interoperating with each other and processing large amounts of network traffic. It is often difficult to determine exactly what is happening when a switch appears not to be operating normally, or what happened when a problem occurs.

A network manager's major task is to monitor the network functions and to deal with problems as they arise. One method for monitoring a switch's activity is by viewing its event messages. These messages can help you identify and solve network problems by providing vital information about system and network activity on an AT-8500 Series switch. The information includes the time and date when an event occurred, the event's severity, the AT-S62 module that generated the event, and an event description.

There are two ways to view a switch's event messages. The first is by viewing the event log in the AT-8500 Series switch. The event log resides in temporary memory and has a maximum storage capacity of 4,000 events. The log is viewable from a local or remote management session of the switch. The log is not a permanent form of storage. All the events are purged whenever the switch is reset or power cycled. For instructions on how to view the log, refer to "Displaying the Event Log" on page 218.

The second method for viewing events is to have the switch send the event messages to a syslog server on your network using the syslog protocol. The advantage to this approach is that a syslog server can function as the central repository for event messages from many different network devices.

In order for a switch to send its events to a syslog server you have to create a syslog server definition. The definition includes the IP address of the syslog server where the messages are to be sent, and other information, such as the types of messages the switch is to send. You can create up to nineteen server definitions on a switch. For instructions on how to create a syslog server definition, refer to "Managing Syslog Server Definitions" on page 225.

Managing the Event Log

The following procedures explain how to view the events in the event log as well as how to enable and disable the log. Procedures include:

Enabling or Disabling the Event Log

This procedure explains how to enable or disable the event log on the switch. If you disable the log, the AT-S62 management software will not store events in its log and will not send events to any syslog servers you may have defined. The default setting for the event log is enabled.

The event log, even when disabled, will log all AT-S62 initialization events that occur whenever the switch is reset or power cycled. Any switch events that occur after AT-S62 initialization are entered into the log only if it is activated.

> **Note**
> Allied Telesyn recommends setting the switch's date and time if you activate the event log so that the messages will have the correct date and time when stored in the event log and sent to a syslog server. For instructions, refer to "Setting the System Time" on page 61.

To enable or disable the event log on a switch, do the following:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **8** to select Event Log.

The Event Log menu is shown in Figure 60.

```
  Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                     11:20:02 02-Jan-2006

                       Event Log
1 - Event Logging..............Enabled
2 - Display Output.............Temporary (Memory)
3 - Display Order..............Chronological
4 - Display Mode...............Normal
5 - Display Severity...........E,W,I
6 - Display Module.............All

C - Clear Log
L - Configure Log Outputs
S - Save Log to File
V - View Log
R - Return to Previous Menu

Enter your selection?
```

Figure 60. Event Log Menu

3. Type **1** to toggle Log Status between the two selections Enabled and Disabled. If you enable the log, the switch immediately begins to add events in the log and send events to defined syslog servers. The default is enabled.

   The other options in this menu are used to display the contents of the event log. For instructions, refer to “Displaying the Event Log” on page 218.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

   To display the events in the log, go to the next procedure. To configure the switch for a syslog server, refer to “Managing Syslog Server Definitions” on page 225.

Displaying the Event Log

This procedure explains how to view the event log from your local or remote management session. To view the event log, do the following:

1. From the Main Menu, type **5** to select System Administration.

2. From the System Administration menu, type **8** to select Event Log.

   The Event Log menu is shown in Figure 60 on page 218.

3. Configure options 2 through 6 in the Event Log menu to specify the types of events you want to view. The options are described below:

2 - Display Output
Selects an event log. This option has only one selection, Temporary. The event log is located in temporary memory.

3- Display Order
Controls the order of the events in the log. Choices are Chronological, which displays the events in the order oldest to newest, and Reverse Chronological, which displays the events newest to oldest. The default is Chronological.

4 - Display Mode
Controls the format of the event log. Choices are Normal, which displays the time, module, severity, and description for each event, and Full, which displays the same information as Normal, plus filename, line number, and event ID. The default is Normal. For an example of the display and definitions of the information, refer to Figure 61 on page 221.

5 - Display Severity
Displays events of a selected severity. Event severity is a predefined value assigned to an event according to its potential impact on switch operation. There are four severity levels, as defined in Table 4. The default is informational, error, and warning. You can specify more than one severity (for example, E,W).

Table 4. Event Log Severity Levels

Value	Severity Level	Description
ALL	-	Selects all severity levels
E	Error	Switch operation is severely impaired.
W	Warning	An issue may require manager attention.
I	Information	Useful information that can be ignored during normal operation.
D	Debug	Messages intended for Technical Support and Software Development.

6 - Display Module
Displays events of a selected AT-S62 module. The AT-S62 management software consists of a number of modules, each responsible for a different part of switch operation. You can instruct the switch to display only those events that apply to selected modules. The default is ALL, which displays the events for all modules. The modules are listed in Table 5.

Table 5. AT-S62 Modules

Module Name	Description
ALL	All modules
ACL	Access control list
CFG	Configuration files
CLASSIFIER	ACL and QoS policy classifiers
CLI	Command line interface commands
DOS	Denial of service defense
ENCO	Encryption keys
ESTACK	Enhanced stacking
EVTLOG	Event log
FILE	File system
GARP	GARP GVRP
HTTP	Web server
IGMPSNOOP	IGMP snooping
IP	Switch IP configuration, DHCP, and BOOTP
LACP	Link Aggregation Control Protocol
MAC	MAC address table
MGMTACL	Management access control list
PACCESS	802.1x port-based access control
PCFG	Port configuration
PKI	Public Key Infrastructure
PMIRR	Port mirroring
POE	Power over Ethernet (AT-8524POE switch only)
PSEC	Port security (MAC address-based)
PTRUNK	Port trunking
QOS	Quality of Service
RADIUS	RADIUS authentication protocol
SNMP	SNMP
SSH	Secure Shell protocol

Table 5. AT-S62 Modules

Module Name	Description
SSL	Secure Sockets Layer protocol
STP	Spanning Tree, Rapid Spanning, and Multiple Spanning Tree protocols
SYSTEM	Hardware status; Manager and Operator log in and log off events.
TACACS	TACACS+ authentication protocol
Telnet	Telnet
TFTP	TFTP
Time	SNTP
VLAN	Port-based and tagged VLANs, and multiple VLAN modes

4. Once you have set the log filters, type **V** to select View Log.

   Figure 61 shows an example of the event log in the Full display mode. The Normal display mode does not include the Filename, Line Number, and Event ID items.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
 User: Manager                                      11:20:02 02-Jan-2006
                               Event Log
S Date     Time      EventID    Source File:Line Number
                     Event
-----------------------------------------------------------------------
I 2/01/04  09:11:02  073001     garpmain.c:259
                     garp: GARP initialized
I 2/01/04  09:55:15  083001     portconfig.c:961
                     pcfg: PortConfig initialized
I 2/01/04  10:22:11  063001     vlanapp.c:444
                     vlan: VLAN initialization succeeded
I 2/01/04  12:24:12  093001     mirrorapp.c:158
                     pmirr: Mirror initialization succeeded
I 2/01/04  12:47:08  043016     macapp.c:1431
                     mac: Delete Dynamic MAC by Port[2] succeeded

Temporary (Memory) Log Events 1 - 5 of 212

P - Previous Page N - Next Page F - First Page L - Last Page
R - Return to Previous Menu

Enter your selection?
```

Figure 61. Event Log Example

The columns in the log are described below:

- S (Severity) - The event's severity. Table 4 defines the different severity levels.
- Date/Time - The date and time the event occurred.
- Event - The module within the AT-S62 software that generated the event followed by a brief description of the event. For a list of the AT-S62 modules, see Table 5 on page 220.
- Event ID - A unique number that identifies the event. (Displayed only in the Full display mode.)
- Filename and Line Number - The subpart of the AT-S62 module and the line number that generated the event. (Displayed only in the Full display mode.)

Modifying the Event Log Full Action

This procedure explains how to control what the log will do once it reaches its maximum capacity of 4,000 events. You have two options. The first is to have the switch delete the oldest entries as it adds new entries to the log. The second is to have the switch stop adding entries, so as to preserve the existing log contents.

This procedure is only relevant when viewing the event log through a local or remote management session. If you defined syslog servers, the switch continues to send events to a syslog server even when the log is full.

To configure the event log, do the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **8** to select Event Log.

   The Event Log menu is shown in Figure 60 on page 218.
3. From the Event Log menu, type **L** to select Configure Log Outputs.

The Configure Log Outputs menu is shown in Figure 62.

```
       Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                              11:20:02 02-Jan-2006

                     Configure Log Outputs
OutputID     Type           Status      Details
-------------------------------------------------------------
1            Temporary      Enabled     Wrap on Full

1 - Create Log Output
2 - Modify Log Output
3 - Delete Log Output
4 - View Log Output Details

R - Return to Previous Menu

Enter your selection?
```

Figure 62. Configure Log Outputs Menu

This menu includes any syslog servers you might have defined.

4. Type **2** to select Modify Log Output.

   The following prompt is displayed:

   ```
   Enter output ID to modify [1 to 20] -> 1
   ```

5. Press Return to accept the default value 1.

   The following prompt is displayed:

   ```
   Enter new log full action (1-Wrap on Full, 2-Halt on Full)
   ->
   ```

6. Type **1** if you want the switch to delete the oldest entries as it adds new entries, or **2** if the switch is to stop adding entries when the log reaches maximum capacity.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Saving the Event Log

The Event Log menu has the selection “S - Save Log to File” for saving the current contents of the log as a file in the file system. Once in the file system, you can either view it or download it to your management workstation.

Before selecting the option, configure options 2 to 7 in the Event Log menu to specify which log entries you want to save. When you select the option, you are asked to specify a filename. The name can be up to 16 alphanumeric characters, followed by the extension “.log”.

For instructions on the AT-S62 file system, refer to Chapter 10, File System.

Clearing the Event Log

To clear all events from the log, perform the following procedure:

1. From the Main menu, type **5** to select System Administration.
2. From the System Administration menu, type **8** to select Event Log.

   The Event Log menu is shown in Figure 60 on page 218.
3. Type **C** to select Clear Log.

   A confirmation prompt is displayed,
4. Type **Y** to clear the log or **N** to cancel the procedure.

   The log, if enabled, immediately begins to learn new events.

Managing Syslog Server Definitions

As explained at the start of this chapter, there are two ways to view the events generated by a switch. One approach is to view the switch's event log through a local or remote management session. The drawbacks to this approach are that you have to establish a management session with the switch before you can view the log and you can view the log of only one switch at a time.

Another way is to have the switch send its events to a syslog server. A syslog server can store the events of many network devices simultaneously. This can make managing your network easier since you can go to one site to see all of the events.

Configuring the switch to send its events to a syslog server involves creating a syslog server definition. The definition contains the IP address of the syslog server along with other information, such as what types of messages you want the switch to send.

Here are the guidelines to observe when using this feature:

- You can define up to 19 syslog servers.
- The event log on the switch must be activated in order for the switch to send events. For instructions, refer to "Enabling or Disabling the Event Log" on page 217.
- The switch must have an IP address and subnet mask. This rule applies to slave switches, which typically do not have an IP address, as well as master switches. If you want a slave switch to send its events to a syslog server, you must assign it an IP address and a subnet mask.
- The syslog server must communicate with the switch through the switch's management VLAN. The AT-S62 management software uses the management VLAN to watch for and transmit management packets. The default management VLAN is Default_VLAN. For further information, refer to "Specifying a Management VLAN" on page 579.

This section contains the following procedures:

Creating a Syslog Server Definition

To create a syslog server definition, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **8** to select Event Log.

   The Event Log menu is shown in Figure 60 on page 218.
3. From the Event Log menu, type **L** to select Configure Log Outputs.

   The Configure Log Outputs menu is shown in Figure 62 on page 223.
4. Type **1** to select Create Log Output.

   The following prompt is displayed:

   ```
   Enter output type (1-SYSLOG) ->
   ```
5. Type **1** to select the SYSLOG option. This is the only available option.

   The Syslog Server Configuration menu is shown in Figure 63.

```
        Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                              11:20:02 02-Jan-2006
                    Syslog Server Configuration
1 - Output ID ..................... <not defined>
2 - Server IP Address ............. 0.0.0.0
3 - Message Generation ............ Disabled
4 - Message Format ................ Extended
5 - Facility Level ................ DEFAULT
6 - Event Severity ................ E,W,I
7 - Event Module .................. All

C - Create Log Output
R - Return to Previous Menu

Enter your selection?
```

Figure 63. Syslog Server Configuration Menu

6. Configure the parameters as needed. The parameters are defined here:

   1 - Output ID
   The ID number for the syslog server definition. The definition will be identified in the Configure Log Outputs menu by this number. The range is 2 to 20. The default is the next available number. You cannot use a number that is already assigned.

   2 - Server IP Address
   The IP address of the syslog server.

3 - Message Generation
This enables and disables the syslog server definition. If set to disabled, which is the default, the switch does not send events to the syslog server. When enabled, the switch sends events. The default is disabled.

4 - Message Format
The information sent with each event. Choices are:

- Normal - sends the severity, module, and description.
- Extended - sends the same as Normal, plus the date, time, and switch's IP address. This is the default.

5 - Facility Level
The facility level to be added to the entries by the switch when it sends them to the syslog server. You can use the facility level to add a numerical code to the entries as they are transmitted to help you group entries on the syslog server according to the management module or switch that produced them. This can help you determine which entries belong to which units when a syslog server is collecting events from several difference network devices. You can specify only one facility level.

There are two approaches to using this parameter. The first is to use the 0-DEFAULT setting. At this setting, the code is based on the functional groupings defined in the RFC 3164 standard. The codes that are applicable to the AT-S62 management software and its modules are shown in Table 6.

Table 6. Applicable RFC 3164 Numerical Code and AT-S62 Module Mappings

Numerical Code	RFC 3164 Facility	AT-S62 Module
4	Security and authorization messages	Security modules: - PSEC - PACCESS - ENCO - PKI - SSH - SSL - MGMTACL - DOS Authentication modules: - SYSTEM - RADIUS - TACACS+

Table 6. Applicable RFC 3164 Numerical Code and AT-S62 Module Mappings

Numerical Code	RFC 3164 Facility	AT-S62 Module
9	Clock daemon	Time- based modules: - TIME (system time and SNTP) - RTC
22	Local use 6	Physical interface and data link modules: - PCFG - PMIRR - PTRUNK - STP - VLAN
23	Local use 7	SYSTEM events related to major exceptions.
16	Local use 0	All other modules and events.

For example, the setting of DEFAULT assigns all port mirroring events a code of 22 and all encryption key events a code of 4.

Your other option is to assign all events from a switch the same numerical code using one of the following facility level settings:

- 1 - LOCAL1
- 2 - LOCAL2
- 3 - LOCAL3
- 4 - LOCAL4
- 5 - LOCAL5
- 6 - LOCAL6
- 7 - LOCAL7

Each setting represents a predefined RFC 3164 numerical code. The code mappings are listed in Table 7.

Table 7. Numerical Code and Facility Level Mappings

Numerical Code	Facility Level Setting
17	LOCAL1
18	LOCAL2
19	LOCAL3

Table 7. Numerical Code and Facility Level Mappings

Numerical Code	Facility Level Setting
20	LOCAL4
21	LOCAL5
22	LOCAL6
23	LOCAL7

For example, selecting LOCAL2 as the facility level assigns the numerical code of 18 to all events sent to the syslog server by the switch.

6 - Event Severity
The severity of events to be sent by the switch to the syslog server. Event severity is a predefined value assigned to an event by the switch according to its potential impact on the switch's operation. You can use this parameter to configure the switch to send only those events that match one or more severity levels. There are four severity levels, as defined in Table 4 on page 219. The default is informational, error, and warning. You can specify more than one severity level (for example, E,W).

7 - Event Module
The originating module of the events to be sent to the syslog server. The AT-S62 management software consists of a number of modules, each responsible for a different part of switch operation. You can use this parameter to instruct the switch to send only those events that originated from selected modules. The default is ALL, which sends the events from all modules. The modules are defined in Table 5 on page 220. You can specify more than one module (for example, CLI,MAC,STP).

7. After you have configured the syslog server definition, type **C** to select Create Log Output.

   The switch immediately begins to send events to the server, if you enabled the definition when you created it, and adds the new syslog server definition to the Configure Log Outputs menu. An example of the menu with a syslog server definition is shown in Figure 64.

```
        Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                              11:20:02 02-Jan-2006

                     Configure Log Outputs
OutputIDType             Status       Details
-------------------------------------------------------------
1     Temporary          Enabled      Wrap on Full
2     Syslog             Enabled      149.44.44.44

1 - Create Log Output
2 - Modify Log Output
3 - Delete Log Output
4 - View Log Output Details

R - Return to Previous Menu

Enter your selection?
```

Figure 64. Configure Log Outputs Menu with a Syslog Server Definition

8. Repeat this procedure starting with step 4 to create additional syslog server definitions, if needed.
9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Syslog Server Definition

To modify a syslog server definition, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **8** to select Event Log.

   The Event Log menu is shown in Figure 60 on page 218.
3. From the Event Log menu, type **L** to select Configure Log Outputs.

   The Configure Log Outputs menu is shown in Figure 62 on page 223.
4. Type **2** to select Modify Log Output.

   The following prompt is displayed:

   ```
   Enter output ID to modify [1 to 20] -> 1
   ```
5. Enter the ID number of the syslog server definition you want to modify.

   The Syslog Server Configuration menu is shown in Figure 63 on page 226. The menu contains the specifications of the selected definition.
6. Modify the settings as needed.

For definitions of the parameters, refer to "Creating a Syslog Server Definition" on page 226. You cannot change a definition's output ID number.

7. When you are finished modifying the settings, type **M** to select Modify Log Output.

   The Configure Log Outputs menu is displayed again.

8. To modify additional definitions, repeat this procedure starting with step 4.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting a Syslog Server Definition

To delete a syslog server definition, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

2. From the System Administration menu, type **8** to select Event Log.

   The Event Log menu is shown in Figure 60 on page 218.

3. From the Event Log menu, type **L** to select Configure Log Outputs.

   The Configure Log Outputs menu is shown in Figure 62 on page 223.

4. Type **3** to select Delete Log Output.

   The following prompt is displayed:

   ```
   Enter output ID to delete [2 to 20] -> 2
   ```

5. Enter the ID number of the syslog server definition you want to modify. You can enter only one ID number at a time.

   The following confirmation prompt is displayed:

   ```
   Are you sure you want to delete output ID 2? [Yes/No] ->
   ```

6. Type **Y** to delete the definition or **N** for no to cancel the procedure.

   The definition is deleted from the Configure Log Outputs menu.

7. To delete additional definitions, repeat this procedure starting with step 4.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying a Syslog Server Definition

To display the details of an existing syslog server definition, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **8** to select Event Log.

   The Event Log menu is shown in Figure 60 on page 218.
3. From the Event Log menu, type **L** to select Configure Log Outputs.

   The Configure Log Outputs menu is shown in Figure 62 on page 223.
4. Type **4** to select View Log Output Details.

   The following prompt is displayed:

   ```
   Enter output ID to view [2 to 20] -> 2
   ```
5. Enter the ID number of the syslog server definition you want to view. You can enter only one ID number at a time.

   The Syslog Server Definition window is displayed with the specifications of the definition. For an explanation of the parameters, refer to “Creating a Syslog Server Definition” on page 226.

Chapter 13

Classifiers

This chapter explains classifiers and how you can create classifiers to define traffic flows. The sections in this chapter include:

Classifier Overview

A classifier defines a *traffic flow*. A traffic flow consists of packets that share one or more characteristics. A traffic flow can range from being very broad to very specific. An example of the former might be all IP traffic while an example of the latter could be packets with specific source and destination MAC addresses.

A classifier consists of a set of criteria. You configure the criteria to match the traffic flow you want the classifier to define. 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 will never use a classifier by itself.

There are two AT-S62 features that use classifiers. They are:

- Access control lists (ACL)
- Quality of Service (QoS) policies

As explained in Chapter 14 on page 251, an ACL filters ingress packets on a port by controlling which packets a port will accept and reject. You can use this feature to improve the security of your network or enhance network performance by creating network paths dedicated to carrying specific types of traffic.

When you create an ACL you must specify the traffic flow you want the ACL to control. You do that by creating one or more classifiers and adding the classifiers to the ACL. The action that the port takes when an ingress packet matches the traffic flow specified by a classifier is contained in the ACL itself. The action will be to either accept packets of the traffic flow or discard them.

The other feature that uses classifiers is QoS policies. You can use this feature to regulate the various traffic flows that pass through the switch. For instance, you might raise or lower the user priority value of a traffic flow or increase or decrease its allotted bandwidth.

As with an ACL, you specify the traffic flow of interest by creating one or more classifiers and applying them to a QoS policy. The action to be taken by a port when it receives a packet that corresponds to the prescribed flow

is dictated by the QoS policy, as explained in Chapter 15 on page 267.

In summary, a classifier is a list of variables that define a traffic flow. You apply a classifier to an ACL or a QoS policy to define the traffic flow you want the ACL or QoS policy to affect or control.

Classifier Criteria

The criteria of a classifier are defined in the following subsections.

Destination MAC Address (Layer 2)
Source MAC Address (Layer 2)

You can identify a traffic flow by specifying the 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. Each source and destination MAC address must be considered as a separate traffic flow, requiring its own classifier.

Ethernet 802.2 and Ethernet II Frame Types (Layer 2)

You can create a classifier that filters packets based on Ethernet frame type and on whether a packet is tagged or untagged within a particular 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 (Layer 2)

A tagged Ethernet frame, as explained in "Tagged VLAN Overview" on page 555, 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 65 illustrates the location of the user priority field within an Ethernet frame.

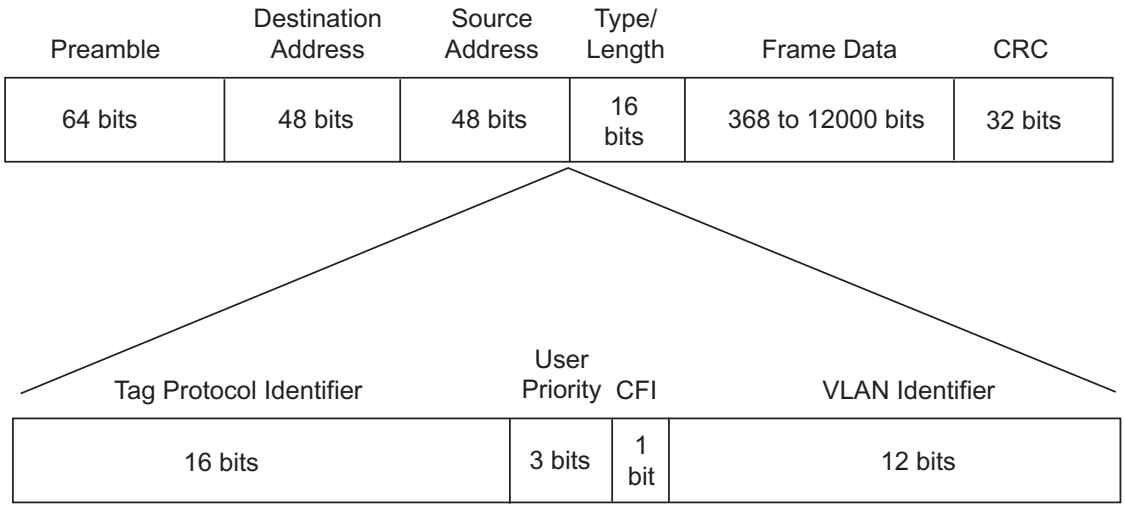

Figure 65. User Priority and VLAN Fields within an Ethernet Frame

You can identify a traffic flow of tagged packets using the user priority value. A classifier for such a traffic flow would instruct a port to watch for tagged packets containing the specified user priority level.

The priority level criterion 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 (Layer 2)

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 65, can be used to identify a traffic flow.

A classifier can contain only one VLAN ID. To create a port ACL or QoS policy that applies to several different VLAN IDs, multiple classifiers are required.

Protocol (Layer 2)

Traffic flows can be identified by any of the following Layer 2 protocols:

- IP
- ARP
- RARP
- Protocol Number

Observe the following guidelines when using this variable:

- This variable must be left blank or set to IP when setting a Layer3 or Layer 4 variable.
- To specify a protocol by its number, you can enter the value in decimal or hexadecimal format. If you choose hexadecimal, precede the number with the prefix "0x".

IP ToS (Type of Service) (Layer 3)

Type of Service (ToS) is a standard field in IP packets. It is used by applications to indicate the priority and Quality of Service for a frame. The range of the value is 0 to 7. The location of the field is shown in Figure 66.

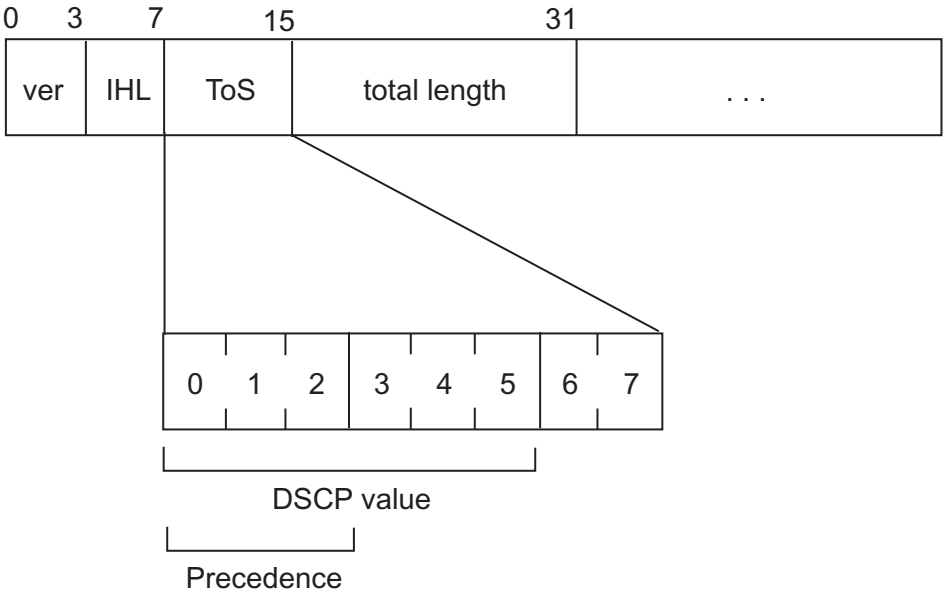

Figure 66. ToS field in an IP Header

Observe these guidelines when using this criterion:

- The Protocol variable must be left blank or set to IP.
- You cannot specify both an IP ToS value and an IP DSCP value in the same classifier.

IP DSCP (DiffServ Code Point) (ToS) (Layer 3)

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 66 on page 237. Routers within the network use this DSCP value to classify packets, and assign QoS appropriately. When a packet leaves the DiffServ domain, the DSCP value can be replaced with a value appropriate for the next DiffServ domain.The range of the value is 0 to 63.

Observe these guidelines when using this criterion:

- The Protocol variable must be left blank or set to IP.
- You cannot specify both an IP ToS value and an IP DSCP value in the same classifier.

IP Protocol (Layer 3)

You can define a traffic flow by the following Layer 3 protocols:

- TCP
- UDP
- ICMP
- IGMP
- IP protocol number

If you choose to specify a Layer 3 protocol by its number, you can enter the value in decimal or hexadecimal format. It you choose the latter, precede the number with the prefix “0x”.

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 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 Class C 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 if you are filtering on the IP address of a specific end node. A mask is required, however, when filtering on a subnet. As with a source IP mask, 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 Class C 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.

TCP Source Ports (Layer 4)
TCP Destination Ports (Layer 4)

Traffic flows can be identified by a source and/or destination TCP port number. A TCP port number is 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.
- The IP Protocol variable must be left blank or set to TCP.
- A classifier cannot contain criteria for both TCP and UDP ports. You may specify one in a classifier, but not both.

UDP Source Ports (Layer 4)
UDP Destination Ports (Layer 4)

Traffic flows can be identified by a source and/or destination UDP port number. A UDP port number is 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.
- The IP Protocol variable must be left blank or set to UDP.
- A classifier cannot contain criteria for both TCP and UDP ports. You may specify only one in a classifier.

TCP Flags

A traffic flow can be based on the following TCP flags:

- URG - Urgent
- ACK - Acknowledgement
- RST - Reset
- PSH - Push
- SYN - Synchronization
- FIN - Finish

Observe the following guidelines when using this criterion:

- The Protocol variable must be left blank or set to IP.
- The IP Protocol variable must be left blank or set to TCP.
- A classifier cannot contain both a TCP flag and a UDP source and/or destination port.

Classifier Guidelines

Here are the guidelines to follow when creating a classifier:

- Each classifier represents a separate traffic flow.
- The variables within a classifier are linked by AND. The more variables specified within a classifier, the more specific it becomes in terms of the defined flow. 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. There are restrictions on which variables can be used together in the same classifier. For the restrictions, refer to “Classifier Criteria” on page 235.
- You can apply the same classifier to more than one ACL or QoS policy.
- A classifier can be used for both ACLs and QoS policies.
- 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 access control lists and QoS policies at any one time will be from 14 to 127. The number depends on several factors, such as the number of ports to which the classifiers are assigned and the types of criteria defined in the classifiers.
- You cannot modify a classifier if it belongs to an ACL or QoS policy that is assigned to a port. You must first remove the port assignments from the ACL or policy and reassign them after you modify the classifier.
- You cannot delete a classifier if it is assigned to an ACL or QoS policy. You must first remove a classifier from its ACLs and QoS policies before you can delete it.

Creating a Classifier

This section contains the procedure for creating a classifier. As explained in "Classifier Overview" on page 234, a classifier is a series of variables that you set to define a traffic flow.

To create a classifier, do the following:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **1** to select Classifier Configuration.

   The Classifier Configuration menu is shown in Figure 67.

```
    Allied Telesyn Ethernet Switch AT-8524M - ATS62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
               Classifier Configuration
1 - Create Classifier
2 - Modify Classifier
3 - Destroy Classifier
4 - Show Classifiers

P - Purge Classifiers
R - Return to Previous Menu

Enter your selection?
```

Figure 67. Classifier Configuration Menu

3. Type **1** to select Create Classifier.

The Create Classifier menu (page 1) is shown in Figure 68.

```
    Allied Telesyn Ethernet Switch AT-8524M - ATS62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
                  Create Classifier
01 - Classifier ID: . 2
02 - Description: ...
03 - Dst MAC: .......
04 - Src MAC: .......
05 - Eth Format .....
06 - Priority: ......
07 - VLAN ID: .......
08 - Protocol: ......
09 - IP ToS: ........
10 - IP DSCP: .......

E - Edit Parameters
C - Create Classifier
N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 68. Create Classifier Menu (Page 1)

This is the first page of the classifier variables. To view the remaining variables, type **N** to select Next Page. The Create Classifier menu (page 2) is shown in Figure 69.

```
    Allied Telesyn Ethernet Switch AT-8524M - ATS62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
                  Create Classifier
11 - IP Protocol: ...
12 - Src IP Addr: ...
13 - Src IP Mask: ...
14 - Dst IP Addr: ...
15 - Dst IP Mask: ...
16 - TCP Src Port: ..
17 - TCP Dst Port: ..
18 - UDP Src Port: ..
19 - UDP Dst Port: ..
20 - TCP Flags: .....

E - Edit Parameters
C - Create Classifier
P - Previous Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 69. Create Classifier Menu (Page 2)

4. To set a variable, type **E** to select Edit Parameters.

   The following prompt is displayed.

   ```
   Enter parameter ID to edit: [1 to 19] ->1
   ```

5. Enter the number of the variable you want to configure. You can configure only one parameter at a time.

6. Adjust the new value for the variable.

   Refer to "Classifier Overview" on page 234 for definitions of the variables.

   Note
   Option 1 is used to assign the classifier an ID number. Each classifier must have a unique number. The range is 1 to 9999. The default is the lowest available number.

   Option 2 is used to assign a description to a classifier. You should assign all your classifiers a description. They can help you identify the different classifiers on the switch. A description can be up to 31 alphanumeric characters. Spaces are allowed. An example might be "IP traffic flow".

7. Repeat steps 5 and 6 to adjust any other variables necessary to define the traffic flow for this classifier.

8. After configuring the necessary variables, type **C** to select Create Classifier.

   The switch creates the classifier. If any of the settings are incompatible, the system displays an error message. Refer to the variable definitions in "Classifier Criteria" on page 235 for assistance in resolving compatibility issues.

9. To create more classifiers, repeat this procedure starting with step 3.

10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

11. To add classifiers to an ACL, refer to "Creating an ACL" on page 259. To add classifiers to a QoS policy, refer to "Managing Flow Groups" on page 283.

Modifying a Classifier

In order to modify a classifier, you need to know its ID number. To view classifier ID numbers, refer to “Displaying Classifiers” on page 248.

You cannot modify a classifier if it belongs to an ACL or QoS policy that is assigned to a port. You must first remove the port assignments from the ACL or policy before you can modify the classifier.

To modify a classifier, do the following:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **1** to select Classifier Configuration.

   The Classifier Configuration menu is shown in Figure 67 on page 241.
3. From the Classifier Configuration menu, type **2** to select Modify Classifier.

   The prompt similar to the following is displayed:

   ```
   Available Classifier(s): 1-11
   Enter Classifier ID :  [1 to 9999] -> 1
   ```
4. Enter the ID number of the classifier you want to modify.

   The Modify Classifier window is displayed. This window is identical to the Create Classifier menus, shown in Figure 68 on page 242 and Figure 69 on page 242.
5. Edit the variables as needed.

   Note the following when modifying a classifier:

   - ❒ You cannot change a classifier’s ID number.
   - ❒ To delete a value from a variable so as to leave it blank, select the criterion and then use the backspace key to delete its default value.
6. Once you have adjusted the variables, type **M** to select Modify Classifier.

   A change to a classifier is immediately activated. If any of the settings are incompatible, the system displays an error message. Refer to the variable definitions in “Classifier Criteria” on page 235 for assistance in resolving any compatibility issues.
7. To modify other classifiers, repeat this process starting with step 3.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

9. To add the modified classifier to an ACL, refer to "Creating an ACL" on page 259 or "Modifying an ACL" on page 261. To add it to a QoS policy, refer to "Managing Flow Groups" on page 283.

Deleting a Classifier

This procedure deletes a classifier from the switch. To delete a classifier, you need to know its ID number. To view the classifier ID numbers, refer to “Displaying Classifiers” on page 248.

You cannot delete a classifier if it belongs to an ACL or QoS policy. You must first remove a classifier from its ACL and QoS policy assignments before you can delete it.

To delete a classifier, do the following:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **1** to select Classifier Configuration.

   The Classifier Configuration menu is shown in Figure 67 on page 241.
3. From the Classifier Configuration menu, type **3** to select Destroy Classifier.

   The following prompt is displayed:

   ```
   Enter Classifier ID :  [1 to 9999] -> 1
   ```
4. Enter the ID number of the classifier to be deleted.

   The management software displays the details of the selected classifier. Use this window to verify that you selected the correct classifier.
5. If this is the correct classifier, type **D** to select Destroy Classifier.

   The classifier is deleted from the switch.
6. To delete additional classifiers, repeat this procedure starting with step 3.
7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting All Classifiers

This procedure deletes all classifiers from the switch. To delete individual classifiers, perform "Deleting a Classifier" on page 246.

You cannot delete the classifiers if any of them belong to an ACL or QoS policy. All classifiers must be removed from their ACL and QoS policy assignments before you can delete them.

To delete all classifiers, do the following:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **1** to select Classifier Configuration.

   The Classifier Configuration menu is shown in Figure 67 on page 241.
3. From the Classifier Configuration menu, type **P** to select Purge Classifiers.

   Caution
   No confirmation prompt is displayed. All classifiers are immediately deleted from the switch.
4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying Classifiers

To display the classifiers on a switch, do the following:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **1** to select Classifier Configuration.

   The Classifier Configuration menu is shown in Figure 67 on page 241.
3. From the Classifier Configuration menu, type **4** to select Show Classifiers.

   An example of the Show Classifiers window is illustrated in Figure 70.

```
          Allied Telesyn Ethernet Switch AT-8524M - ATS62
                        Production Switch
User: Manager                                11:20:02 02-Jan-2006
                         Show Classifiers
Number of classifiers: 5

                              Number of      Number of
ID   Description              References     Active Associations
-------------------------------------------------------------
1    IP flow                  4              3
2    Dst149.11.11.0           1              1
3    TCP flow                 1              0
4    Src149.22.22.49          1              1
5    ToS 6                    2              2

D - Detail Classifier Display
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 70. Show Classifiers Menu

The columns in the window are defined here:

- ID - The ID number of the classifier.
- Description - A description of the classifier.
- Number of References - The number of current assignments of a classifier to ACLs and QoS policies. For instance, a value of 8 would mean that a classifier is assigned to 8 ACLs and/or QoS policies. This number includes both active and inactive ACLs and QoS policies. An active ACL or policy is assigned to a switch port while an inactive ACL or policy is not. If this number is 0 (zero), the classifier has not been assigned to any ACLs or policies.

- ❒ Number of Active Associations - The number of current assignments of a classifier to only active ACLs and QoS policy.

4. To view the details of a classifier, type **D** to select Detail Classifier Display.

   The following prompt is displayed:

   ```
   Enter Classifier ID :  [1 to 9999] -> 1
   ```

5. Enter the ID number of the classifier you want to display.

   The details of the specified classifier are displayed. For examples of the windows, refer to Figure 68 on page 242 and Figure 69 on page 242. For definitions of the variables, refer to "Classifier Criteria" on page 235.

Chapter 14

Access Control Lists

This chapter explains access control lists (ACL) and how you can use this feature to improve network security and performance. This chapter contains the following sections:

Access Control List (ACL) Overview

An ACL is a filter that controls the ingress packets on a port. You can use this feature to control which ingress packets a port will accept and which it will reject. Packets are filtered based on the criteria defined in the classifiers assigned to an ACL.

There are several benefits of this feature. One is that it can add to your network security. You can create ACLs to protect parts of a network from unauthorized access by allowing only permitted traffic to enter the ports of a switch.

You can also use ACLs to enhance network performance by creating data links dedicated to carrying specific types of traffic. This provides the permitted traffic a higher priority by virtue of having its own dedicated network path.

This feature can also be used to achieve load-balancing by creating dedicated links for different types or categories of traffic. This too can result in enhanced network performance by distributing different types of network traffic across multiple physical links.

Note
This feature is not related to the management ACL feature, described in Chapter 35, “Management Access Control List” on page 759. They perform different functions and are configured in different ways.

The heart of an ACL is a classifier. A classifier, as explained “Classifier Overview” on page 234, defines packets that share a common trait. Packets that share a trait are referred to as a traffic flow. A traffic flow can be very broad, such as all IP packets, or very specific, such as packets from a specific end node destined for another specific node. You specify the traffic using different criteria, such as source and destination MAC addresses or protocol.

When you create an ACL, you are asked to specify the classifier that defines the traffic flow you want to permit or deny on a port.

There are two kinds of ACLs based on the two actions that an ACL can perform. One is called a permit ACL. Packets that meet the criteria in a permit ACL are accepted by a port.

The second type of ACL is a deny ACL. This type of ACL will deny entry to packets that meet the criteria of its classifiers, unless the packet also meets the criteria of a permit ACL on the same port, in which case the packet is accepted. This is because a permit ACL overrides a deny ACL.

Here is an overview of how the process works.

1. When an ingress packet arrives on a port, the switch checks it against the criteria in the classifiers of all the ACLs, both permit and deny, assigned to that port.
2. If the packet matches the criteria of a permit ACL, the port immediately accepts it. The packet is accepted even if it matches a deny ACL on the same port because a permit ACL overrides a deny ACL.
3. If a packet meets the criteria of a deny ACL but not any permit ACLs on the port, then the packet is discarded.
4. Finally, if a packet does not meet the criteria of any ACLs on a port, it is accepted by the port.

Parts of an ACL

To create an ACL, you must provide the following information:

- Name - An ACL needs a name. The name should reflect the type of traffic flow the ACL will be filtering and, perhaps, also the action. An example might be "HTTPS flow - permit." The more specific the name, the easier it will be for you to identify the different ACLs.
- Action - An ACL can have one of two actions: permit or deny. An action of permit means that the ingress packets matching the criteria in the classifiers are to be accepted by the switch port. An action of deny means any ingress packets matching the criteria are to be discarded, unless the packets match a permit ACL on the port, in which case the packets are accepted.
- Classifiers - An ACL needs one or more classifiers to define the traffic flow whose packets you want the port to accept or reject. Each classifier defines a different traffic flow. An ACL can have more than one classifier to filter multiple traffic flows.
- Port Lists - You need to specify the ports to which an ACL is to be assigned.

Guidelines

Here are rules for ACLs:

- A port can have multiple permit and deny ACLs.
- An ACL must have at least one classifier.
- An ACL can be assigned to more than one switch port.
- An ACL filters ingress traffic, but not egress traffic.
- The action of a ACL can be either permit or deny. A permit ACL overrides a deny ACL on the same port.
- The order in which ACLs are assigned to a port is unimportant. An ingress packet is compared against all of a port's ACLs.

- A classifier can be assigned to multiple ACLs. However, a classifier cannot be assigned more than once to a port. Put another way, ACLs that have the same classifier cannot be assigned to the same port.
- The switch can store up to 64 ACLs.

Examples

This section contains several examples of ACLs.

In this example, port 4 is assigned a deny ACL for the subnet 149.11.11.0. This ACL prevents the port from accepting any traffic originating from that subnet. Since this is the only ACL applied to the port, all other traffic is accepted. As explained earlier, a port automatically accepts all packets that do not meet the criteria of the classifiers assigned to its ACLs.

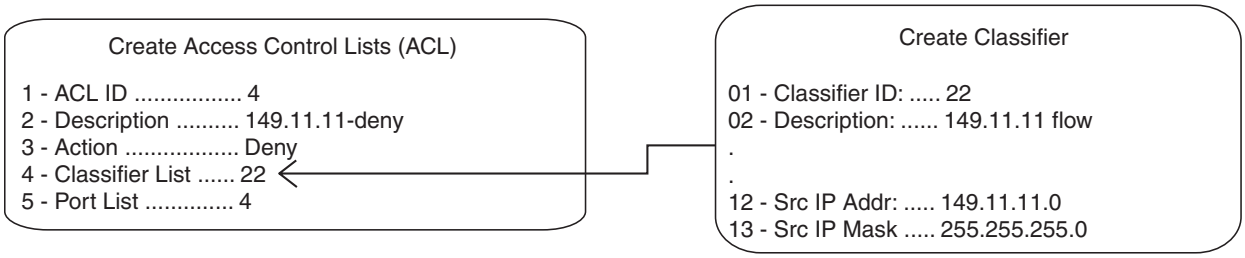

Figure 71. ACL Example 1

To deny traffic from several subnets on the same port, you can create multiple classifiers and apply them to the same ACL. This example denies traffic on port 4 from three subnets using three classifiers, one for each subnet, assigned to the same ACL.

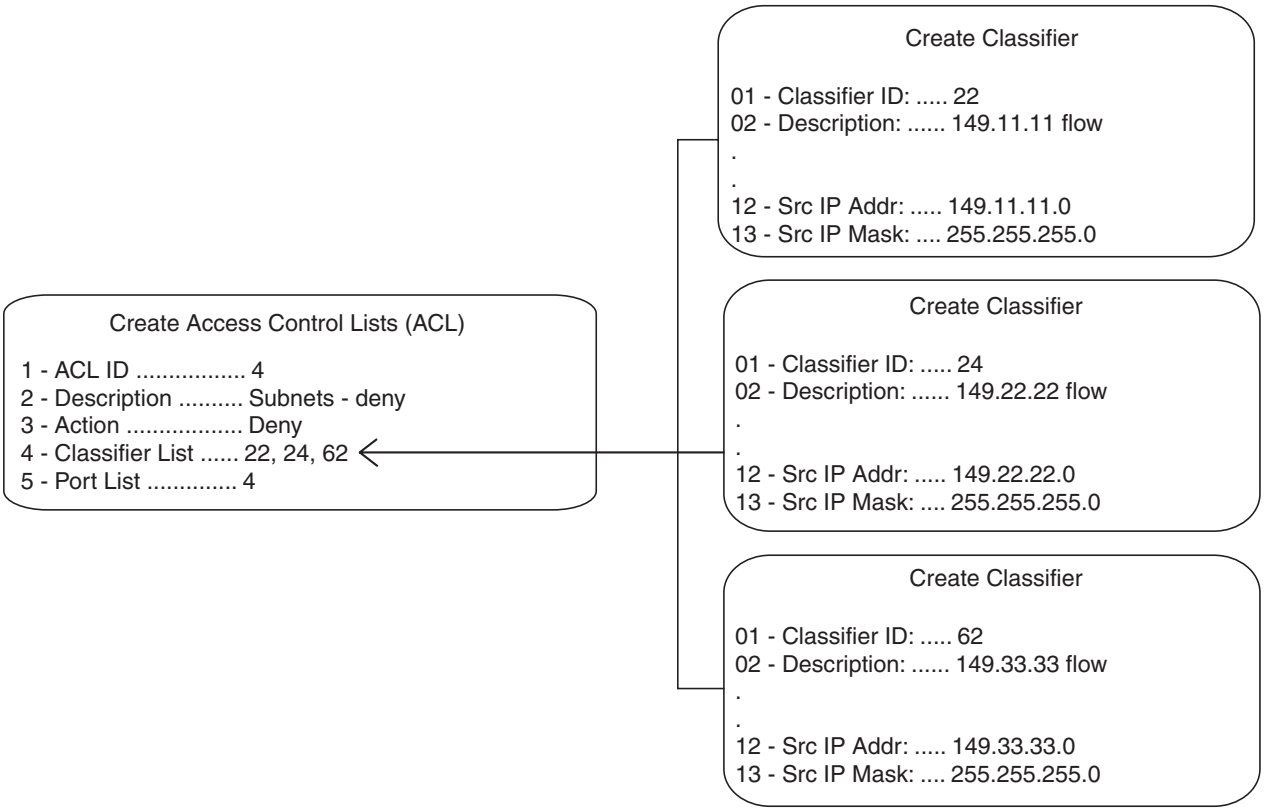

Figure 72. ACL Example 2

You can achieve the same result by assigning each classifier to a different ACL and assigning the ACLs to the same port, as in this example, again for port 4.

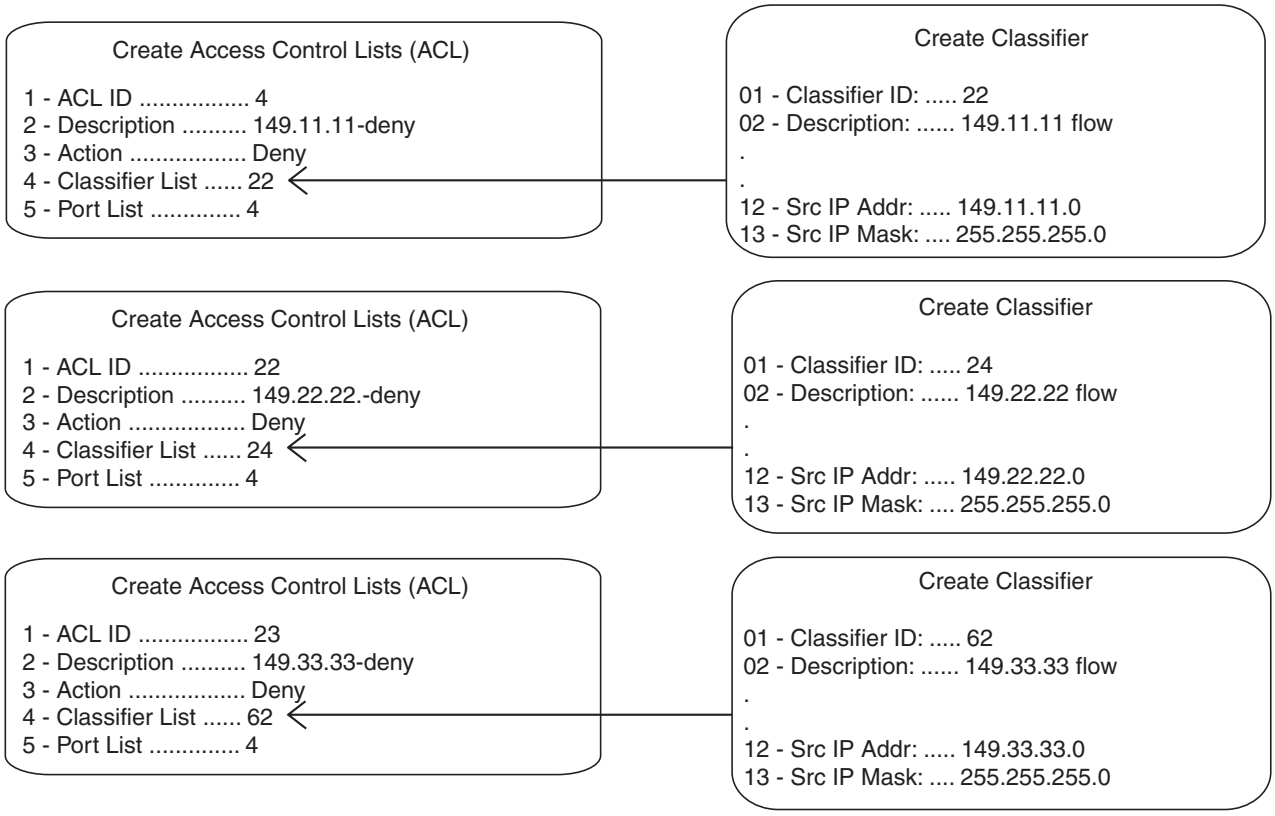

Figure 73. ACL Example 3

In this example, the traffic on ports 14 and 15 is restricted to packets from the source subnet 149.44.44.0. All other IP traffic is denied. Classifier ID 11 defines the authorized traffic flow for the ports and is assigned to an ACL with a permit action. Classifier ID 17 defines all IP traffic and is assigned to an ACL with a deny action. Since a permit ACL overrides a deny ACL, the ports accept the traffic from the 149.44.44.0 subnet even though that traffic meets the criteria of the deny ACL.

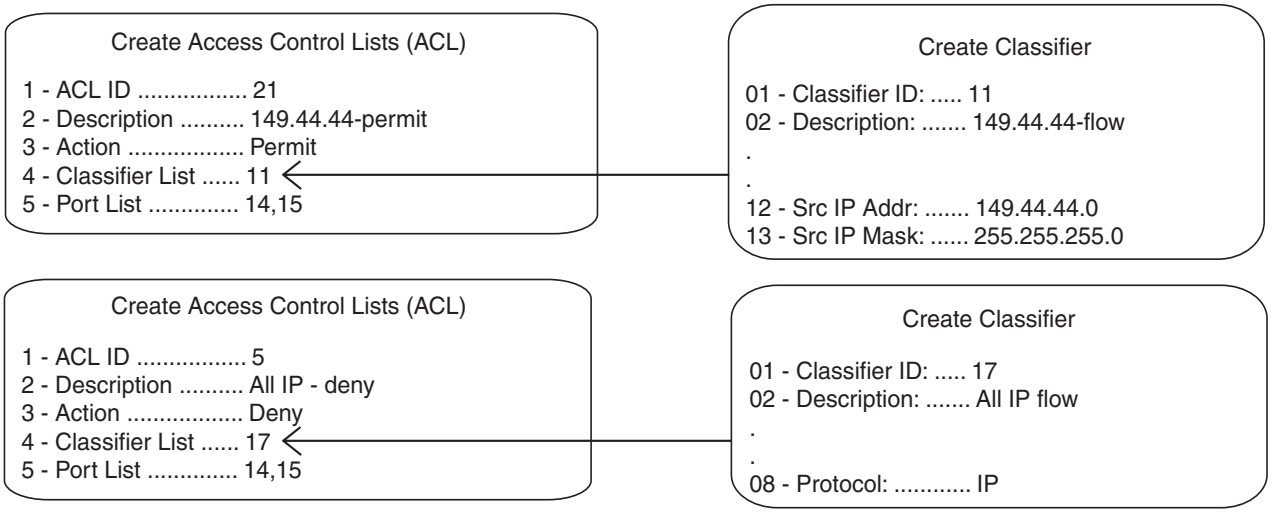

Figure 74. ACL Example 4

This example limits the traffic on port 22 to HTTPS web traffic directed to an end node with the IP address 149.55.55.55. All other IP traffic is rejected. (The Dst IP Mask field in classifier 6 is left empty because a mask is unnecessary when specifying a source or destination IP address of an end node. If you included the mask, it would be 255.255.255.255.)

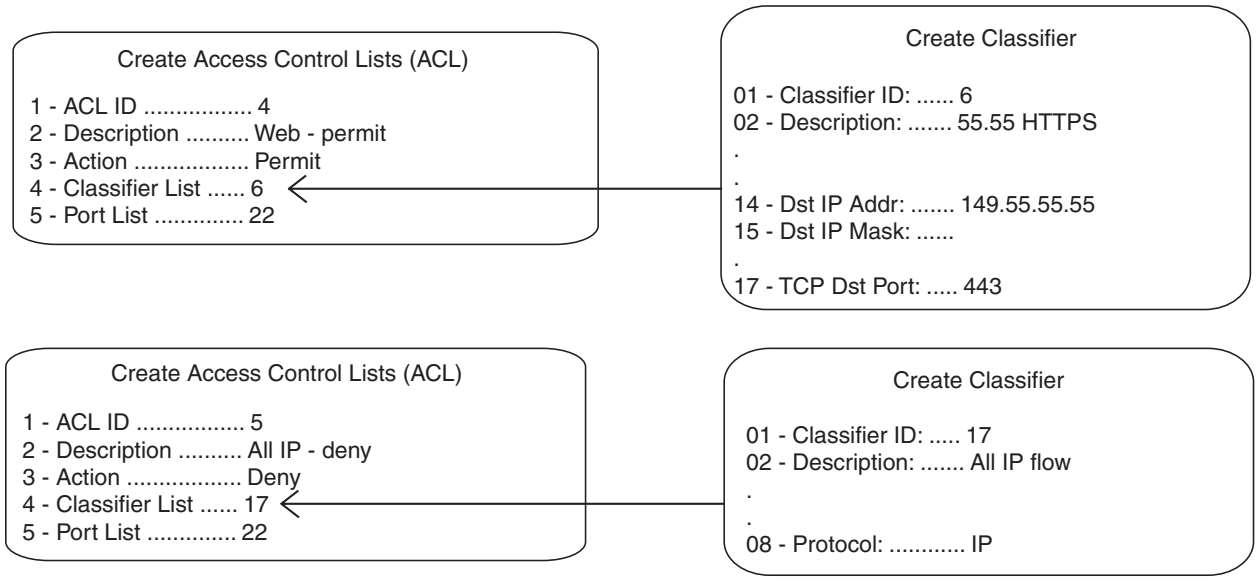

Figure 75. ACL Example 5

The next example limits the ingress traffic on port 17 to IP packets from the subnet 149.22.11.0 and a Type of Service setting of 6, destined to the end node with the IP address 149.22.22.22. All other IP traffic and ARP packets are prohibited.

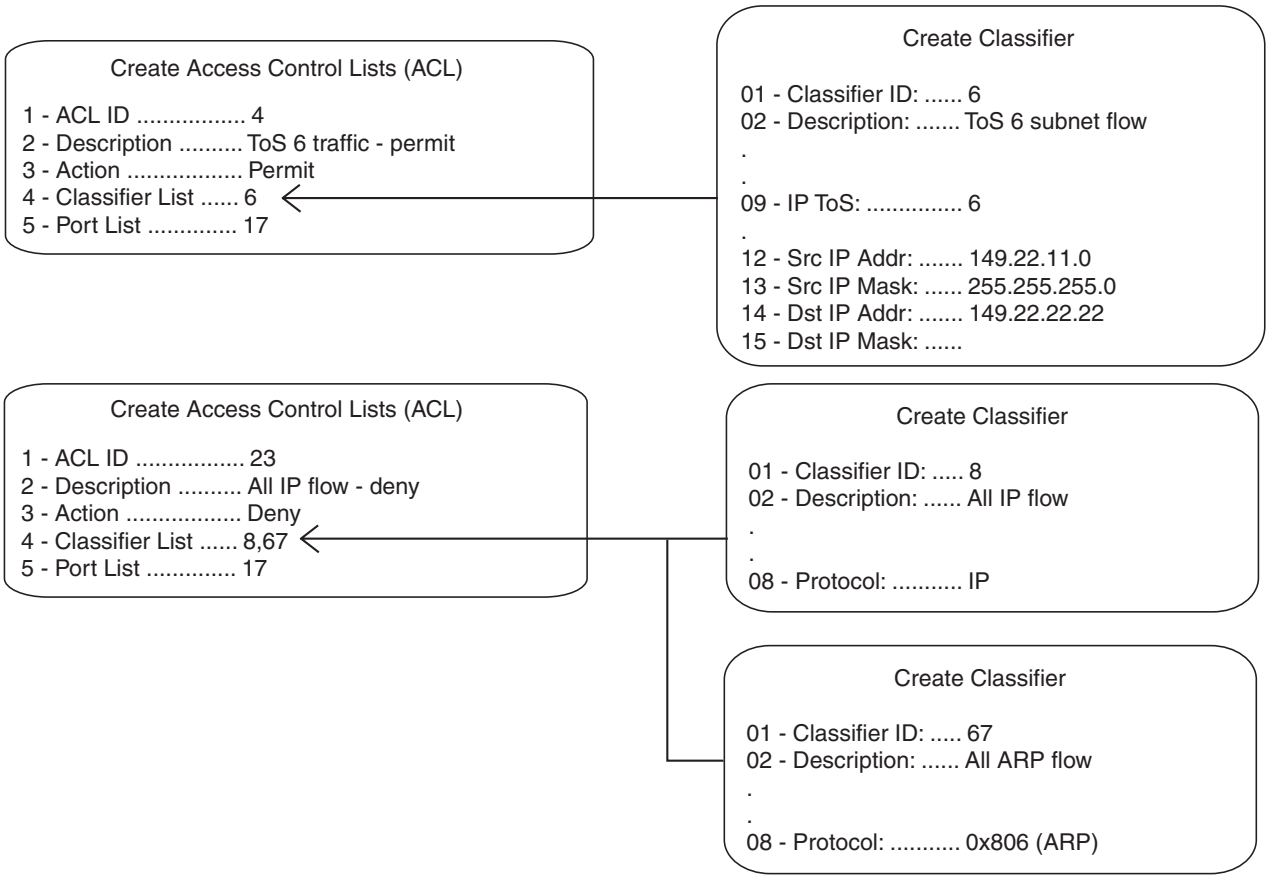

Figure 76. ACL Example 6

Creating an ACL

This procedure explains how to create an ACL. In order to perform this procedure, you need to know the ID numbers of the classifiers you want to assign to the ACL. To view classifier ID numbers, refer to "Displaying Classifiers" on page 248.

To create an ACL, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **4** to select Access Control Lists. The Access Control Lists (ACL) menu is shown in Figure 77.

```
  Allied Telesyn Ethernet Switch AT-8524M - ATS62
                Production Switch
User: Manager                  11:20:02 02-Jan-2006
            Access Control Lists (ACL)
1 - Create ACL
2 - Modify ACL
3 - Destroy ACL
4 - Show ACLs

P - Purge ACLs
R - Return to Previous Menu

Enter your selection?
```

Figure 77. Access Control Lists (ACL) Menu

3. Type **1** to select Create ACL. The Create ACL menu is shown in Figure 78.

```
  Allied Telesyn Ethernet Switch AT-8524M - ATS62
                Production Switch
User: Manager                 11:20:02 02-Jan-2006
                     Create ACL
1 - ACL ID ........... 0
2 - Description .......
3 - Action ............ Deny
4 - Classifier List ...
5 - Port List .........

C - Create ACL
R - Return to Previous Menu

Enter your selection?
```

Figure 78. Create ACL Menu

4. Type **1** to select ACL ID and, when prompted, enter an ID number for the ACL. Every ACL on the switch must have a unique ID number. The range is 0 to 255. The default is the lowest unused number. This parameter is required.

5. Type **2** to select Description and enter a description for the ACL. A description can be up to 31 alphanumeric characters. Spaces are allowed. This parameter is optional, though recommended. Assigning the ACLs different names will make it easier for you to identify them.

6. Type **3** to select Action.

   The following prompt is displayed:

   ```
   Enter Value [0 - Deny, 1 - Permit] : [0 to 1] -> 0
   ```

7. Type **0** if you want the ACL to discard ingress packets that meet the criteria in the classifiers to be assigned to the ACL or **1** if the packets are to be accepted. The default setting is Deny.

8. Type **4** to select Classifier List from the Create ACL menu and, when prompted, enter the classifiers to be assigned to the ACL. The prompt includes the ID numbers of the classifiers on the switch. You can assign more than one classifier to an ACL. Multiple classifiers are separated by a comma (for example, 4,7,2). The order in which you specify the classifiers is not important.

   When entering classifiers, keep in mind the action that you specified for this ACL in step 7. The action and the traffic flows defined by the classifiers should correspond. For instance, an ACL with an action of permit should be assigned those classifiers that define the traffic flow you want the ports to accept.

9. Type **5** to select Port List and, when prompted, enter the ports where you want to assign the ACL. You can assign an ACL to just one port or to more than one port. When entering multiple ports, the ports can be listed individually (e.g., 2,5,7), as a range (e.g., 8-12) or both (e.g., 1-4,6,8).

10. Type **C** to select Create ACL.

    The ACL is created on the switch and immediately activated on the specified ports.

11. To create additional ACLs, repeat this procedure starting with step 3.

12. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an ACL

This procedure explains how to modify an ACL. In order to perform this procedure, you need to know the ID number of the ACL you want to modify. To display ACL ID numbers, refer to "Displaying ACLs" on page 266. If you plan to add classifiers to the ACL, you also need to know the ID numbers of the classifiers. To view classifier ID numbers, refer to "Displaying Classifiers" on page 248.

To modify an ACL, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **4** to select Access Control Lists.

   The Access Control Lists (ACL) menu is shown in Figure 77 on page 259.
3. From the Access Control Lists (ACL) menu, type **2** to selection Modify ACL.

   The following prompt is displayed:

   ```
   Available ACL(s): 0-15
   Enter ACL ID :  [0 to 255] -> 0
   ```
4. Enter the ID number of the ACL you want to modify. You can modify only one ACL at a time.

   The Modify ACL window is displayed with the specifications of the selected ACL. An example of the window is shown in Figure 79.

```
  Allied Telesyn Ethernet Switch AT-8524M - ATS62
                 Production Switch
User: Manager                 11:20:02 02-Jan-2006
                     Modify ACL
1 - ACL ID ............ 12
2 - Description ....... HTTP - permit
3 - Action ............ Permit
4 - Classifier List ... 18,22
5 - Port List ......... 7,10-14

M - Modify ACL
R - Return to Previous Menu

Enter your selection?
```

Figure 79. Modify ACL Menu

You cannot change an ACL's ID number.

5. To change the description of the ACL, type **2** to select Description and enter a new description for the ACL. The description can be up to 31 alphanumeric characters. Spaces are allowed. This parameter is optional, though recommended. Assigning each ACL a name will make it easier for you to identify them.

6. To change the ACL's action, type **3** to select Action.

   The following prompt is displayed:

   ```
   Enter Value [0 - Deny, 1 - Permit] : [0 to 1] -> 0
   ```

7. Type **0** if you want the ACL to discard ingress packets that meet the criteria in the classifiers to be assigned to the ACL or **1** if the packets are to be accepted. The default setting is Deny.

8. To change the classifiers assigned to the ACL, type **4** to select Classifier List and, when prompted, enter the classifiers. The prompt includes the ID numbers of the classifiers on the switch. You can assign more than one classifier to an ACL. Multiple classifiers are separated by a comma (for example, 2,4,7). The order in which you specify the classifiers is not important.

   When entering classifiers, keep in mind the action you specified for this ACL in step 7. The action and the traffic flows defined by the classifiers should correspond. For instance, an ACL with an action of permit should be assigned those classifiers that define the traffic flow you want ports to accept.

9. To change the ports to which the ACL is assigned, type **5** to select Port List and, when prompted, enter the ports where you want to assign the ACL. You can assign an ACL to more than one port. Ports can be listed individually (e.g., 2,5,7), as a range (e.g., 8-12) or both (e.g., 1-4,6,8).

10. Type **M** to select Modify ACL.

    The ACL is modified on the switch. Modifications take affect immediately.

11. To modify additional ACLs, repeat this procedure starting with step 3.

12. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an ACL

This procedure deletes an ACL from the switch. In order to perform this procedure, you need to know the ID number of the ACL you want to delete. To display ACL ID numbers, refer to "Displaying ACLs" on page 266.

To delete an ACL, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **4** to select Access Control Lists.

   The Access Control Lists (ACL) menu is shown in Figure 77 on page 259.
3. From the Access Control Lists (ACL) menu, type **3** to selection Destroy ACL.

   The following prompt is displayed:

   ```
   Available ACL(s): 0-15
   Enter ACL ID :  [0 to 255] -> 0
   ```
4. Enter the ID number of the ACL you want to modify. You can modify only one ACL at a time.

   The Destroy ACL window is displayed with the specifications of the selected ACL. You can use this window to confirm that you are deleting the correct ACL. An example of the window is shown in Figure 80.

```
  Allied Telesyn Ethernet Switch AT-8524M - ATS62
                Production Switch
User: Manager                 11:20:02 02-Jan-2006
                    Destroy ACL
1 - ACL ID ............ 25
2 - Description ....... UDP-deny
3 - Action ............ Deny
4 - Classifier List ... 32
5 - Port List ......... 15,22

D - Destroy ACL
R - Return to Previous Menu

Enter your selection?
```

Figure 80. Destroy ACL Menu

5. To delete the ACL, type **D** to select Destroy ACL. To cancel the procedure, type **R** to select Return to Previous Menu.

   A deleted ACL is immediately removed from the switch.

6. To delete additional ACLs, repeat this procedure starting with step 3.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting All ACLs

This procedure deletes all ACLs from the switch.

To delete all ACLs, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **4** to select Access Control Lists.

   The Access Control Lists (ACL) menu is shown in Figure 77 on page 259.
3. From the Access Control Lists (ACL) menu, type **P** to selection Purge ACLs.

   Caution
   No confirmation prompt is displayed. All ACLs are immediately deleted from the switch.
4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying ACLs

To display the ACLs on a switch, perform this procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **4** to select Access Control Lists.

   The Access Control Lists (ACL) menu is shown in Figure 77 on page 259.
3. From the Access Control Lists (ACL) menu, type **4** to selection Show ACLs.

   An example of the Show ACLs window is illustrated in Figure 81.

```
  Allied Telesyn Ethernet Switch AT-8524M - ATS62
                Production Switch
User: Manager                 11:20:02 02-Jan-2006
                     Show ACLs

Number of ACLs: 12
ID  Description
-----------------------------------------------
1   IP - deny
2   HTTP - permit
3   TCP - deny
4   Src22.49 - deny
5   P-149.22.22.22
6   Dst22.50
7   ARP packets - deny

D - Detail ACL Display
N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 81. Show Classifiers Menu

4. To view the details of an ACL, type **D** to select Detail Classifier Display.

   The following prompt is displayed:

   ```
   Enter ACL ID :  [0 to 250] -> 0
   ```
5. Enter the ID number of the ACL you want to display.

   The details of the selected ACL are displayed.

Chapter 15

Quality of Service

This chapter describes Quality of Service (QoS). Sections in the chapter include:

Quality of Service Overview

Quality of Service allows you to prioritize traffic and/or limit the bandwidth available to it. The concept of QoS is a departure from the original networking protocols, which treated all traffic on the Internet or within a LAN the same. Without QoS, every different traffic type is equally likely to be dropped if a link becomes oversubscribed. This approach is now inadequate in many networks, because traffic levels have increased and networks transport time-critical applications such as streams of video data. QoS also enables service providers to easily supply different customers with different amounts of bandwidth.

Configuring Quality of Service involves two separate stages:

1. Classifying traffic into flows, according to a wide range of criteria.

   Classification is performed by the switch's packet classifiers, described in Chapter 13, "Classifiers" on page 233.

2. Acting on these traffic flows.

Quality of Service is a broadly used term that encompasses as a minimum both Layer 2 and Layer 3 in the OSI model. QoS is typically demonstrated by how the switch accomplishes the following:

- ❒ Assigns priority to incoming frames, if they do not carry priority information
- ❒ Maps prioritized frames to traffic classes, or maps frames to traffic classes based upon other criteria
- ❒ Maps traffic classes to egress queues, or maps prioritized frames to egress queues
- ❒ Provides maximum bandwidth limiting for traffic classes, egress queues and/or ports
- ❒ Schedules frames in egress queues for transmission (for example, empty queues in strict priority or sample each queue)
- ❒ Relabels the priority of frames
- ❒ Determines which frames to drop if the network becomes congested
- ❒ Reserves memory for switching/routing or QoS operation (e.g. reserving buffers for egress queues, or buffers to store packets with particular characteristics)

Note
QoS is only performed on packets which are switched at wirespeed. This includes IP, IP multicast, IPX, and Layer 2 traffic within VLANs.

The QoS functionality described by this chapter sorts packets into various flows, according to the QoS policy that applies to the port the traffic is received on. The switch then allocates resources to direct this traffic according to bandwidth or priority settings in the policy. Each policy is built up out of traffic classes, flow groups and classifiers. In summary, to configure QoS:

- Create classifiers to sort packets into traffic flows.
- Create flow groups and add classifiers to them. Flow groups are groups of classifiers which group together similar traffic flows. You can apply QoS prioritization to flow groups and/or replace the traffic's DiffServ Code Point.
- Create *traffic classes* and add flow groups to them. Traffic classes are groups of flow groups and are central to QoS. You can apply bandwidth limits and QoS prioritization to traffic classes, and/or replace the traffic's DiffServ Code Point.
- Create *policies* and add traffic classes to them. Policies are groups of traffic classes. A policy defines a complete QoS solution for a port or group of ports.
- Associate policies with ports.

Note
These steps are listed above in a conceptually logical order, but the switch cannot check a policy for errors until the policy is attached to a port. You can simplify error diagnosis by determining your QoS configuration on paper first, and then entering it into the switch starting with classifiers.

Policies, traffic classes, and flow groups are created as individual entities. When a traffic class is added to a policy, a logical link is created between the two entities. Destroying the policy will only unlink the traffic class, leaving the traffic class in an unassigned state. Destroying a policy will not destroy any of the underlying entities. Similarly, destroying traffic classes will simply unlink flow groups and destroying flow groups will simply unlink classifiers.

Classifiers

Classifiers are used to identify a particular traffic flow, and range from general to specific. (See Chapter 13, "Classifiers" on page 233 for more information.) Note that a single classifier should not be used in different flows that will end up, via traffic classes, assigned to the same policy. A classifier should only be used once per policy. Traffic is matched in the order of classifiers. For example, if a flow group has classifiers 1, 3, 2 and 5, that is the order in which the packets are matched.

Flow Groups

Flow groups are used to group similar traffic flows together, and allow more specific QoS controls to be used, in preference to those specified by the traffic class. Flow groups consist of a small set of QoS parameters and a group of classifiers. Once a flow group has been added to a traffic class it cannot be added to another traffic class. A traffic class may have many flow groups. Traffic is matched in the order of the flow groups. For example, if a traffic class has flow groups 1, 3, 2 and 5, this is the order in which the packets are matched.

QoS controls at the flow group level provide a QoS hierarchy. Non-default flow group settings are always used, but if no setting is specified for a flow group, the flow group uses the settings for the traffic class to which it belongs. For example, you can use a traffic class to limit the bandwidth available to web and FTP traffic combined. Within that traffic class, you can create two different flow groups with different priorities, to give web traffic a higher priority than FTP. Web traffic would then be given preferential access to bandwidth, but would be limited to the bandwidth limit of the traffic class.

Traffic Classes

Traffic classes are the central component of the QoS solution. They provide most of the QoS controls that allow a QoS solution to be deployed. A traffic class can be assigned to only one policy. Once assigned, it cannot be used by any other policies. Traffic classes consist of a set of QoS parameters and a group of QoS *Flow Groups*. Traffic can be prioritized, marked (IP TOS or DSCP field set), and bandwidth limited. Traffic is matched in the order of traffic class. For example, if a policy has traffic classes 1, 3, 2 and 5, this is the order in which the packets are matched.

Policies

QoS policies consist of a collection of user defined traffic classes. A policy can be assigned to more than one port, but a port may only have one policy.

Note that the switch can only perform error checking of parameters and parameter values for the policy and its traffic classes and flow groups when the policy is set on a port.

QoS controls are applied to ingress traffic on ports. Therefore, to control a particular type of traffic, an appropriate QoS policy must be attached to each port that type of traffic ingresses.

Although a policy can be applied to an egress port, the classifiers and the QoS controls are actually applied by the switch on the ingress ports of the traffic. This means the parameters used to classify the traffic and the actions specified by the policy are checked and applied on the ingress traffic of every port, before the traffic reaches an egress queue. As a consequence, a policy is never applied to the whole aggregate traffic of a designated egress port, but rather to the individual ingress flows destined to the port.

The effects of this behavior become evident when using the maximum bandwidth feature of QoS. Here is an example. Suppose you have a policy that assigns 5 Mbps of maximum bandwidth to an egress port. Now assume there are 10 ports on the switch where ingress traffic matches the criteria specified in the classifier assigned to the policy of the egress port. Since the policy considers each ingress flow separately, the result would be a maximum bandwidth of 50 Mbps (10 x 5 Mbps) on the egress port, because there are 10 flows, one from each ingress port, directed to the egress port.

An additional factor to consider when specifying an egress port in a policy is that if the destination MAC address of the traffic flow has not been learned by the egress port or, alternatively, added as a static address to the port, the policy remains inactive. This is because the ingress ports consider the traffic as unknown traffic and flood the traffic to all the ports. This applies equally to unknown unicast and unknown multicast traffic, as well as broadcast traffic.

QoS Policy Guidelines

- A classifier may be assigned to many flow groups. However, assigning a classifier more than once within the same policy may lead to undesirable results. A classifier may be used successfully in many different policies.
- A flow group must be assigned at least one classifier but may have many classifiers.
- A flow group may only be assigned to one traffic class.
- A traffic class may have many flow groups.
- A traffic class may only be assigned to one policy.
- A policy may have many traffic classes.
- A policy may be assigned to many ports.
- A port may only have one policy.
- You can create a policy without assigning it to a port, but the policy will be inactive.
- A policy must have at least one action defined in the flow group, traffic class, or the policy itself. A policy without an action is invalid.
- The switch can store up to 64 flow groups.
- The switch can store up to 64 traffic classes.
- The switch can store up to 64 policies.

Packet Processing

The switch's QoS tools can be used to perform any combination of the following functions on a packet flow:

- Limiting bandwidth
- Prioritizing packets, to determine the level of precedence the switch will give to the packet for processing

- Replacing the VLAN tag User Priority, to enable the next switch in the network to process the packet correctly
- Replacing the TOS precedence or DSCP value, to enable the next switch in the network to process the packet correctly.

Bandwidth Allocation

Bandwidth limiting is configured at the level of traffic classes, and encompasses the flow groups contained in the traffic class. Traffic classes can be assigned maximum bandwidths, specified in kbps, Mbps or Gbps.

Packet Prioritization

The switch has four Class of Service (CoS) egress queues, numbered from 0 to 3. Queue 3 has the highest priority. When the switch becomes congested, it gives high priority queues precedence over lower-priority queues. When the switch has information about a packet's priority, it sends the packet to the appropriate queue. You can specify the queue where the switch sends traffic, how much precedence each queue has, and whether priority remapping is written into the packet's header for the next hop to use.

Prioritizing packets cannot improve your network's performance when bandwidth is sufficiently over-subscribed so that egress queues are always full. If one type of traffic is causing the congestion, you can limit its bandwidth. Other solutions in this situation are to increase bandwidth or decrease traffic.

You can set a packet's priority by configuring a priority in the flow group or traffic class to which the packet belongs. The packet is put in the appropriate CoS queue for that priority. If the flow group and traffic class do not include a priority, the switch can determine the priority from the VLAN tag User Priority field of incoming tagged packets. The packet is put in the appropriate CoS queue for its VLAN tag User Priority field. If neither the traffic class / flow group priority nor the VLAN tag User Priority is set, the packet is sent to the default queue, queue 1.

Both the VLAN tag User Priority and the traffic class / flow group priority setting allow eight different priority values (0-7). These eight priorities are mapped to the switch's four CoS queues. The switch's default mapping is shown in Table 8 on page 309. Note that priority 0 is mapped to CoS queue 1 instead of CoS queue 0 because tagged traffic that has never been prioritized has a VLAN tag User Priority of 0. If priority 0 was mapped to CoS queue 0, this default traffic goes to the lowest queue, which is probably undesirable. This mapping also makes it possible to give some traffic a lower priority than the default traffic.

Replacing Priorities

The traffic class or flow group priority (if set) determines the egress queue a packet is sent to when it egresses this switch, but by default has no effect on how the rest of the network processes the packet. To permanently change the packet's priority, you need to replace one of two priority fields in the packet header:

- The User Priority field of the VLAN tag header. Replacing this field relabels VLAN-tagged traffic, so that downstream switches can process it appropriately. Replacing this field is most useful outside DiffServ domains.
- The DSCP value of the IP header's TOS byte (Figure 66 on page 237). Replacing this field may be required as part of the configuration of a DiffServ domain. See "DiffServ Domains" on page 273 for information on using the QoS policy model and the DSCP value to configure a DiffServ domain.

VLAN Tag User Priorities

Within a flow group or traffic class, the VLAN tag User Priority value of incoming packets can be replaced with the priority specified in the flow group or traffic class. Replacement occurs before the packet is queued, so this priority also sets the queue priority.

DSCP Values

There are three methods of replacing the DSCP byte of an incoming packet. You can use these methods together or separately. They are described in the order in which the switch performs them.

1. The DSCP value can be overwritten at ingress, for all traffic in a policy.
2. The DSCP value in the packet can be replaced at the traffic class or flow group level.

   You can use these two replacements together at the edge of a DiffServ domain, to initialize incoming traffic.
3. The DSCP value in a flow of packets can replaced if the bandwidth allocated to that traffic class is exceeded, using the command. This option allows the next switch in the network to identify traffic that exceeded the bandwidth allocation.

DiffServ Domains

Differentiated Services (DiffServ) is a method of dividing IP traffic into classes of service, without requiring that every router in a network remember detailed information about traffic flows. DiffServ operates within a *DiffServ domain*, a network or subnet is managed as a single QoS unit. Packets are classified according to user-specified criteria at the edge of the network, divided into classes, and assigned the required class of service. Then packets are marked with a Differentiated Services Code Point (DSCP) tag to indicate the class of service to which they belong. The DSCP value is written into the TOS field of the IP header. Routers within the network then use this DSCP value to classify packets, and assign QoS appropriately. When a packet leaves the DiffServ domain, the DSCP value can be replaced with a value appropriate for the next DiffServ domain.

A simple example of this process is shown in Figure 82, for limiting the amount of bandwidth used by traffic from a particular IP address. In the domain shown, this bandwidth limit is supplied by the class of service represented by a DSCP value of 40. In the next DiffServ domain, this traffic is assigned to the class of service represented by a DSCP value of 3.

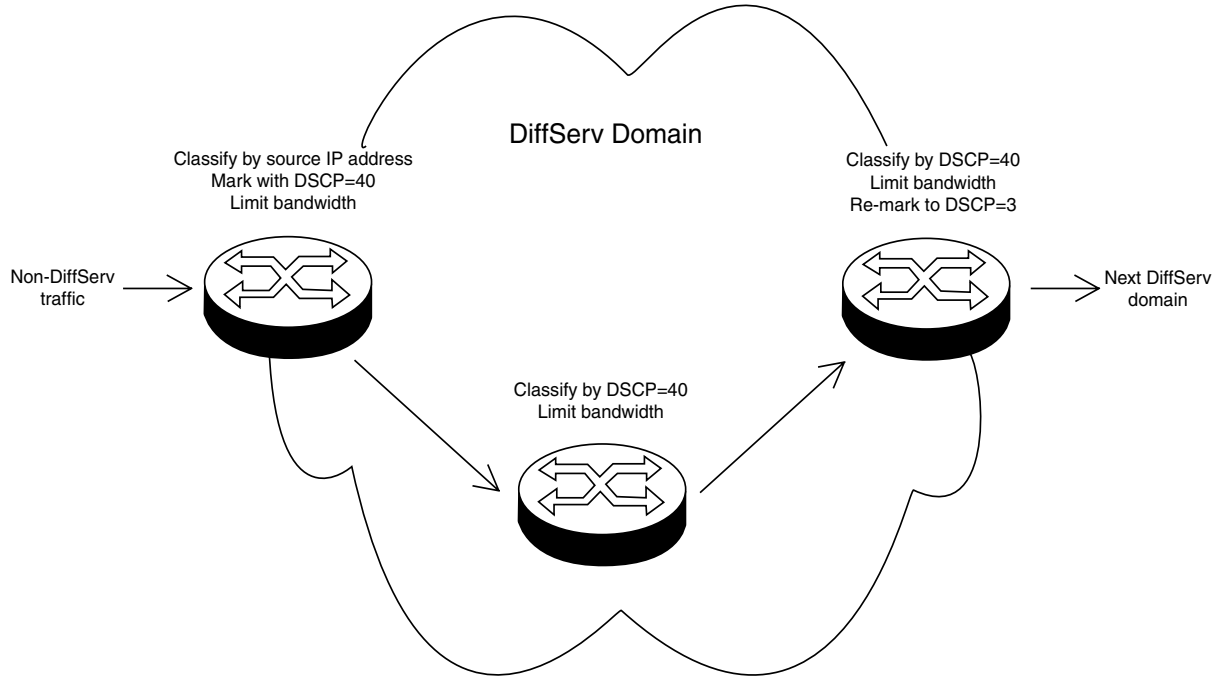

Figure 82. DiffServ Domain Example

To use the QoS tool set to configure a DiffServ domain:

1. As packets come into the domain at edge switches, replace their DSCP value, if required.

   Classify the packets according to the required characteristics. For available options, see Chapter 13, “Classifiers” on page 233.

   Assign the classifiers to flow groups and the flow groups to traffic classes, with a different traffic class for each DiffServ code point grouping within the DiffServ domain.

   Give each traffic class the priority and/or bandwidth limiting controls that are required for that type of packet within this part of the domain.

   Assign a DSCP value to each traffic class, to be written into the TOS field of the packet header.

2. On switches and routers within the DiffServ domain, classify packets according to the DSCP values that were assigned to traffic classes on the edge switches.

Assign the classifiers to flow groups and the flow groups to traffic classes, with a different traffic class for each DiffServ code point grouping within the DiffServ domain.

Give each traffic class the priority and/or bandwidth limiting controls that are required for that type of packet within this part of the domain. These QoS controls need not be the same for each switch.

3. As packets leave the DiffServ domain, classify them according to the DSCP values.

   Assign the classifiers to flow groups and the flow groups to traffic classes, with a different traffic class for each DiffServ code point grouping within the DiffServ domain.

   Give each traffic class the priority and/or bandwidth limiting controls required for transmission of that type of packet to its next destination, in accordance with any Service Level Agreement (SLA) with the providers of that destination.

   If necessary, assign a different DSCP value to each traffic class, to be written into the TOS field of the packet header, to match the DSCP or TOS priority values of the destination network.

Examples Voice Applications

Voice applications typically require a small bandwidth but it must be consistent. They are sensitive to *latency* (interpacket delay) and *jitter* (delivery delay). Voice applications can be set up to have the highest priority.

This example creates two policies that ensure low latency for all traffic sent by and destined to a voice application located on a node with the IP address 149.44.44.44. The policies raise the priority level of the packets to 7, the highest level. Policy 6 is for traffic from the application that enters the switch on port 1. Policy 11 is for traffic arriving on port 8 going to the application. The components of the policies are shown in Figure 83.

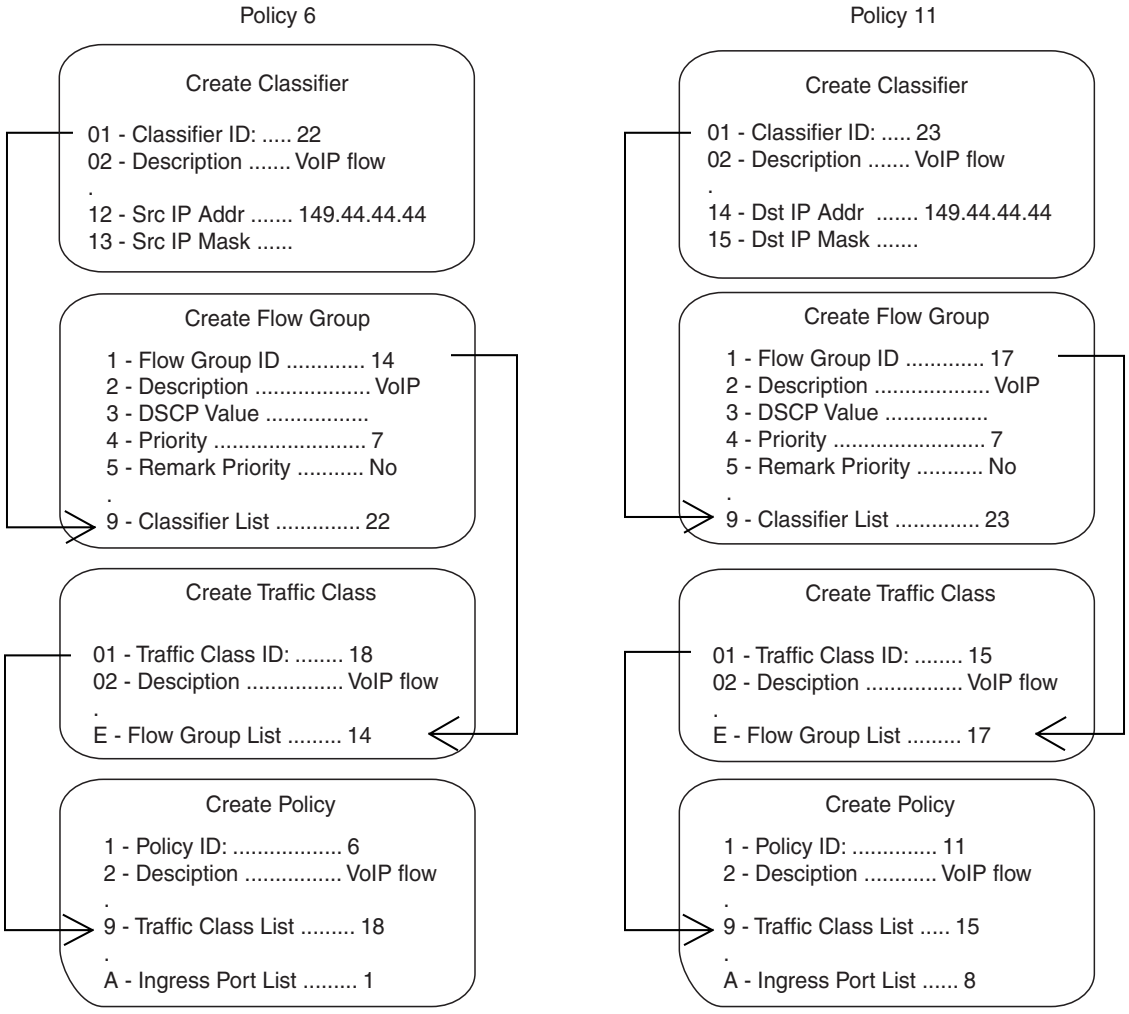

Figure 83. QoS Voice Application Example

The parts of the policies are:

- Classifier - Defines the traffic flow by specifying the IP address of the node with the voice application. The classifier for Policy 6 specifies the address as a source address since this classifier is part of a policy for packets coming from the application. The classifier for Policy 11 specifies the address as a destination address since this classifier is part of a policy for packets going to the application.
- Flow Group - Specifies the new priority level of 7 for the packets. It should be noted that in this example the packets leave the switch with the same priority level they had when they entered. The new priority level is relevant only as the packets traverse the switch. To alter the packets so that they leave containing the new level, you would change option 5, Remark Priority, to Yes.
- Traffic Class - No action is taken by the traffic class, other than to specify the flow group. Traffic class has a priority setting that can be used to override the priority level of packets, just as in a flow group. If you enter a priority value in both places, the setting in the flow group overrides the setting in the traffic class.
- Policy - Specifies the traffic class and the port to which the policy is to be assigned. Policy 6 is applied to port 1 since this is where the application is located. Policy 11 is applied to port 8 since this is where traffic going to the application will be received.

Video Applications

Video applications typically require a larger bandwidth than voice applications. Video applications can be set up to have a high priority and buffering, depending on the application.

This example creates policies with low latency and jitter for video streams (for example, net conference calls). The policies in Figure 84 assign the packets a priority level of 4 and limit the bandwidth to 5 Mbps. The node containing the application has the IP address 149.44.44.44. Policy 17 is assigned to port 1, where the application is located, and Policy 32 is assigned to port 8 where packets destined to the application enter the switch.

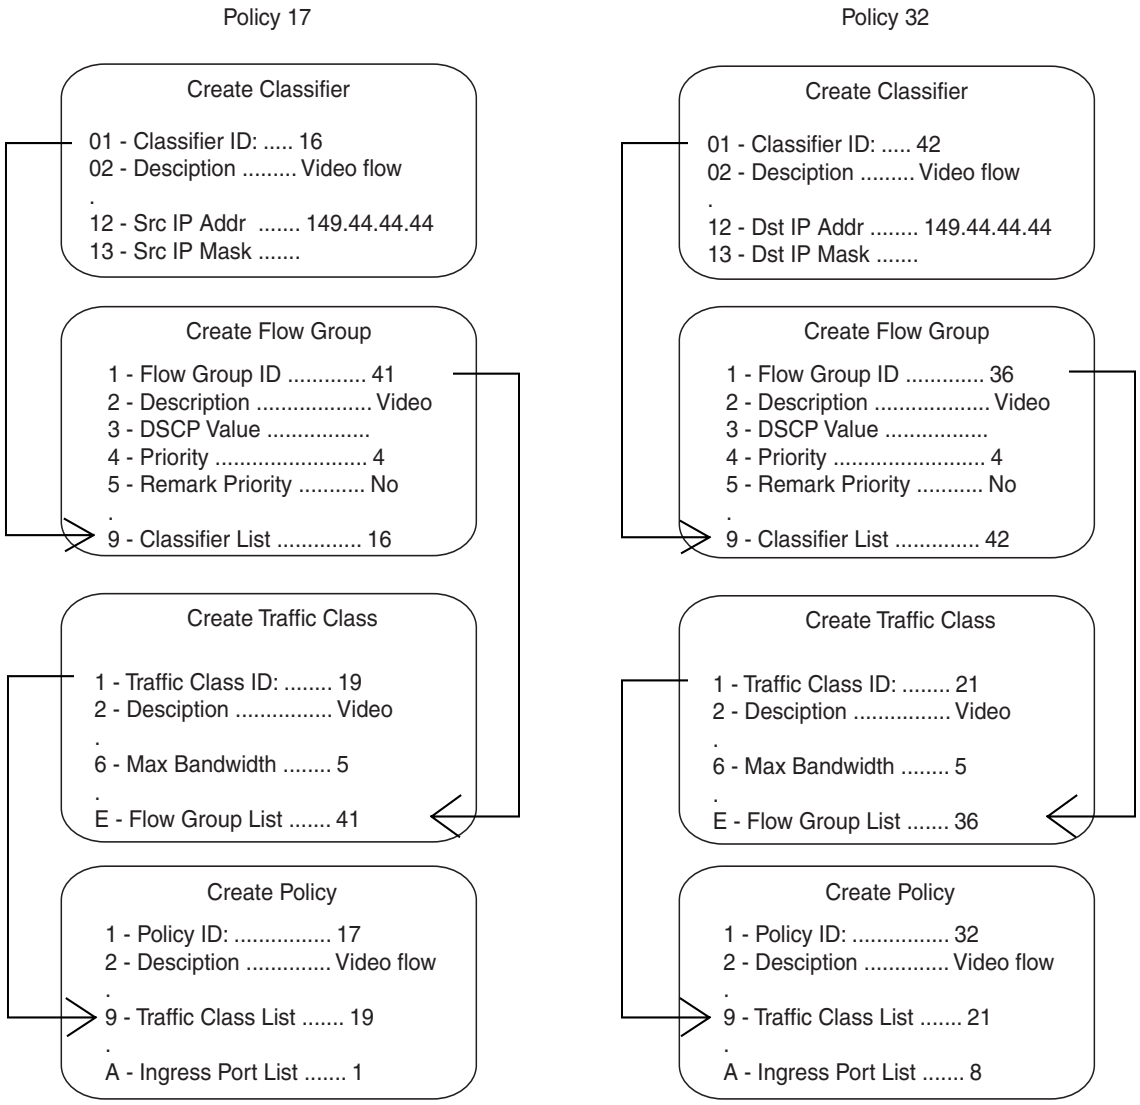

Figure 84. QoS Video Application Example

The parts of the policies are:

- Classifier - Specifies the IP address of the node with a video application. The classifier for Policy 17 specifies the address as a source address since this classifier is part of a policy concerning packets coming from the application. The classifier for Policy 32 specifies the address as a destination address since this classifier is part of a policy concerning packets going to the application.
- Flow Group - Specifies the new priority level of 4 for the packets. As with the previous example, the packets leave the switch with the same priority level they had when they entered. The new priority level is relevant only while the packets traverse the switch. To alter the packets so that they leave containing the new level, you would change option 5, Remark Priority, to Yes.
- Traffic Class - The packet stream is assigned a maximum bandwidth of 5 Mbps. Bandwidth assignment can only be made at the traffic class level.
- Policy - Specifies the traffic class and the port where the policy is to be assigned.

Critical Database

Critical databases typically require a high bandwidth. They also typically require less priority than either voice or video.

The policies in Figure 85 assign 50 Mbps bandwidth, with no change to priority, to traffic going to and from a database. The database is located on a node with the IP address 149.44.44.44 on port 1 of the switch.

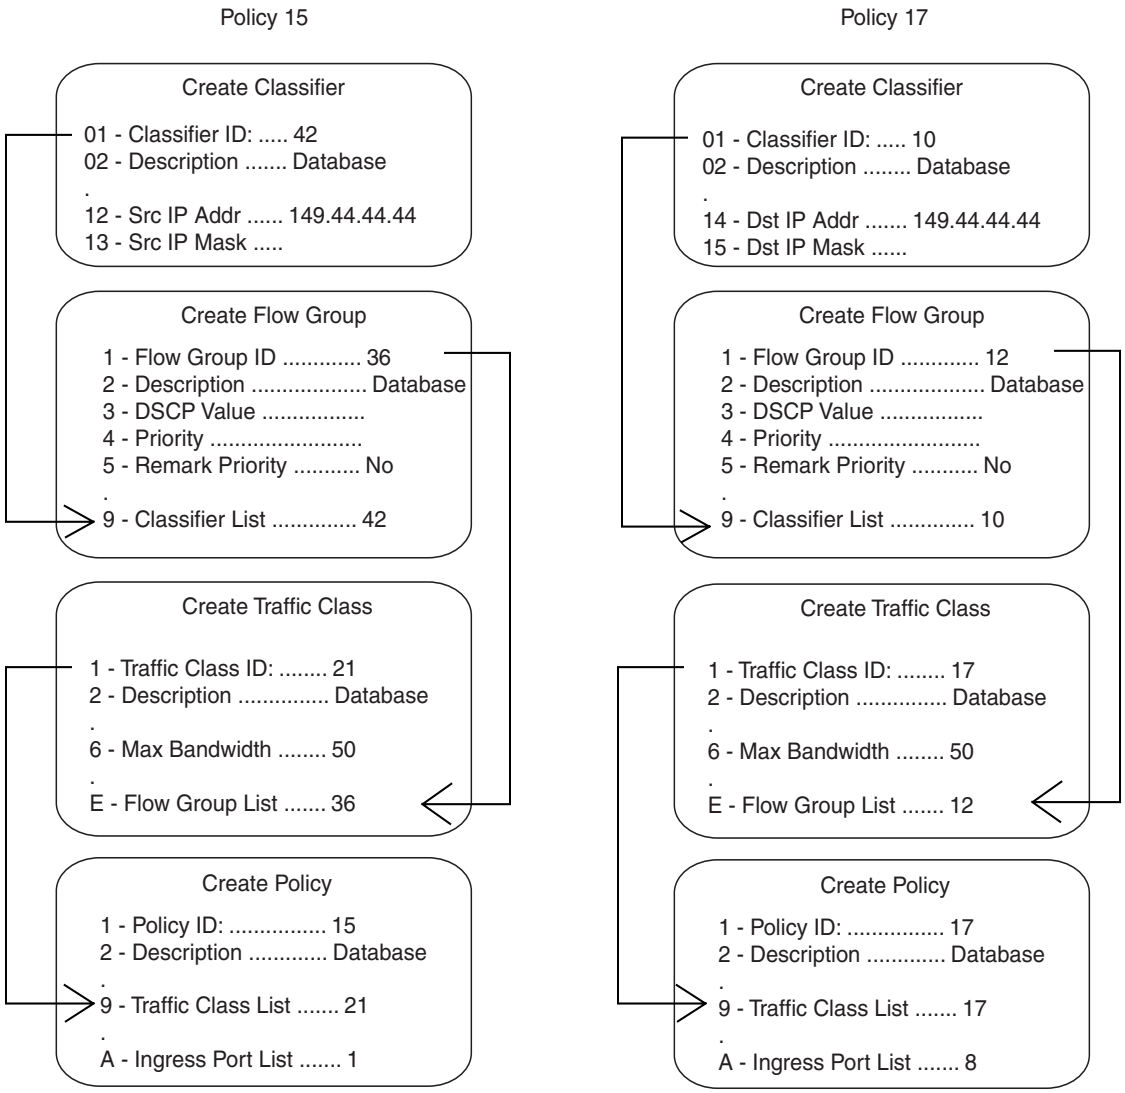

Figure 85. QoS Critical Database Example

Policy Component Hierarchy

The purpose of this example is to illustrate the hierarchy that exists among the components of a QoS policy and how that hierarchy needs to be taken into account when assigning new priority and DSCP values. A new priority can be set at the flow group and traffic class levels, while a new DSCP value can be set at all three levels -- flow group, traffic class and policy. The basic rules are:

- A new setting in a flow group takes precedence over a corresponding setting in a traffic class or policy.
- A new setting in a traffic class takes precedence over a corresponding setting in a policy.
- A new setting in a policy is used only if there is no corresponding setting in a flow group or traffic class.

This concept is illustrated in Figure 86 on page 282. It shows a policy for a series of traffic flows consisting of subnets defined by their destination IP addresses. New DSCP values for the traffic flows are established at different levels within the policy.

Traffic flows 149.11.11.0 and 149.22.22.0, defined by classifiers 1 and 2, are attached to a flow group, traffic class, and policy that contain new DSCP values. Since a setting in a flow group takes precedence over that of a traffic class or policy, the value in the flow group is used. The result is that the DSCP value in the two traffic flows is changed to 10.

The flow group for traffic flows 149.33.33.0 and 149.44.44.0, defined in classifiers 3 and 4, does not contain a new DSCP value. Consequently, the new value in the traffic class is used, in this case 30. The policy also has a DSCP setting, but it is not used for these traffic flows because a new DSCP setting in a traffic class takes precedence over that of a policy.

Finally, the new DSCP value for traffic flows 149.55.55.0 and 149.66.66.0, defined in classifiers 5 and 6, is set at the policy level to a value of 55 because the flow group and traffic class do not specify a new value.

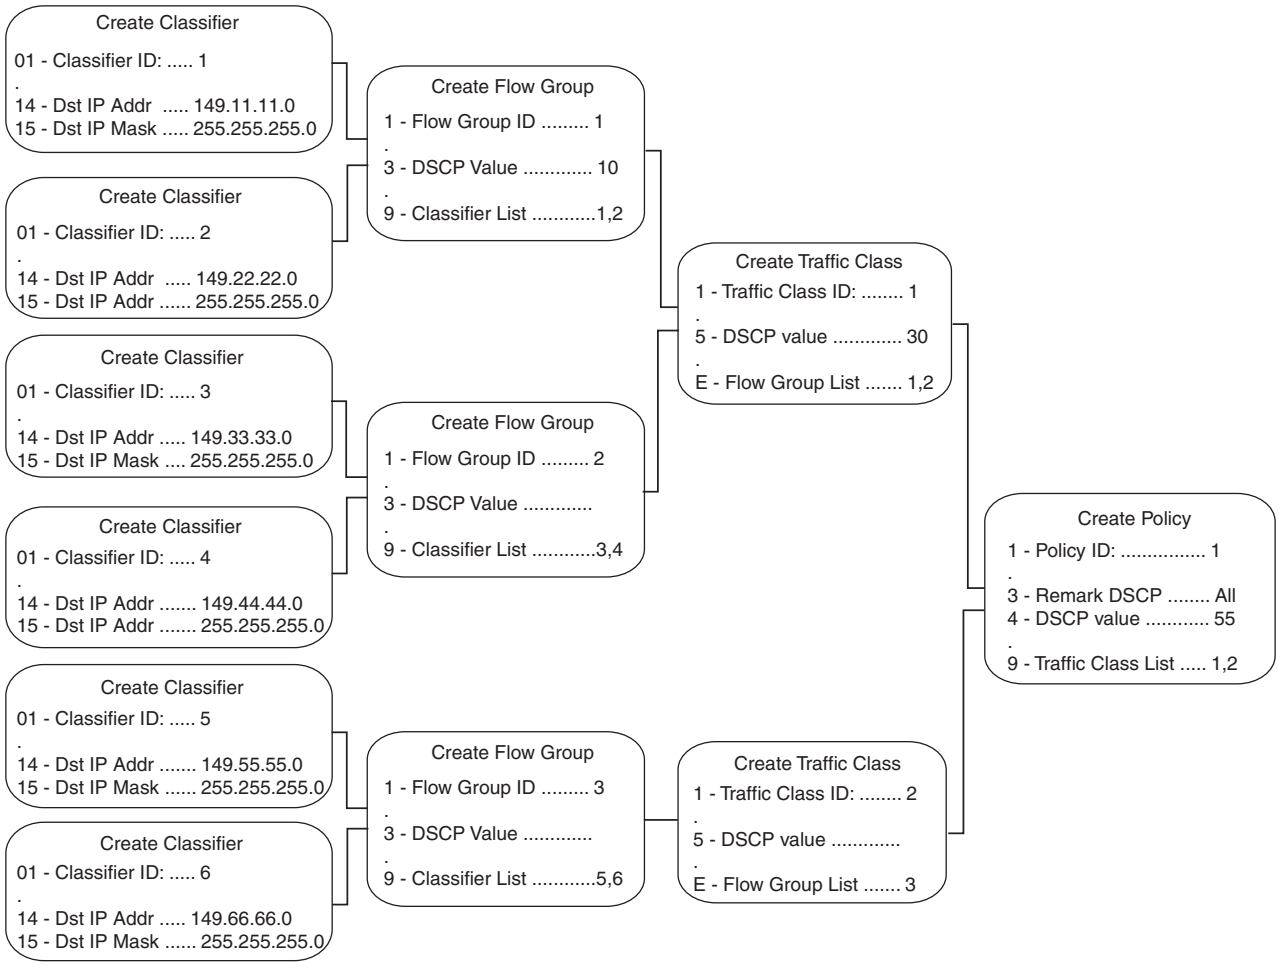

Figure 86. Policy Component Hierarchy Example

Managing Flow Groups

This section contains the following procedures:

Creating a Flow Group

To create a flow group, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                      11:20:02 02-Jan-2006
                 Quality of Service (QoS)
1 - Flow Group Configuration
2 - Traffic Class Configuration
3 - Policy Configuration

R - Return to Previous Menu

Enter your selection?
```

Figure 87. Quality of Service (QoS) menu

3. From the Quality of Service (QoS) menu, type **1** to select Flow Group Configuration.

   The Flow Group Configuration menu is shown in Figure 88.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                      11:20:02 02-Jan-2006
                Flow Group Configuration
1 - Create Flow Group
2 - Modify Flow Group
3 - Destroy Flow Group
4 - Show Flow Groups

R - Return to Previous Menu

Enter your selection?
```

Figure 88. Flow Group Configuration Menu

4. Type **1** to select Create Flow Group.

   The Create Flow Group menu is shown in Figure 89.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                          11:20:02 02-Mar-2006
                    Create Flow Group
1 - FlowGroup ID ............... 0
2 - Description ................
3 - DSCP value .................
4 - Priority ...................
5 - Remark Priority ............ No
6 - ToS ........................
7 - Move ToS to Priority ...... No
8 - Move Priority to ToS ...... No
9 - Classifier List ............

C - Create Flow Group
R - Return to Previous Menu

Enter your selection?
```

Figure 89. Create Flow Group Menu

5. Configure the parameters as needed. The parameters are described below:

   1 - Flow Group ID
   Specifies an ID number for the flow group. Each flow group on the switch must have a unique number. The range is 0 to 1023. The default is 0. This parameter is required.

   2 - Description
   Specifies a description for the flow group. The description can be from 1 to 15 alphanumeric characters including spaces. This parameter is optional, but recommended. Names can help you identify the groups on the switch.

   3 - DSCP value
   Specifies a replacement value to write into the DSCP (TOS) field of the packets. The range is 0 to 63.

   A new DSCP value can be set at all three levels: flow group, traffic class, and policy. A DSCP value specified in a flow group overrides a DSCP value specified at the traffic class or policy level.

   4 - Priority
   Specifies a new user priority value for the packets. The range is 0 to 7. If you specify a new user priority value here and in Traffic Class, the value here overrides the value in Traffic Class. If you want the packets to retain the new value when they exit the switch, change option 5, Remark Priority, to Yes.

5 - Remark Priority
If set to Yes, replaces the user priority value in the packets with the new value specified in option 4, Priority. If set to No, which is the default, the packets retain their preexisting priority level.

6 - ToS
Specifies a replacement value to write into the Type of Service (ToS) field of IPv4 packets. The range is 0 to 7.

A new ToS value can be set at all three levels: flow group, traffic class, and policy. A ToS value specified in a flow group overrides a ToS value specified at the traffic class or policy level.

7 - Move ToS to Priority
If set to Yes, replaces the value in the 802.1p priority field with the value in the ToS priority field on IPv4 packets. If set to No, which is the default, the packets retain their preexisting 802.1p priority level.

8 - Move Priority to ToS
If set to Yes, replaces the value in the ToS priority field with the value in the 802.1p priority field on IPv4 packets. If set to No, which is the default, the packets retain their preexisting ToS priority level.

9 - Classifier List
Specifies the classifiers to be assigned to the policy. The specified classifiers must already exist. Separate multiple classifier IDs with commas (e.g., 4,11,13).

6. After configuring the parameters, type **C** to select Create Flow Group.
7. To create another flow group, repeat this procedure starting with step 4. To assign the flow group to a traffic class, go to "Managing Traffic Classes" on page 290.
8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Flow Group

To modify a flow group, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.
3. From the Quality of Service (QoS) menu, type **1** to select Flow Group Configuration.

   The Flow Group Configuration menu is shown in Figure 88 on page 283.

4. Type **2** to select Modify Flow Group.

   The following prompt is displayed:

   ```
   Available Flow Group(s): 0-10
   Enter Flow Group ID :  [0 to 1023] -> 0
   ```

5. Enter the ID number of the flow group you want to modify. You can modify only one flow group at a time.

   The selected flow group is displayed in the Modify Flow Group menu. An example is shown in Figure 90.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                            11:20:02 02-Mar-2006
                    Modify Flow Group

1 - Flow Group ID ..............  2
2 - Description ................  Video1
3 - DSCP value .................  0
4 - Priority ...................  6
5 - Remark Priority ............  No
6 - ToS ........................
7 - Move ToS to Priority .......  No
8 - Move Priority to ToS .......  No
9 - Classifier List ............  11

M - Modify Flow Group
R - Return to Previous Menu

Enter your selection?
```

Figure 90. Modify Flow Group Menu

6. Modify the settings as needed. For parameter definitions, refer to “Creating a Flow Group” on page 283.

   When modifying a flow group, note the following:

   - ❒ You cannot change a flow group’s ID number.
   - ❒ To delete a value from a variable so as to leave it blank, select the variable and then use the backspace key to delete its default value.
   - ❒ Specifying an invalid value for a parameter that already has a value causes the parameter to revert to its default value.

7. After you have finished modifying the parameter settings, type **M** to select Modify Flow Group.

8. To modify another flow group, repeat this procedure starting with step 4. To assign the flow group to a traffic class, go to “Managing Traffic Classes” on page 290.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting a Flow Group

To delete a flow group, do the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.
3. From the Quality of Service (QoS) menu, type **1** to select Flow Group Configuration.

   The Flow Group Configuration menu is shown in Figure 88 on page 283.
4. Type **3** to select Destroy Flow Group.

   The following prompt is displayed:

   ```
   Available Flow Group(s): 0-10
   Enter Flow Group ID :  [0 to 1023] -> 0
   ```
5. Enter the ID number of the flow group you want to delete. You can delete only one flow group at a time.

   The selected flow group is displayed in the Destroy Flow Group menu. You can use the menu to verify that you are deleting the correct group. An example is shown in Figure 91.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                          11:20:02 02-Mar-2006
                     Destroy Flow Group

1 - Flow Group ID ..............  2
2 - Description ................  Video1
3 - DSCP value .................  0
4 - Priority ...................  6
5 - Remark Priority ............  No
6 - ToS ........................
7 - Move ToS to Priority .......  No
8 - Move Priority to ToS .......  No
9 - Classifier List ............  11

D - Destroy Flow Group
R - Return to Previous Menu

Enter your selection?
```

Figure 91. Destroy Flow Group Menu

6. Type **D** to delete the flow group.

   The flow group is deleted from the switch. The group is removed from any traffic classes to which it is assigned.

7. To delete another flow group, repeat this procedure starting with step 4.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying Flow Groups

To display flow groups, do the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.

3. From the Quality of Service (QoS) menu, type **1** to select Flow Group Configuration.

   The Flow Group Configuration menu is shown in Figure 88 on page 283.

4. Type **4** to select Show Flow Groups.

   The Show Flow Groups menu is shown in Figure 92.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                              11:20:02 02-Mar-2006
                           Show Flow Groups
Number of Flow Groups: 5
                                    Parent
ID    Description                   Traffic Class ID     Active
-------------------------------------------------------------
0     Dev database                  22                   Yes
1     Inv database                  5                    No
2     Video1                        14                   Yes
3     Video2                        2                    Yes
4     Demo dev                      1                    Yes

D - Display Flow Group Details
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 92. Show Flow Groups Menu

The Show Flow Groups menu provides the following information:

ID
The flow group's ID number.

Description
A description of the flow group.

Parent Traffic Class ID
The ID number of the traffic class to which the flow group is assigned. A flow group can belong to only one traffic class at a time.

Active
The status of the flow group. If the flow group is part of a QoS policy that is assigned to one or more ports, the flow group is deemed active. If the flow group has not been assigned to a policy or if the policy has not been assigned to any ports, the flow group is deemed inactive.

5. To display the specifics of a flow group, type **D** to select Display Flow Group Details.

   The following prompt is displayed:

```
Available Flow Group(s): 0-10
Enter Flow Group ID :  [0 to 1023] -> 0
```

6. Enter the ID number of the flow group you want to view. You can display only one flow group at a time.

   The specifications of the selected flow group are displayed in the Display Flow Group Details menu. An example is shown in Figure 93. For definitions of the parameters, refer to "Creating a Flow Group" on page 283.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                           11:20:02 02-Mar-2006
                Display Flow Group Details

1 - Flow Group ID ..............  2
2 - Description ................  Video1
3 - DSCP value .................  0
4 - Priority ...................  7
5 - Remark Priority ............  No
6 - ToS ........................
7 - Move ToS to Priority .......  No
8 - Move Priority to ToS .......  No
9 - Classifier List ............  11

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 93. Display Flow Group Detail Menu

Managing Traffic Classes

This section contains the following procedures:

Creating a Traffic Class

To create a traffic class, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.
3. Type **2** to select Traffic Class Configuration.

   The Traffic Class Configuration menu is shown in Figure 94.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                      11:20:02 02-Jan-2006
               Traffic Class Configuration
1 - Create Traffic Class
2 - Modify Traffic Class
3 - Destroy Traffic Class
4 - Show Traffic Classes

R - Return to Previous Menu

Enter your selection?
```

Figure 94. Traffic Class Configuration Menu

4. Type **1** to select Create Traffic Class.

The Create Traffic Class menu is shown in Figure 95.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                           11:20:02 02-Mar-2006
                    Create Traffic Class
1 - Traffic Class ID ..... 0
2 - Description ..............
3 - Exceed Action ............. Drop
4 - Exceed Remark Value ....... 0
5 - DSCP value ...............
6 - Max bandwidth ............
7 - Burst Size ...............
8 - Priority .................
9 - Remark Priority ........... No
A - ToS ......................
B - Move ToS to Priority ...... No
D - Move Priority to ToS ...... No
E - Flow Group List ...........

C - Create Traffic Class
R - Return to Previous Menu

Enter your selection?
```

Figure 95. Create Traffic Class Menu

5. Configure the parameters as needed. The parameters are described below:

   1 - Traffic Class ID
   Specifies an ID number for the traffic class. Each traffic class on the switch must be assigned a unique number. The range is 0 to 511. The default is 0. This parameter is required.

   2 - Description
   Specifies a description for the traffic class. The description can be from 1 to 15 alphanumeric characters. Spaces are allowed. This parameter is optional, but recommended. Names can help you identify the traffic classes on the switch.

   3 - Exceed Action
   Specifies the action to be taken if the traffic of the traffic class exceeds the maximum bandwidth, specified in option 6. There are two possible exceed actions, drop and remark. If drop is selected, traffic exceeding the bandwidth is discarded. If remark is selected, the packets are forwarded after replacing the DSCP value with the new value specified in option 4, Exceed Remark Value. The default is drop.

   4 - Exceed Remark Value
   Specifies the DSCP replacement value for traffic that exceeds the maximum bandwidth. This value takes precedence over the DSCP value set with option 5, DSCP Value. The default is 0.

5 - DSCP value
Specifies a replacement value to write into the DSCP (TOS) field of the packets. The range is 0 to 63.

A new DSCP value can be set at all three levels: flow group, traffic class, and policy. A DSCP value specified in a flow group overrides a DSCP value specified at the traffic class or policy level. A DSCP value specified at the traffic class level is used only if no value has been specified at the flow group level. It will override any value set at the policy level.

6 - Max Bandwidth
Specifies the maximum bandwidth available to the traffic class. This parameter determines the maximum rate at which the ingress port accepts data belonging to this traffic class before either dropping or remarking occurs, depending on option 3, Exceed Action. If the sum of the maximum bandwidth for all traffic classes on a policy exceeds the (ingress) bandwidth of the port to which the policy is assigned, the bandwidth for the port takes precedence and the port discards packets before they can be classified. The range is 0 to 1016 Mbps.

The value for this parameter is rounded up to the nearest Mbps value when this traffic class is assigned to a policy on a 10/100 port, and up to the nearest 8 Mbps value when assigned to a policy on a gigabit port (for example, on a gigabit port, 1 Mbps is rounded to 8 Mbps, and 9 is rounded to 16).

Note
If this option is set to 0 (zero), all traffic that matches that traffic class is dropped. However, an access control list can be created to match the traffic that is marked for dropping, or a subset of it, and given an action of permit, to override this. This functionality can be used to discard all but a certain type of traffic. For more information about configuring access control lists, see Chapter 14, “Access Control Lists” on page 251.

7 - Burst Size
Specifies the size of a token bucket for the traffic class. The token bucket is used in situations where you have set a maximum bandwidth for a class, but where traffic activity may periodically exceed the maximum. A token bucket can provide a buffer for those periods where the maximum bandwidth is exceeded.

Tokens are added to the bucket at the same rate as the traffic class’ maximum bandwidth, set with option 6, Max Bandwidth. For example, a maximum bandwidth of 50 Mbps adds tokens to the bucket at that rate.

If the amount of traffic flow matches the maximum bandwidth, no traffic is dropped because the number of tokens added to the bucket

matches the number being used by the traffic. However, no unused tokens will accumulate in the bucket. If the traffic increases, the excess traffic will be discarded since no tokens are available for handling the increase.

If the traffic is below the maximum bandwidth, unused tokens will accumulate in the bucket since the actual bandwidth falls below the specified maximum. The unused tokens will be available for handling excess traffic should the traffic exceed the maximum bandwidth. Should an increase in traffic continue to the point where all the unused tokens are used up, packets will be discarded.

Unused tokens accumulate in the bucket until the bucket reaches maximum capacity, set by this parameter. When the maximum capacity of the bucket is reached, no extra tokens are added. The range is 4 to 512 Kbps. The default is 512 Kbps.

Note
To use this parameter you must specify a maximum bandwidth using item 6 - Max Bandwidth. Specifying a token bucket size without also specifying a maximum bandwidth serves no function.

8 - Priority
Specifies the priority value in the IEEE 802.1p tag control field that traffic belonging to this traffic class is assigned. Priority values range from 0 to 7 with 0 being the lowest priority and 7 being the highest priority. Incoming frames are mapped into one of eight Class of Service (CoS) queues based on the priority value.

If you want the packets to retain the new value when they exit the switch, change option 9, Remark Priority, to Yes.

If you specify a new user priority value here and in Flow Group, the value in Flow Group overwrites the value here.

9 - Remark Priority
Replaces the user priority value in the packets with the new value specified in option 4, Priority, if set to Yes. If set to No, which is the default, the packets retain their preexisting priority level when they leave the switch.

A - ToS
Specifies a replacement value to write into the Type of Service (ToS) field of IPv4 packets. The range is 0 to 7.

If you specify a new ToS value here and in Flow Group, the value in Flow Group overwrites the value here.

B - Move ToS to Priority
If set to yes, replaces the value in the 802.1p priority field with the value in the ToS priority field for IPv4 packet. If set to No, which is the default, the packets retain their preexisting 802.1p priority level.

D - Move Priority to ToS
If set to yes, replaces the value in the ToS priority field with the value in the 802.1p priority field on IPv4 packets. If set to No, which is the default, the packets retain their preexisting ToS priority level.

E- Flow Group List
Specifies the flow groups to be assigned to the traffic class. The specified flow groups must already exist. Separate multiple IDs with commas (e.g., 4,11,13).

6. After configuring the parameters, type **C** to select Create Traffic Class.
7. To create another traffic class, repeat this procedure starting with step 3. To assign the traffic class to a policy, go to “Managing Policies” on page 299.
8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Traffic Class

To modify a traffic class, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.
3. Type **2** to select Traffic Class Configuration.

   The Traffic Class Configuration menu is shown in Figure 94 on page 290.
4. From the Traffic Class Configuration menu, type **2** to select Modify Traffic Class.

   The following prompt is displayed:

```
Available Traffic Class(es): 0-7
Enter Traffic Class ID :  [0 to 511] -> 0
```

5. Enter the ID number of the traffic class you want to modify. You can modify only one traffic class at a time.

The selected traffic class is displayed in the Modify Traffic Class menu. An example is shown in Figure 96.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                           11:20:02 02-Mar-2006
                    Modify Traffic Class

1 - Traffic Class ID ..... 15
2 - Description ..... Video2
3 - Exceed Action ..... Drop
4 - Exceed Remark Value ..... 0
5 - DSCP value ..... 0
6 - Max bandwidth ..... 0
7 - Burst Size ..... 0
8 - Priority ..... 0
9 - Remark Priority ..... No
A - ToS .....
B - Move ToS to Priority ..... No
D - Move Priority to ToS ..... No
E - Flow Group List ..... 17

M - Modify Traffic Class
R - Return to Previous Menu

Enter your selection?
```

Figure 96. Modify Traffic Class Menu

6. Modify the settings as needed. For parameter definitions, refer to "Creating a Traffic Class" on page 290.

   When modifying a traffic class, note the following:

   - ❒ You cannot change a traffic class' ID number.
   - ❒ To delete a value from a variable so as to leave it blank, select the variable and then use the backspace key to delete its default value.
   - ❒ Specifying an invalid value for a parameter that already has a value causes the parameter to revert to its default value.

7. After you have finished modifying the parameter settings, type **M** to select Modify Traffic Class.

8. To modify another traffic class, repeat this procedure starting with step 4. To assign the traffic class to a policy, go to "Managing Policies" on page 299.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting a Traffic Class

To delete a traffic class, do the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.
3. From the Quality of Service (QoS) menu, type **2** to select Traffic Class Configuration.

   The Traffic Class Configuration menu is shown in Figure 94 on page 290.
4. Type **3** to select Destroy Traffic Class.

   The following prompt is displayed:

   ```
   Available Traffic Class(es): 0-7
   Enter Traffic Class ID :  [0 to 511] -> 0
   ```
5. Enter the ID number of the traffic class you want to delete. You can delete only one traffic class at a time.

   The selected traffic class is displayed in the Destroy Traffic Class menu. An example is shown in Figure 97. You can use the menu to verify that you are deleting the correct traffic class.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                          11:20:02 02-Mar-2006
                   Destroy Traffic Class

1 - Traffic Class ID .......... 15
2 - Description ............... Video2
3 - Exceed Action ............. Drop
4 - Exceed Remark Value ....... 0
5 - DSCP value ................ 0
6 - Max bandwidth ............. 0
7 - Burst Size ................ 0
8 - Priority .................. 0
9 - Remark Priority ........... No
A - ToS .......................
B - Move ToS to Priority ...... No
D - Move Priority to ToS ...... No
E - Flow Group List ........... 17

D - Destroy Traffic Class
R - Return to Previous Menu

Enter your selection?
```

Figure 97. Destroy Traffic Class Menu

6. Type **D** to delete the traffic class.

   The traffic class is deleted from the switch. The class is removed from any policies to which it is assigned.

7. To delete another traffic class, repeat this procedure starting with step 4.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying Traffic Classes

To display the traffic classes, do the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.

3. From the Quality of Service (QoS) menu, type **2** to select Traffic Class Configuration.

   The Traffic Class Configuration menu is shown in Figure 94 on page 290.

4. Type **4** to select Show Traffic Classes.

   The Show Traffic Classes menu is shown in Figure 98.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                              11:20:02 02-Mar-2006
                         Show Traffic Classes

Number of Traffic Classes: 5
                                 Parent
ID   Description                 Policy ID         Active
-------------------------------------------------------------
0    Dev database                6                 Yes
1    Inv database                12                No
2    Video1                      4                 Yes
3    Video2                      5                 Yes
4    Demo dev                    2                 Yes

D - Display Traffic Class Details
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 98. Show Traffic Classes Menu

The Show Traffic Classes menu provides the following information:

ID
The traffic class' ID number.

Description
A description of the traffic class.

Parent Policy ID
The ID number of the policy where the traffic class is assigned. A traffic class can belong to only one policy at a time.

Active
The status of the traffic class. If the traffic class is part of a QoS policy that is assigned to one or more ports, the traffic class is deemed active. If the traffic class has not been assigned to a policy or if the policy has not been assigned to any ports, the traffic class is deemed inactive.

5. To display the specifics of a traffic class, type **D** to select Detail Traffic Class Display.

6. When prompted, enter the ID number of the traffic class you want to view. You can display only one traffic class at a time.

   The specifics of the traffic class are displayed in the Detail Traffic Class Display. For definitions of the parameters, refer to "Creating a Traffic Class" on page 290.

Managing Policies

This section contains the following procedure:

Creating a Policy

To create a policy, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.
3. Type **3** to select Policy Configuration.

   The Policy Configuration menu is shown in Figure 99.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                      11:20:02 02-Jan-2006
                   Policy Configuration
1 - Create Policy
2 - Modify Policy
3 - Destroy Policy
4 - Show Policies

R - Return to Previous Menu

Enter your selection?
```

Figure 99. Policy Configuration Menu

4. Type **1** to select Create Policy.

The Create Policy menu is shown in Figure 100.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                            11:20:02 02-Mar-2006
                       Create Policy
1 - Policy ID ................. 0
2 - Description ...............
3 - Remark DSCP ............... None
4 - DSCP value ................
5 - ToS .......................
6 - Move ToS to Priority ..... No
7 - Move Priority to ToS ..... No
8 - Send to Mirror Port ...... No
9 - Traffic Class List .......
A - Ingress Port List ........
B - Egress Port ..............

C - Create Policy
R - Return to Previous Menu

Enter your selection?
```

Figure 100. Create Policy Menu

5. Configure the following parameters as needed:

   1 - Policy ID
   Specifies an ID number for the policy. Each policy on the switch must be assigned a unique number. The range is 0 to 255. The default is 0. This parameter is required.

   2 - Description
   Specifies a description for the policy. The description can be from 1 to 15 alphanumeric characters. Spaces are allowed. This parameter is optional, but recommended. Names can help you identify the policies on the switch.

   3- Remark DSCP
   Specifies the conditions under which the ingress DSCP value is overwritten. If All is specified, all packets are remarked. If None is specified, the function is disabled. The default is None.

   4 - DSCP value
   Specifies a replacement value to write into the DSCP (TOS) field of the packets. The range is 0 to 63.

   A new DSCP value can be set at all three levels: flow group, traffic class, and policy. A DSCP value specified in a flow group overrides a DSCP value specified at the traffic class or policy level. A DSCP value specified at the policy level is used only if no value has been specified at the flow group and traffic class levels.

5 - ToS
Specifies a replacement value to write into the Type of Service (ToS) field of IPv4 packets. The range is 0 to 7. A ToS value specified at the policy level is used only if no value has been specified at the flow group and traffic class levels.

6 - Move ToS to Priority
If set to yes, replaces the value in the 802.1p priority field with the value in the ToS priority field on IPv4 packets. If set to No, which is the default, the packets retain their preexisting 802.1p priority level.

7 - Move Priority to ToS
If set to yes, replaces the value in the ToS priority field with the value in the 802.1p priority field on IPv4 packets. If set to No, which is the default, the packets retain their preexisting ToS priority level.

8 - Send to Mirror Port
Copies the traffic that meets the criteria of the classifiers to a destination mirror port. If you set this to yes, you must specify the destination port by creating a port mirror, as explained in Chapter 9, "Port Mirroring" on page 165.

9 - Traffic Class List
Specifies the traffic classes to assign to the policy. The specified traffic classes must already exist. Separate multiple IDs with commas (e.g., 4,11,13).

A - Ingress Port List
Specifies the ingress ports to which the policy is to be assigned. Ports can be identified individually (for example, 5,7,22), as a range (for example, 18-23), or both (for example, 1,5,14-22).

A port can be an ingress port of only one policy at a time. If a port is already an ingress port of a policy, you must remove the port from its current policy assignment before adding it to another policy.

B - Egress Port
Specifies the egress port to which the policy is to be assigned. You can specify only one egress port.

A port can be an egress port of only one policy at a time. If a port is already an egress port of a policy, you must remove the port from its current policy assignment before adding it to another policy.

6. After configuring the parameters, type **C** to select Create Policy.

   The new policy is immediately activated on the specified ports.

7. To create another policy, repeat this procedure starting with step 3.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Policy

To modify a policy, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.
3. Type **3** to select Policy Configuration.

   The Policy Configuration menu is shown in Figure 99 on page 299.
4. From the Policy Configuration menu, type **2** to select Modify Policy.

   The following prompt is displayed:

   ```
   Available Policy(ies): 0-4
   Enter Policy ID :  [0 to 255] -> 0
   ```
5. Enter the ID number of the policy you want to modify. You can modify only one policy at a time.

   The selected policy is displayed in the Modify Policy menu. An example is shown in Figure 101.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                             11:20:02 02-Mar-2006
                        Modify Policy

1 - Policy ID ................. 4
2 - Description ............... VoIP flow
3 - Remark DSCP ............... None
4 - DSCP value .................
5 - ToS ........................
6 - Move ToS to Priority ..... No
7 - Move Priority to ToS ..... No
8 - Send to Mirror Port ...... No
9 - Traffic Class List ....... 15
A - Ingress Port List ........ 7
B - Egress Port ..............

M - Modify Policy
R - Return to Previous Menu

Enter your selection?
```

Figure 101. Modify Policy Menu

6. Modify the settings as needed. For parameter definitions, refer to “Creating a Policy” on page 299.

   When modifying a policy, note the following:

   - You cannot change a traffic class’ ID number.

- ❒ To delete a value from a variable so as to leave it blank, select the variable and then use the backspace key to delete its default value.
- ❒ Specifying an invalid value for a parameter that already has a value causes the parameter to revert to its default value.

7. After you have finished modifying the parameter settings, type **M** to select Modify Policy.

   Modifications to a policy are immediately activated on the ports where the policy is assigned.

8. To modify another policy, repeat this procedure starting with step 4.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting a Policy

To delete a policy, do the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.

3. From the Quality of Service (QoS) menu, type **3** to select Policy Configuration.

   The Policy Configuration menu is shown in Figure 99 on page 299.

4. Type **3** to select Destroy Policy.

   The following prompt is displayed:

   ```
   Available Policy(ies): 0-4
   Enter Policy ID :  [0 to 255] -> 0
   ```

5. Enter the ID number of the policy you want to delete. You can delete only one policy at a time.

   The Destroy Policy menu is displayed. The menu contains the specifications of the selected policy. Use this menu to confirm that you are deleting the correct policy.

6. Type **D** to delete the policy.

   The policy is deleted from the switch.

7. To delete another policy, repeat this procedure starting with step 4.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying Policies

To display policies, do the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **6** to select Quality of Service.

   The Quality of Service (QoS) menu is shown in Figure 87 on page 283.
3. From the Quality of Service (QoS) menu, type **3** to select Policy Configuration.

   The Policy Configuration menu is shown in Figure 99 on page 299.
4. Type **4** to select Show Policies.

   The Show Policies menu is shown in Figure 102.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                          11:20:02 02-Mar-2006
                        Show Policies

Number of Policies: 4
ID    Description             Active
-----------------------------------------------
0     P1-4 database           No
1     Main video              Yes
2     Dev eng                 Yes
3     Alt video               Yes

D - Display Policy Details
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 102. Show Policies Menu

The Show Policies menu provides the following information:

ID
The policy's ID number.

Description
A description of the policy.

Active
The status of the policy. A policy that is assigned to one or more ports is deemed active while a policy that is not assigned to any ports is deemed inactive.

5. To display the specifics of a policy, type **D** to select Detail Policy Display.

6. When prompted, enter the ID number of the policy you want to view. You can display only one policy at a time.

   The specifications of the policy are displayed in the Detail Policy Display. For definitions of the parameters, refer to "Creating a Policy" on page 299.

Chapter 16

Class of Service

This chapter contains the procedures for configuring Class of Service (CoS). Sections in the chapter include:

Class of Service 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, resulting in the delay of packets 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 of no consequence to a network 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 CoS is 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.

CoS applies primarily to tagged packets. A tagged packet, as explained in “Tagged VLAN Overview” on page 555, contains information within it that specifies the VLAN to which the packet belongs.

A tagged packet can also contain a priority level. This priority level is used by network switches and other networking devices to know how important (delay sensitive) that packet is compared 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 priority. The priorities are 0 to 7, with 0 the lowest priority and 7 the highest.

When a tagged packet is received on a port on the switch, it is examined by the AT-S62 software for its priority. The switch software uses the priority to determine which egress priority queue the packet should be directed to on the egress port.

Each switch port has four egress queues, labeled Q0, Q1, Q2, and Q3. Q0 is the lowest priority queue and Q3 is the highest. A packet in a high priority egress queue is typically transmitted out a port sooner than a packet in a low priority queue.

Table 8 lists the mappings between the eight CoS priority levels and the four egress queues of a switch port.

Table 8. Default Mappings of IEEE 802.1p Priority Levels to Priority Queues

IEEE 802.1p Priority Level	Port Priority Queue
0	Q1
1	Q0
2	Q0
3	Q1
4	Q2
5	Q2
6	Q3
7	Q3

For example, if a tagged packet with a priority level of 3 entered a port on the switch, the switch would store the packet in Q1 queue on the egress port.

Priority 0 is mapped to CoS queue 1 instead of CoS queue 0 because tagged traffic that has never been prioritized has a VLAN tag User Priority of 0. If priority 0 was mapped to CoS queue 0, this default traffic goes to the lowest queue, which is probably undesirable. This mapping also makes it possible to give some traffic a lower priority than the default traffic.

You can change these mappings. For example, you might decide that packets with a priority of 5 need to be handled by egress queue Q3 and packets with a priority of 2 should be handled in Q1. The result is shown in Table 9.

Table 9. Customized Mappings of IEEE 802.1p Priority Levels to Priority Queues

IEEE 802.1p Priority Level	Port Priority Queue
0	Q1
1	Q0
2	Q1
3	Q1

Table 9. Customized Mappings of IEEE 802.1p Priority Levels to Priority Queues

IEEE 802.1p Priority Level	Port Priority Queue
4	Q2
5	Q3
6	Q3
7	Q3

The procedure for changing the default mappings is found in “Mapping CoS Priorities to Egress Queues” on page 316. 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 completely ignore the priority levels in its tagged packets and store all the packets in the same egress queue. For instance, perhaps you decide that all tagged packets received on port 4 should be stored in an egress port’s Q3 egress queue, regardless of the priority level in the packets themselves. The procedure for overriding priority levels is explained in “Configuring CoS” on page 313.

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 313.

One last thing to note is that the AT-S62 software 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.

Scheduling

A switch port needs a mechanism for knowing the order in which it should handle the packets in its four egress queues. For example, if all the queues contain packets, should the port transmit all packets from Q3, the highest priority queue, before moving on to the other queues, or should it instead just do a few packets from each queue and, if so, how many?

This control mechanism is called *scheduling*. Scheduling determines the order in which a port handles the packets in its egress queues. The AT-S62 software has two types of scheduling:

- Strict priority
- Weighted round robin priority

Note
Scheduling is set at the switch level. You cannot set this on a per-port basis.

Strict Priority Scheduling

With this type of scheduling, a port transmits all packets out of higher priority queues before transmitting any from the lower priority queues. For instance, as long as there are packets in Q3 it does not handle any packets in Q2.

The value to this type of scheduling is that high priority packets are always handled before low priority packets.

The problem with this method is that some low priority packets might never be transmitted out the port because a port might never get to the low priority queues. A port handling a large volume of high priority traffic may be so busy transmitting that traffic that it never has an opportunity to get to any packets that are stored in its low priority queues.

Weighted Round Robin Priority Scheduling

The weighted round robin scheduling method functions as its name implies. The port transmits a set number of packets from each queue, in a round robin fashion, so that each has a chance to transmit traffic. This method guarantees that every queue receives some attention from the port for transmitting packets.

To use this scheduling method, you need to specify the maximum number of packets a port should transmit from a queue before moving to the next queue. This is referred to as specifying the "weight" of a queue. In all likelihood, you will want to give greater weight to the packets in the higher priority queues over the lower queues.

Table 10 shows an example.

Table 10. Example of Weighted Round Robin Priority

Port Egress Queue	Maximum Number of Packets
Q3	15
Q2	10
Q1	5
Q0	1

In this example, the port transmits a maximum number of 15 packets from Q3 before moving to Q2, from which it transmits up to 10 packets, and so forth.

Configuring CoS

As explained in "Class of Service Overview" on page 308, a tagged packet received on 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 8 on page 309.

You can override the mappings at the port level by assigning the packets a temporary priority level. Note that this assignment is made when a packet is received on the ingress port and 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 switch port so that all ingress frames are assigned a temporary priority level of 5, regardless of the actual priority levels that might be in the frames themselves, as found in tagged frames.

A temporary priority level applies only while a frame traverses the switching matrix. Tagged frames, which can contain a priority level, 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 Main Menu, type **7** to select Security and Services. The Security and Services menu is shown in Figure 103.

```
Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                11:20:02 02-Jan-2006
              Security and Services
1 - Classifier Configuration
2 - Port Access Control (802.1X)
3 - Denial of Service (DoS)
4 - Access Control Lists (ACL)
5 - Class of Service (CoS)
6 - Quality of Service (QoS)
7 - Keys/Certificates Configuration
8 - Secure Shell (SSH)
9 - Secure Socket Layer (SSL)

R - Return to Previous Menu

Enter your selection?
```

Figure 103. Security and Services Menu

2. From the Security and Services menu, type **5** to select Class of Service (CoS).

   The Class of Service (CoS) menu is shown in Figure 104.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                11:20:02 02-Jan-2006
              Class of Service (CoS)
Number of CoS Queues: 4
1 - Configure Port CoS Priorities
2 - Map CoS Priority to Egress Queue
3 - Configure Egress Scheduling
4 - Show Port CoS Priorities

R - Return to Previous Menu
Enter your selection?
```

Figure 104. Class of Service (CoS) Menu

   The “Number of CoS Queues” line indicates the number of egress queues each port has. On the AT-8500 Series switch, there are four queues per port. This value cannot be changed.

3. From the Class of Service menu, type **1** to select Configure Port CoS Priorities.

   The following prompt is displayed:

```
Enter port number -> [1 to 26] ->
```

4. Enter the number of the port on the switch where you want to configure CoS. You can specify only one port at a time.

   The Configure Port COS Priorities menu is shown in Figure 105.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                   11:20:02 02-Jan-2006
           Configure Port COS Priorities
1 - Port Number ..................... 1
2 - Priority (0-7) 0=Low 7=High ... 0
3 - Override Priority (Y/N) ....... N

C - Configure Port COS Priorities
R - Return to Previous Menu

Enter your selection?
```

Figure 105. Configure Port COS Priorities Menu

Menu option 1 cannot be changed.

5. Type **2** to select Priority (0 - 7). The following prompt is displayed:

   ```
   Enter new value -> [0 to 7]
   ```

6. Enter the new temporary priority value of 0 to 7 for the untagged frames received on the port. For example, to assign a temporary priority level of 4 to the ingress untagged packets, enter 4. The default is 0. (If you perform Step 7 and override the priority level in ingress tagged packets, this temporary priority value will also apply to those packets as well.)

7. If you are configuring a tagged port and you want the switch to ignore the priority tag in ingress tagged frames, type **3** to select Override Priority and type **Y**.

   All ingress tagged frames use the temporary priority level specified in Step 6.

   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.

   The default for this parameter is No, meaning that the priority level of tagged frames is determined by the priority level specified in the frame itself.

8. Type **C** to select Configure Port COS Priorities. A change to a CoS setting is immediately activated on the port.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Mapping CoS Priorities to Egress Queues

This procedure explains how to change the default mappings of CoS priorities to egress priority queues, shown in Table 10 on page 311. This is set at the switch level. You cannot set this at the per-port level.

To change the mappings, perform the following procedure.

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **5** to select Class of Service (CoS).

   The Class of Service (CoS) menu is shown in Figure 104 on page 314.

3. From the Class of Service (CoS) menu, type **2** to select Map CoS Priority to Egress Queue.

   The Map CoS Priority to Egress Queue menu is shown in Figure 106.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                11:20:02 02-Jan-2006
         Map CoS Priority to Egress Queue
1 - CoS 0 Priority Queue ..... Q1
2 - CoS 1 Priority Queue ..... Q0
3 - CoS 2 Priority Queue ..... Q0
4 - CoS 3 Priority Queue ..... Q1
5 - CoS 4 Priority Queue ..... Q2
6 - CoS 5 Priority Queue ..... Q2
7 - CoS 6 Priority Queue ..... Q3
8 - CoS 7 Priority Queue ..... Q3

R - Return to Previous Menu
Enter your selection?
```

Figure 106. Map CoS Priority to Egress Queue Menu

4. Type the number of the CoS priority whose queue assignment you want to change. This toggles the queue value through the possible queue settings.

   For example, to direct all tagged packets with a CoS priority of 5 to egress queue Q3, you would toggle 6 until the CoS 5 Priority Queue value reads Q3.

5. If desired, repeat Step 3 to change the queue assignments of other CoS priorities.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring Egress Scheduling

This procedure explains how to select and configure a scheduling method for Class of Service. Scheduling determines the order in which the ports handle packets in their egress queues. For an explanation of the two scheduling methods, refer to “Scheduling” on page 310. Scheduling is set at the switch level. You cannot set this on a per-port basis.

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **5** to select Class of Service (CoS).

   The Class of Service (CoS) menu is shown in Figure 104 on page 314.

3. From the Class of Service (CoS) menu, type **3** to select Configure Egress Scheduling.

   The Configure Egress Scheduling menu is shown in Figure 107.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                  11:20:02 02-Jan-2006
            Configure Egress Scheduling
1 - Scheduling Mode ............. Strict Priority
2 - Queue 0 Weight .............. 0
3 - Queue 1 Weight .............. 0
4 - Queue 2 Weight .............. 0
5 - Queue 3 Weight .............. 0

R - Return to Previous Menu
Enter your selection?
```

Figure 107. Configure Egress Scheduling Menu

4. Type **1** to toggle Scheduling Mode between its two possible settings. The default setting is Strict Priority.

   If you select Strict Priority, skip the next step. Options 2 through 5 in the menu do not apply to Strict Priority scheduling.

5. If you select Weighted Round Robin Priority as the scheduling method, select menu options 2 through 5 and specify the maximum number of packets you want a port to transmit from each queue before it moves to the next queue. The range is 0 to 255. For an example, refer to Table 10 on page 311. The default value of 1 for each queue gives all egress queues the same weight.

6. Return to the Main Menu and type **S** to select Save Configuration Changes.

Displaying Port CoS Priorities

The following procedure displays a menu that lists the current CoS priority level settings for each port.

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **5** to select Class of Service (CoS).

   The Class of Service (CoS) menu is shown in Figure 104 on page 314.

3. From the Class of Service (CoS) menu, type **4** to select Show Port CoS Priorities.

   The Show Port CoS Priorities menu is shown in Figure 108.

```
 Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                Production Switch
User: Manager                11:20:02 02-Jan-2006
             Show Port CoS Priorities
Port   PVID    Priority   Override Priority
-------------------------------------------

01     1       0          No
02     1       0          No
03     1       0          No
04     1       0          No
05     1       0          No
06     1       0          No
07     1       0          No

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 108. Show Port CoS Priorities Menu

The PVID column displays the identifier of the VLAN where the port is an untagged member.

The Priority column displays the temporary priority level assigned to ingress untagged packets on the port.

The Override Priority column indicates whether the priority level in ingress tagged frames is being used or not. If No, the override is

deactivated and the port is using the priority levels contained within the frames. If Yes, the override is activated and the tagged packets are assigned the temporary priority level shown in the Priority column.

Chapter 17
IGMP Snooping

This chapter explains how to activate and configure the Internet Group Management Protocol (IGMP) snooping feature on the switch. Sections in the chapter include:

IGMP Snooping Overview

IGMP enables 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 wanting to become a member of a particular multicast group responds to a query by sending a *report*. A report 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 out 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 the appropriate membership list. The router also stops sending multicast packets out 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 through the use of *Group-Source report* and *Group-Source leave* messages.

The IGMP snooping feature on the AT-8500 Series switch supports all three IGMP versions. It enables the switch to monitor 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 switch 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 switch ports connected to host nodes.

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

The AT-8500 Series switch maintains its 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
The default setting for IGMP snooping on the switch is disabled.

Configuring IGMP Snooping

To configure the IGMP snooping parameters on the switch, perform the following procedure:

1. From the Main Menu, type **6** to select Advanced Configuration.

   The Multicast Configuration menu is shown in Figure 109.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                      11:20:02 02-Jan-2006
                  Advanced Configuration
1 - IGMP Snooping Configuration

R - Return to Previous Menu

Enter your selection?
```

Figure 109. Advanced Configuration Menu

2. From the Advanced Configuration menu, type **1** to select IGMP Snooping Configuration.

   The IGMP Snooping Configuration menu is shown in Figure 110.

```
         Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                         Production Switch
User: Manager                                 11:20:02 02-Jan-2006
                     IGMP Snooping Configuration
1 - IGMP Snooping Status ......... Disabled
2 - Multicast Host Topology ...... Single-Host/Port (Edge)
3 - Host/Router Timeout Interval . 260 seconds
4 - Maximum Multicast Groups ..... 64
5 - Multicast Router Port(s) ..... Auto Detect
6 - View Multicast Hosts List
7 - View Multicast Routers List

R - Return to Previous Men

Enter your selection?
```

Figure 110. IGMP Snooping Configuration Menu

The options in the menu are defined below:

1 - IGMP Snooping Status
Enables and disables IGMP snooping on the switch. After selecting this option, type **E** to enable or **D** to disable this feature.

2 - Multicast Host Topology
Defines whether there is only one host node per switch port or multiple host nodes per port. Possible settings are Single-Host/Port (Edge) and Multiple Host/Ports (Intermediate).

The Single-Host/Port 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 out 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. The switch responds by immediately ceasing the transmission of additional multicast packets out the port where the host node is connected.

The Multi-Host 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 out 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 will continue to receive the multicast packets. Only after all the host nodes connected to a switch port have transmitted leave requests or have timed out will the switch stop sending multicast packets out the port.

If a switch has a mixture of host nodes, that is, some connected directly to the switch and others through an Ethernet hub, you should select the Multi-Host Port (Intermediate) selection.

3 - Host/Router Timeout Interval
Specifies the time period in seconds at which the switch determines that a host node has become inactive. An inactive host node is a node that has not sent an IGMP report during the specified time interval. The range is from 0 second to 86,400 seconds (24 hours). The default is 260 seconds.

This parameter also specifies the time interval used by the switch in determining whether a multicast router is still active. The switch makes the determination by watching for queries from the router. If the switch does not detect any queries from a multicast router during the specified time interval, it assumes that the router is no longer active on the port.

The actual timeout may be ten seconds less that the specified value. For example, a setting of 25 seconds can result in the switch classifying a host node or multicast router as being inactive after just 15 seconds. A setting of 10 seconds or less can result in the immediate timeout of an inactive host node or router.

A value of 0 disables the timer. A switch with a disabled timer never times out inactive host nodes or multicast routers.

4 - Maximum Multicast Groups
Specifies the maximum number of multicast groups the switch will learn. This parameter is useful with networks that contain a large number of multicast groups. You can use the parameter to prevent the switch's MAC address table from filling up with multicast addresses, leaving no room for dynamic or static MAC addresses. The range is 1 to 255 groups. The default is 64 multicast groups.

5 - Multicast Router Port(s)
Specifies the port on the switch to which a multicast router is detected. You can let the switch determine this automatically by selecting Auto Detect, or you can specify the port yourself by entering a port number. To select Auto Detect, enter "0" (zero) for this parameter. You can specify more than one port.

Your changes are immediately activated on the switch.

Note
Option "6 - View Multicast Hosts List" is described in "Displaying a List of Host Nodes", next. Option "7 - View Multicast Routers List" is described in "Displaying a List of Multicast Routers" on page 331.

---

3. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying a List of Host Nodes

To view a list of the multicast groups and host nodes on a switch, perform the following procedure:

1. From the Main Menu, type **6** to select Advanced Configuration.

   The Advanced Configuration menu is shown in Figure 109 on page 326.

2. From the Advanced Configuration menu, type **1** to select IGMP Snooping Configuration.

   The IGMP Snooping Configuration menu is shown in Figure 110 on page 326.

3. From the IGMP Snooping Configuration menu, type **6** to select View Multicast Host List.

   The View Multicast Host List is shown in Figure 111.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
 User: Manager                                  11:20:02 02-Jan-2006
                       View Multicast Hosts List
Number of Multicast Groups: 4

                     VLAN   Port/                     IGMP   Exp.
MulticastGroup       ID     TrunkID   HostIP         Ver    Time
-------------------------------------------------------------------
01:00:5E:00:01:01    1      6/-       172.16.10.51    v2    21
01:00:5E:7F:FF:FA    1      5/-       149.35.200.75   v2    11
01:00:5E:00:00:02    1      17/-      149.35.200.69   v2    34
01:00:5E:00:00:09    1      14/-      172.16.10.51    v2    32

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 111. View Multicast Hosts List Menu

The information in this menu is for viewing purposes only. The columns are defined below:

Multicast Group - The multicast address of the group.

VLAN ID - The VID of the VLAN where the port is an untagged member.

Port/TrunkID - The port on the switch where a host node of the multicast group is connected. If the host node is connected to the switch through a trunk, the trunk ID number, not the port number, is displayed.

HostIP - The IP address of the host node connected to the port.

IGMP Ver. - The version of IGMP being used by the host.

Exp. Time - The number of seconds remaining before the host is timed out if no further IGMP reports are received from it.

Displaying a List of Multicast Routers

A multicast router is a router that is receiving multicast packets from a multicast application and transmitting the packets to host nodes. You can use the AT-S62 software to display a list of the multicast routers that are connected to the switch. To display a list of the multicast routers, perform the following procedure:

1. From the Main Menu, type **6** to select Advanced Configuration.

   The Advanced Configuration menu is shown in Figure 109 on page 326.

2. From the Advanced Configuration menu, type **1** to select IGMP Snooping Configuration.

   The IGMP Snooping Configuration menu is shown in Figure 110 on page 326.

3. From the IGMP Snooping Configuration menu, type **7** to select View Multicast Routers List. The View Multicast Routers List menu is shown in Figure 112.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
              View Multicast Routers List
VLAN    Port/TrunkID     RouterIP
-------------------------------------------------

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 112. View Multicast Routers List Menu

The information in this menu is for viewing purposes only. The columns are defined below:

VLAN
The VID of the VLAN where the port is an untagged member.

Port/Trunk ID
The port on the switch where the multicast router is connected. If the switch learned the router on a port trunk, the trunk ID number, not the port number, is displayed.

Router IP
The IP address of the multicast router.

Chapter 18
Denial of Service Defenses

This chapter contains procedures on how to configure the switch to protect your network against Denial of Service (DoS) attacks. Sections in the chapter include:

Denial of Service Defense Overview

The AT-S62 management software can help protect your network against the following types of Denial of Service attacks.

- SYN Flood Attack
- SMURF Attack
- Land Attack
- Teardrop Attack
- Ping of Death Attack
- IP Options Attack

The following subsections describe each type of attack and the mechanism employed by the AT-S62 management software to protect your network.

> **Note**
> Be sure to read the following descriptions before implementing a DoS defense on a switch. Some defense mechanisms are CPU intensive and can impact switch behavior.

SYN Flood Attack

In this type of attack, an attacker sends a large number of TCP connection requests (TCP SYN packets) with bogus source addresses to the victim. The victim responds with acknowledgements (SYN ACK packets), but since the original source addresses are bogus, the victim node does not receive any replies. If the attacker sends enough requests in a short enough period, the victim may freeze operations when the number of requests exceeds the capacity of its connections queue.

To defend against this form of attack, a switch port monitors the number of ingress TCP connection requests it receives. If a port receives more than 60 requests per second, it assumes that an attack might be occurring. The switch does the following:

- It sends a SNMP trap to the management workstations
- The port discards all ingress TCP-SYN packets for one minute. However, the port continues to allow existing TCP connections to go through.

This defense mechanism does not involve the switch's CPU. You can activate it on as many ports as you want without it impacting switch performance.

SMURF Attack

This DoS attack is instigated by an attacker sending a ICMP Echo (Ping) request containing a broadcast address as the destination address and the address of the victim as the source of the ICMP Echo (Ping) request. This overwhelms the victim with a large number of ICMP Echo (Ping) replies from the other network nodes.

A switch port defends against this form of attack by examining the destination addresses of ingress ICMP Echo (Ping) request packets and discarding those that contain a broadcast address as a destination address.

Implementing this defense requires providing an IP address of a node on your network and a subnet mask. The switch uses the two to determine the broadcast address of your network.

This defense mechanism does not involve the switch's CPU. You can activate it on as many ports as you want without it impacting switch performance.

Land Attack

In this attack, an attacker sends a bogus IP packet where the source and destination IP addresses are the same. This leaves the victim thinking that it is sending a message to itself.

The most direct approach for defending against this form of attack would be for the AT-S62 management software to check the source and destination IP addresses in the IP packets, searching for and discarding those with identical source and destination addresses. But this would require too much processing by the switch's CPU, and could adversely impact switch performance.

Instead, the switch examines the IP packets that are entering and leaving your network. IP packets generated within your network and containing a local IP address as the destination address are not allowed to leave the network, and IP packets generated outside the network but containing a local IP address as the source address are not allowed into the network.

In order for this defense mechanism to work, you need to specify an uplink port. This is the port on the switch that is connected to a device that leads outside your network, such as a DSL router. You can specify only one uplink port.

You also need to enter the IP address of one of your network devices as well as a mask which the switch uses to differentiate between the network portion and node portion of the address. The switch uses the IP address and mask to determine which IP addresses are local to your network and which are from outside you network.

The following is a overview of how the process takes place. This example assumes that you have activated the feature on port 4, which is connected to a device local to your network, and specified port 1 as the uplink port,

which is connected to the device that leads outside your network. The steps below review what happens when an ingress IP packet from the local device arrives on port 4:

1. When port 4 receives an ingress IP packet with a destination MAC address learned on uplink port 1, it examines the packet's source IP address.

2. If the source IP address is not local to the network, it discards the packet because it assumes that a packet with an IP address that is not local to the network should not be appearing on a port that is not an uplink port. This protects against the possibility of a Land attack originating from within your network.

3. If the source IP address is local to the network, the port forwards the packet to uplink port 1.

Below is a review of how the process takes place when an ingress IP packet arrives on uplink port 1 that is destined for port 4:

1. When uplink port 1 receives an ingress IP packet with a destination MAC address that was learned on port 4, it examines the packet's source IP address before forwarding the packet.

2. If the source IP address is local to the network, uplink port 1 does not forward the packet to port 4 because it assumes that a packet with a source IP address that is local to the network should not be entering the network from outside the network on the uplink port.

3. If the source IP address is not local to the network, port 1 forwards the packet to port 4.

The following guidelines apply to using this defense mechanism:

- ❒ If you choose to use it, Allied Telesyn recommends activating it on all ports on the switch, including the uplink port.
- ❒ You can specify only one uplink port.
- ❒ You must specify the IP address of one of the network nodes, preferably the lowest IP address, and a mask.

Note
You should not use this defense mechanism on a switch if it is not connected to a device that leads outside your network.

This form of defense is not CPU intensive. Activating it on all ports should not affect switch behavior.

Teardrop Attack

An attacker sends an IP packet in several fragments with a bogus offset value, used to reconstruct the packet, in one of the fragments to a victim. The victim is unable to reassemble the packet, possibly causing it to freeze operations.

The defense mechanism for this type of attack has all ingress IP traffic received on a port sent to the switch's CPU. The CPU samples related, consecutive fragments, checking for fragments with invalid offset values.

If one is found, the following occurs:

- The switch sends a SNMP trap to the management workstations.
- The switch port discards the fragment with the invalid offset and, for a one minute period, discards all ingress fragmented IP traffic.

Because the CPU only samples the ingress IP traffic, this defense mechanism may catch some, though not necessarily all of this form of attack.

Caution

This defense is extremely CPU intensive; use with caution. Unrestricted use can overwhelm the switch's CPU with IP traffic, causing the unit to halt operations. Even restricted use can impact the switch's handling of CPU events, such as the processing of IGMP packets and spanning tree BPDUs.To prevent this, Allied Telesyn recommends activating this defense on only one switch port at a time.

Ping of Death Attack

The attacker sends an oversized, fragmented ICMP Echo (Ping) request (greater than 65,535 bits) to the victim, which, if lacking a policy for handling oversized packets, may freeze.

To defend against this form of attack, a switch port searches for the last fragment of a fragmented ICMP Echo (Ping) request and examines its offset to determine if the packet size is greater than 63,488 bits. If it is, the fragment is forwarded to the switch's CPU for final packet size determination. If the switch determines that the packet is oversized, the following occurs:

- The switch sends a SNMP trap to the management workstations.
- The switch port discards the fragment and, for one minute, discards all fragmented ingress ICMP Echo (Ping) requests.

Note
This defense mechanism requires some involvement by the switch's CPU, though not as much as the Teardrop defense. This will not impact the forwarding of traffic between the switch ports, but it can affect the handling of CPU events, such as the processing of IGMP packets and spanning tree BPDUs. For this reason, Allied Telesyn recommends limiting the use of this defense, activating it only on those ports where an attack is most likely to originate.

Also note that an attacker can circumvent the defense by sending a stream of ICMP Echo (Ping) requests with a size of 63,488 to 65,534 bits. A large number of requests could overwhelm the switch's CPU.

IP Options Attack

In the basic scenario of an IP attack, an attacker sends packets containing bad IP options. There are several different types of IP option attacks and the AT-S62 management software does not distinguish between them.

The defense mechanism counts the number of ingress IP packets containing IP options received on a port. If the number exceeds 20 packets per second, the switch considers this a possible IP options attack and does the following occurs:

- It sends a SNMP trap to the management workstations.
- The switch port discards all ingress packets containing IP options for one minute.

This defense mechanism does not involve the switch's CPU. You can activate it on as many ports as you want without it impacting switch performance.

Note
This defense does not actually check IP packets for bad IP options. Consequently, it can only alert you to a *possible* attack.

Mirroring Traffic

The Land, Teardrop, Ping of Death, and IP Options defense mechanisms can copy the examined traffic to a mirror port for further analysis with a data sniffer or analyzer. This feature differs slightly from port mirroring in that prior to an actual violation of a defense mechanism, only the packets examined by a defense mechanism, rather than all packets, are mirrored to the destination port. Should a violation occur, then all ingress packets on the port where the violation occurred are mirrored.

As an example, activating the mirroring feature in conjunction with the Teardrop defense on a port sends all examined ingress fragmented IP traffic to the destination mirror port. If the switch detects a violation, all ingress packets on the port are copied to the mirror port during the 60 seconds that the port is blocked.

Implementing this feature requires configuring the port mirroring feature as follows:

- Activate port mirroring.
- Specify a destination port.
- Do not specify any source ports. The source ports are defined by the Denial of Service defense mechanism.

For instructions, refer to "Creating a Port Mirror" on page 167.

Denial of Service Defense Guidelines

Observe these guidelines when using this feature:

- A switch port can support more than one DoS defense at a time.
- The Teardrop and the Ping of Death defenses are CPU intensive. Use these defenses with caution.

Enabling or Disabling Denial of Service Prevention

To configure DoS defense, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **3** to select Denial of Service (DoS).

   The Denial of Service (DoS) Menu is shown in Figure 113.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch

User: Manager                       11:20:02 02-Jan-2006

                          DoS Menu

1 - Lan IP Subnet
2 - SYN Flood Configuration
3 - Smurf Configuration
4 - Land Configuration
5 - Teardrop Configuration
6 - Ping Of Death Configuration
7 - IP Option Configuration

R - Return to Previous Menu

Enter your selection?
```

Figure 113. Denial of Service (DoS) Menu

3. If you are implementing the SMURF or Land defense, you must provide the IP address of a node connected to the switch and a subnet mask. For the Land defense, you must also specify an uplink port. To do this, complete the following steps. Otherwise, skip ahead to Step 4.

   a. Type **1** to select Lan IP Subnet.

The LAN IP Subnet menu is shown in Figure 114.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch

User: Manager                       11:20:02 02-Jan-2006

                        Lan IP Subnet

1 - IP Address ................. 0.0.0.0
2 - Subnet Mask ................ 0.0.0.0
3 - Uplink Port ................ 26

R - Return to Previous Menu

Enter your selection?
```

Figure 114. LAN IP Subnet Menu

b. Type **1** to select IP Address and, when prompted, enter the IP address of one of the devices connected to the switch, preferably the lowest IP address.

c. Type **2** to select Subnet Mask and enter the mask. A binary "1" indicates the switch should filter on the corresponding bit of the IP address, while a "0" indicates that it should not. As an example, assume that the devices connected to a switch are using the IP address range 149.11.11.1 to 149.11.11.50. The mask would be 0.0.0.63.

d. If you are activating the Land defense, type **3** to select Uplink Port and enter the number of the port connected to the device (e.g., DSL router) that leads outside your network. You can specify only one uplink port. The default is the highest numbered existing port in the switch. For example, the default uplink port for an AT-8524M switch with no installed expansion modules would be Port 24.

e. Type **R** to return to the Denial of Service (DoS) Configuration menu and continue with the next step.

4. Type the number of the DoS attack that you want to enable or disable.

5. When prompted, enter the port(s) where you want to enable or disable a defense mechanism.

Note
If you plan to use the Teardrop defense, Allied Telesyn recommends activating it on only the uplink port and one other port. The defense is CPU intensive and can overwhelm the switch's CPU.

A menu is displayed containing either one or two options, depending on the DoS defense you selected. An example of the menu is shown in Figure 115.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                        11:20:02 02-Jan-2006

                SYN Flood Configuration

Configuring DoS for Port 2

1 - DoS Status ................. Disabled

R - Return to Previous Menu

Enter your selection?
```

Figure 115. SYN Flood Configuration Menu

6. Adjust the parameter settings as needed. The parameters are defined below.

   DoS Status
   Enables and disables the selected DoS defense on the selected ports. The default is disabled.

   2 - Mirroring Attack Pkt
   This option is displayed with the Land, Tear Drop, Ping of Death, and IP options defense mechanisms. You can use this option to mirror the traffic examined by a defense mechanism to another port on the switch. For background information, refer to “Mirroring Traffic” on page 338. To use this feature, you must activate port mirroring on the switch and specify a destination mirror port, as explained in “Creating a Port Mirror” on page 167. Mirroring traffic is not required.

7. Repeat this procedure starting with Step 3 to configure other DoS defenses.

8. Return to the Main Menu and type **S** to select Save Configuration Changes.

Chapter 19

Power Over Ethernet

This chapter contains the procedures for configuring Power over Ethernet (PoE) for the AT-8524POE switch. Sections in the chapter include:

Note
This chapter applies only to the AT-8524POE switch.

Power Over Ethernet Overview

The twisted pair ports on the AT-8524POE switch offer the same features as the twisted pair ports on the other switches in the series. As such, they can operate at 10 or 100 Mbps, feature Auto-Negotiation and Auto-MDI/MDI-X, and so forth.

These ports, however, also offer Power over Ethernet (PoE). PoE is a mechanism for supplying power to network devices over the same twisted pair cables used to carry network traffic. This can simplify network installation and maintenance by allowing you to use the switch as a central power source to other network devices.

A device that receives its power over an Ethernet cable is called a *powered device*. Examples can be wireless access points, IP telephones, webcams, and even other Ethernet switches, such as the unmanaged AT-FS705PD Ethernet switch from Allied Telesyn. A powered device connected to a port on the switch receives both network traffic and power over the same twisted pair cable.

There are several advantages that the PoE feature of the AT-8524POE switch adds to the installation and maintenance of your network. Since the switch acts as the central power source for your powered devices, adding a redundant power supply (RPS) or uninterruptible power source (UPS) to the switch increases the protection not just to the switch from possible power source problems but also to all of the powered devices connected to it. This can increase the reliability of your network by minimizing the impact to network operations from a power failure.

PoE can also simplify the installation of your network. The selection of a location for a network device is often limited by whether there is a power source nearby. This often limits equipment placement or requires the added time and cost of having additional electrical sources installed. With PoE, you can install PoE-compatible devices wherever they are needed without having to worry about whether there are power sources nearby.

This feature requires little configuration or management. The switch automatically determines whether a device connected to a port is a powered device or not.

A port on the switch connected to a powered device can supply up to 15.4 watts of power to the device, while at the same time furnishing standard 10/100 Mbps Ethernet functionality. A port 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. The PoE feature remains activated on the port but no power is delivered to the device.

PoE Implementation on the AT-8524POE Switch

A standard Ethernet twisted pair cable contains four pairs of strands for a total of eight strands. 10/100 Mbps network traffic requires only four strands, leaving four strands in the cable unused. The strands that carry the network traffic are 1, 2, 3, and 6, and the spare strands are 4, 5, 7, and 8.

The IEEE 802.3af standard, which is the IEEE standard for PoE, describes two methods for implementing PoE over twisted pair cabling. One method uses the same strands that carry the network traffic and the other the spare strands.

The PoE implementation on the AT-8524POE switch transmits power over the same strands that carry the network traffic. The power transfer does not interfere with the network traffic. The power and the network traffic can coexist on the same strands simultaneously.

Powered devices that comply with the IEEE 802.3af standard typically support both methods of power delivery methods. So you should not need to be concerned about whether a powered device is compatible with the switch's power delivery method. So long as a powered device is compliant with the standard, it should be able to receive its power from the switch.

The PoE feature on the switch should also work with most legacy powered devices as well. A legacy device is a node that was manufactured before the IEEE 802.3af standard was completed and, consequently, may not adhere to the standard.

Power Budgeting

The power supply in the AT-8524POE switch can provide up to a total of 400 watts (W) of PoE to Ports 1 to 24 on the switch. (PoE is not supported on expansion modules.) In a maximum load configuration, where all ports are connected to a powered device and all devices require the maximum of 15.4 W, the total power requirement would be approximately 370 W. This is below the maximum power available.

The fact that the maximum possible power requirement falls below the maximum amount of power available means that you can connect powered devices to all the ports on the switch (excluding optional expansion ports) without exceeding the available power, even when all the powered devices require the maximum of 15.4 W.

You can, using the AT-S62 management software, disable PoE on a per-port basis. You can also reduce the maximum amount of power a port can receive, from the maximum of 15.4 W. However, configuring PoE on an AT-8524POE switch will probably not be necessary. As already mentioned, the power supply in the switch can provide enough power to meet the needs of all 24 base ports, even when all are all connected to power devices requiring the maximum of 15.4 W. Additionally, a switch port can automatically determine for itself whether the device connected to it is PoE-compliant or not and, if it is, how much power is required.

The default setting for PoE on the switch is enabled on all ports.

Port Prioritization

This section explains port prioritization, a mechanism by which the switch determines which ports are to receive PoE in the event the needs of the powered devices exceed the available power resources of the switch. This discussion does not apply to the AT-8524POE switch since its power supply can deliver the maximum of 15.4 W to all 24 based ports simultaneously. This discussion becomes relevant only if, at some later date, Allied Telesyn releases an AT-8500 Series switch with PoE capability with a power supply that cannot service all ports simultaneously.

If the powered devices connected to a switch require more power than the switch is capable of delivering, the switch will deny power to some ports based on a system called port prioritization. You can use this mechanism to ensure that powered devices critical to the operations of your network are given preferential treatment by the switch in the distribution of power should the demands of the devices exceed the available capacity.

There are three priority levels:

- Critical
- High
- Low

The Critical level is the highest priority level. Ports set to this level are guaranteed power before any ports assigned to the other two priority levels. Ports assigned to the other priority levels receive power only if all the Critical ports are receiving power. Your most critical powered devices should be assigned to this level. If there is not enough power to support all the ports set to the Critical priority level, power is provided to the ports based on port number, in ascending order.

The High level is the second highest level. Ports set to this level receive power only if all the ports set to the Critical level are already receiving power. If there is not enough power to support all of the ports set to the High priority level, power is provided to the ports based on port number, in ascending order.

The lowest priority level is Low. This is the default setting. Ports set to this level only receive power if all the ports assigned to the other two levels are already receiving power. As with the other levels, if there is not enough power to support all of the ports set to the Low priority level, power is provided to the ports based on port number, in ascending order.

Power allocation is dynamic. Ports supplying power to powered devices may cease power transmission if the switch's power budget has reached maximum usage and new powered devices, connected to ports with a higher priority, become active.

PoE Device Classes

The IEEE 802.3af standard specifies four levels of classes for powered devices. The classes are defined by power usage. The classes are:

- 0 - 0.44 W to 12.95 W
- 1 - 0.44 W to 3.84 W
- 2 - 3.84 W to 6.49 W
- 3 - 6.49 W to 12.95 W

(The standard actually specifies five levels; the fifth is reserved for future use.)

You cannot adjust this on a powered device. It is set by the manufacturer. This is mentioned here because you can view the class of a powered device through the switch's management software. To view this information, refer to "Displaying PoE Status and Settings" on page 352.

You might notice that according to the IEEE standard the maximum amount of power a powered device should consume is 12.95 W. So why does the switch offer up to 15.4 W per port? It has to do with line loss. Some power is lost on the twisted pair cable as it travels from the switch to the device. For those devices needing 12.95 W, the extra watts act as compensation for the possible loss.

Setting the PoE Threshold

The PoE threshold is a percentage of the total maximum PoE power on the switch, which for the AT-8524POE switch is 400 W. If the total power requirements of the powered devices exceed this threshold, the switch sends an SNMP trap to your management workstation and enters an event in the event log. At the default setting of 95%, the threshold is exceeded when the PoE devices require more than 380 W, which is 95% of 400 W. The threshold is adjustable. Of course, for your management workstations to receive traps from the switch, you must configure SNMP on the switch by specifying the IP address of the workstations.

To configure the PoE threshold, perform the following procedure:

1. From the Main Menu, type **6** to select Advanced Configuration.
2. From the Advanced Configuration menu, type **2** to select Power Over Ethernet Configuration menu.

   The Power Over Ethernet Configuration menu is shown in Figure 116.

```
  Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
        Power Over Ethernet (PoE) Configuration
1 - PoE Global Configuration
2 - PoE Port Configuration
3 - PoE Status

R - Return to Previous Menu

Enter your selection?
```

Figure 116. Power Over Ethernet Configuration Menu

3. From the Power Over Ethernet Configuration menu, type **1** to select PoE Global Configuration. The PoE Global Configuration menu is shown in Figure 117.

```
  Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
               PoE Global Configuration
1 - Power Threshold ................. 95 percent
2 - Maximum Available Power ........ 400W

R - Return to Previous Menu

Enter your selection?
```

Figure 117. PoE Global Configuration Menu

Options 2, Maximum Available Power, displays the maximum amount of PoE for the switch. For the AT-8524POE switch, this value is 400W. This value cannot be changed.

4. From the PoE Global Configuration menu, type **1** to select Power Threshold.

   The following prompt is displayed:

   ```
   Enter percentage of power limit threshold : [1 to 100] -
   > 95
   ```

   Enter the new threshold as a percentage of the total available PoE power on the switch. As an example, to configure the switch to enter an event in the event log and send an SNMP trap when power consumption exceeds 300 W, you would enter 75, for 75%.

   The new threshold is immediately activated on the switch.

5. After making the change, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring PoE Port Settings

This procedure enables and disables PoE on a port. This procedure also sets a port's priority level and its maximum power usage.

To configure PoE port settings, do the following:

1. From the Main Menu, type **6** to select Advanced Configuration.
2. From the Advanced Configuration menu, type **2** to select Power Over Ethernet Configuration.

   The Power Over Ethernet Configuration menu is shown in Figure 116 on page 348.
3. From the Power Over Ethernet Configuration menu, type **2** to select PoE Port Configuration.

   The following prompt is displayed:

   ```
   Enter port-list:
   ```
4. Enter the port you want to configure. You can specify more than one port at a time.

   The PoE Port Configuration menu is shown in Figure 118.

```
  Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                  Production Switch
User: Manager                       11:20:02 02-Jan-2006
                PoE Port Configuration
Port 4
1 - PoE Function ....... ENABLED
2 - Power Priority ..... LOW
3 - Power Limit ........ 15,400 mW

R - Return to Previous Menu

Enter your selection?
```

Figure 118. PoE Port Configuration Menu

   If you are configuring multiple ports, the management software displays the settings of the lowest numbered port.
5. To enable or disable PoE on the port, type **1** to select PoE Function and, when prompted, type **E** to enable PoE or **D** to disable it. The default is Enabled.

6. To change the port's priority, type **2** to select Power Priority and, when prompted, type **C** for Critical, **H** for High, or **L** for Low. A port can belong to only one priority level at a time. The default is Low. For an explanation of this parameter, refer to "Port Prioritization" on page 346.

7. To change the maximum amount of power the port can supply to the device, type **3** to select Power Limit and enter a new value in milliwatts. The default value is 15,400 mW.

   A change to a parameter value is immediately activated on the switch.

8. After making your changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying PoE Status and Settings

Use this procedure to display PoE status and settings at the switch or port level.

To display PoE information, do the following:

1. From the Main Menu, type **6** to select Advanced Configuration.
2. From the Advanced Configuration menu, type **2** to select Power Over Ethernet Configuration.

   The Power Over Ethernet Configuration menu is shown in Figure 116 on page 348.
3. From the Power Over Ethernet Configuration menu, type **3** to select PoE Status.

   The PoE Status menu is shown in Figure 119.

```
  Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                  Production Switch
User: Manager                   11:20:02 02-Jan-2006
                        PoE Status
1 - PoE Global Status
2 - PoE Summary Ports Status
3 - PoE Detailed Ports Status
4 - PoE Device Information

R - Return to Previous Menu

Enter your selection?
```

Figure 119. PoE Status Menu

The selections are defined below.

1 - PoE Global Status Menu

This selection displays the following window:

```
  Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
                  PoE Global Status
Max Available Power ...... 400 W
Consumed Power ........... 25 W
Available Power .......... 375W
Power Usage .............. 6.25 percent
Min Shutdown Voltage ..... 44.0 V
Max Shutdown Voltage ..... 57.0 V

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 120. PoE Global Status Menu

The selections in this window are for viewing purposes only. These parameters are not adjustable. The selections are described below.

Max Available Power
The total available power for PoE supplied by the switch. This value is 400 W for the AT-8524POE switch.

Consumed Power
The amount of power being used by the powered devices.

Available Power
The amount of unused power available for additional powered devices.

Power Usage
The amount of power currently consumed by the powered devices connected to the switch. The value is give as a percentage of the total amount of power available, which for the AT-8524POE switch is 400 W.

Min Shutdown Voltage
The minimum threshold voltage at which the switch shuts down PoE. If the power supply in the switch experiences a problem and the output voltage drops below this value, the switch shuts down PoE on all ports. This value is not adjustable.

Max Shutdown Voltage
The maximum threshold voltage at which the switch shuts down PoE. If the power supply in the switch experiences a problem and the output voltage exceeds this value, the switch shuts down PoE on all ports. This value is not adjustable.

2 - Summary All Ports Status Menu

This selection display an abbreviated status report of PoE on the individual switch ports. For more detailed information, refer to selection 3.

This selection displays the following window:

```
            Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                          Production Switch
User: Manager                                        11:20:02 02-Jan-2006

                          PoE Summary Ports Status

Port PoE Function    Consumed Power (mW)  Power Status
-------------------------------------------------
1    ENABLED         1,900                ON - Valid PD detected
2    ENABLED         1,900                ON - Valid PD detected
3    ENABLED         1,900                ON - Valid PD detected
4    ENABLED         0                    OFF - Detection is in progress
5    ENABLED         0                    OFF - Detection is in progress

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 121. PoE Summary Ports Status Menu

The selections in this window are for viewing purposes only. Each column is described below.

Port
Port number.

PoE Function
Whether PoE is enabled or disabled on the port. The default setting is enabled. To enable or disable PoE on a port, refer to “Configuring PoE Port Settings” on page 350.

Consumed Power
The amount of power in milliwatts currently consumed by the powered device connected to the port. If the port is not connected to a powered device, this value will be 0 (zero).

Power Status
Whether power is being supplied to the device. ON means that the port is providing power to a powered device. OFF means the device is not a powered device or PoE has been disabled on the port.

3 - Detailed Ports Status Menu

When you select this option, you are prompted to enter the port(s) you want to view. You can specify more than one port at a time. Once you have specified the port, the selection displays the following window:

```
     Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                      Production Switch
User: Manager                          11:20:02 02-Jan-2006

                   PoE Detailed Port Status

Port: 4

PoE Function ........... ENABLED
Power Status ........... ON - Valid PD detected
Power Consumed ......... 1,900 mW
Power Limit ............ 15,400 mW
Power Priority ......... Low
Power Class ............ 1
Voltage ................ 48.6V
Current ................ 40 mA

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 122. PoE Summary Ports Status Menu

The selections in this window are for viewing purposes only. Each selection is described below.

Port
Port number.

PoE Function
Whether PoE is enabled or disabled on the port. The default setting is enabled. To enable or disable PoE on a port, refer to "Configuring PoE Port Settings" on page 350.

Power Status
Whether power is being supplied to the device. ON means that the port is providing power to a powered device. OFF means the device is not a powered device, PoE has been disabled on the port, or no device is connected to the port.

Power Consumed
The amount of power in milliwatts currently consumed by the powered device connected to the port. If the port is not connected to a powered device, this value will be 0 (zero).

Power Limit
The maximum amount of power allowed by the port for the device. The default is 15,400 milliwatts (15.4 W). To adjust this value for a port, refer to "Configuring PoE Port Settings" on page 350.

Power Priority
The port priority. This can be Critical, High, or Low. For an explanation of this parameter, refer to “Port Prioritization” on page 346. To adjust this value, refer to “Configuring PoE Port Settings” on page 350.

Power Class
The IEEE 802.3af class of the device. For an explanation of this parameter, refer to “PoE Device Classes” on page 347. This parameter cannot be changed.

Voltage
The voltage being delivered to the powered device

Current
The current drawn by the powered device.

4 - PoE Device Information

This selection displays the hardware and firmware version numbers of the PoE chipset used in the switch. This selection is intended for troubleshooting purposes and displays the following window:

```
  Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
                PoE Device Information

MCU Device Info:
Hardware Version ......... 0
Firmware Version ......... 0290
Build Number ............. 13
Serial Number ............ 0000000

PSE Devices Info:
Device 0 Hardware Version .... 1
Device 1 Hardware Version .... 1

R - Return to Previous Menu

Enter your selection?
```

Figure 123. PoE Device Information

Chapter 20

Networking Stack

The AT-S62 management software allows you to perform a few basic functions on the switch's TCP/IP stack. The functions include viewing the switch's Address Resolution Protocol (ARP) table and routing table. The switch uses these tables when performing a management function that requires it to interact with another network device. You can also view the TCP connections table, which lists the active Telnet, SSH, and web browser management sessions, and a global TCP table, which displays basic TCP status and statistics.

Sections in the chapter include:

Managing the Address Resolution Protocol Table

The switch has an Address Resolution Protocol (ARP) table for storing IP addresses of network devices and their corresponding MAC addresses. The switch uses the table whenever you issue a management command that requires the switch's AT-S62 management software to interact with another device on the network. An example would be if you instructed the switch to ping another network device or download a new AT-S62 image file or configuration file from a network server.

The value of the ARP table is that it eliminates the need of the switch to issue unnecessary ARP broadcast packets when performing some management functions. This can improve the switch's response time as well as reduce the number of broadcast packets on your network.

The table can hold up to 10 entries. There are two types of entries. One type is permanent. There is only one permanent entry. It is used by the switch for internal diagnostics and can never be removed from the table.

The other type is a temporary entry, of which there can be up to nine. The switch adds a temporary entry whenever its management software interacts with another network device during a management function. When you enter a management command that contains an IP address not in the table, the switch sends out an ARP broadcast packet. When the remote device responds with its MAC address, the switch adds the device's IP address and MAC address as a new temporary entry to the table.

A temporary entry remains in the table only while active. An entry remains active so long as it is periodically used by the switch for management functions. If an entry is inactive for a specified period of time, referred to as the ARP cache timeout, it is automatically removed from the table. This value is adjustable, as explained in "Configuring the ARP Table Timeout Value" on page 364. The default is 400 seconds. If the table becomes full, the management software continues to add new temporary entries by deleting the oldest entries.

The management software allows you to view the contents of the table. You can also delete individual table entries or delete all the entries. These functions are explained in the following subsections:

Note
The switch does not use the ARP table to move packets through its switching matrix. The switch refers to the table only when performing a management function that involves interaction with another network node.

Displaying the ARP Table

To view the switch's ARP table, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. From the System Utilities menu, type **6** to select Networking Stack.

   The Networking Stack menu is shown in Figure 124.

```
  Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                  Production Switch

User: Manager                      11:20:02 02-Jan-2006

                    Networking Stack

1 - Display ARP Table
2 - Delete ARP Entry
3 - Reset ARP Table
4 - Display Route Table
5 - Display TCP Connections
6 - Display TCP Global Information
7 - Delete TCP Connection

R - Return to Previous Menu

Enter your selection?
```

Figure 124. Networking Stack Menu

4. From the Networking Stack menu, type **1** to select Display ARP Table.

The Display ARP Table menu is shown in Figure 125.

```
        Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                          Production Switch

User: Manager                                   11:20:02 02-Jan-2006

                          Display ARP Table

Interface     IP Address     MAC Address              Type
-------------------------------------------------------------
loopback      127.0.0.1      00:00:00:00:00:00        PERMANENT
eth0          149.22.22.22   00:30:84:32:8A:5B        TEMPORARY
eth0          149.22.22.1    00:30:84:32:12:42        TEMPORARY
eth0          149.22.22.101  00:30:84:32:8A:1B        TEMPORARY
eth0          149.22.22.27   00:30:84:32:6A:11        TEMPORARY
eth0          149.22.22.86   00:30:84:32:81:22        TEMPORARY

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 125. Display ARP Table Menu

The information in this table is for viewing purposes only. The columns in the menu are defined here:

Interface
The network interface of a table entry. The switch has two network interfaces. The “loopback” designation represents the interface used by the switch for internal diagnostics. The “eth0” designation represents the Ethernet network interface.

IP Address and
MAC Address
The IP addresses and their corresponding MAC addresses.

Type
The type of ARP entry. An entry can be permanent, meaning it can never be deleted from the table, or temporary. Only the “loopback” entry is permanent. All “eth0” entries are temporary.

Deleting an ARP Entry

To delete an entry from the ARP table, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. From the System Utilities menu, type **6** to select Networking Stack.

   The Networking Stack menu is shown in Figure 124 on page 361.

4. From the Networking Stack menu, type **2** to select Delete ARP Entry.

   The following prompt is displayed:

   ```
   Enter IP address of the ARP entry to delete:
   ```

5. Enter the IP address of the entry you want to delete. You cannot delete the first entry in the table with the interface designation "loopback."

   The entry is immediately removed from the switch.

6. Repeat steps 4 and 5 to delete additional ARP table entries.

   You do not need to return to the main menu to save the changes made with this procedure.

Deleting All ARP Entries

To delete all entries from the ARP table, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. From the System Utilities menu, type **6** to select Networking Stack.

   The Networking Stack menu is shown in Figure 124 on page 361.

4. From the Networking Stack menu, type **3** to select Reset ARP Table.

Note
No confirmation prompt is displayed. All entries in the ARP table are immediately deleted, with the exception of the "loopback" entry, which cannot be deleted.

The switch begins to add new entries to the table as it performs new management functions in conjunction with other network devices.

Configuring the ARP Table Timeout Value

Inactive temporary entries in the ARP table are timed out according to the ARP cache timeout value. This parameter prevents the table from becoming full with inactive entries. The default setting is 400 seconds. To set this value, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **2** to select System Configuration.

   The System Configuration menu is shown in Figure 5 on page 53.

3. From the System Configuration menu, type **A** to select ARP Cache Timeout.

   The following prompt is displayed:

   ```
   Enter your new value -> [1 to 260000] -> 400 seconds
   ```

4. Enter the new timeout value in seconds. The range is 1 to 260,000 seconds. The default is 400 seconds.

   A new timeout value takes affect immediately on the switch.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying the Routing Table

The routing table is used by the switch when the IP address of a remote node specified in a management command is not on the same physical network as the switch. The table contains the IP address of the next hop to reaching the remote network or device. For example, the switch might refer to the table if you instructed it to download a new AT-S62 image file from a network server that was on a different physical network.

To view the route table, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. From the System Utilities menu, type **6** to select Networking Stack.

   The Networking Stack menu is shown in Figure 124 on page 361.

4. From the Networking Stack menu, type **4** to select Display Route Table.

   The Display Route Table is shown in Figure 126.

```
          Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                          Production Switch

User: Manager                                    11:20:02 02-Jan-2006

                          Display Route Table

Destination          Mask                 Next Hop          Interface
-------------------------------------------------------------------
127.0.0.0            255.0.0.0            127.0.0.1         loopback
169.254.0.0          255.255.0.0          169.254.37.1      eth0
169.254.37.1         255.255.255.255      127.0.0.1         loopback

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 126. Display Route Table

The information in this table is for viewing purposes only. The columns are defined here:

Destination
The IP address of a destination network, subnetwork, or end node.

Mask
A filter used to designate the active part of the destination IP address. A binary 1 in the mask indicates an active bit in the address while a binary 0 indicates an inactive corresponding bit.

Next Hop
The IP address of the next intermediary device to reaching the destination network, subnetwork, or end node.

Interface
The interface on the switch where the next hop is located. The switch has two interfaces. The interface “loopback” is for internal diagnostics only. The other interface is “eth0.”

Displaying the TCP Connections Table

The TCP connections table lists the active Telnet, SSH, and web browser management sessions on a switch and includes the IP addresses of the management stations. You can use the table to determine the number of active, remote active management sessions open on a switch, as well as identify the management stations.

To view the TCP Connections Table, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. From the System Utilities menu, type **6** to select Networking Stack.

   The Networking Stack menu is shown in Figure 124 on page 361.

4. From the Networking Stack menu, type **5** to select Display TCP Connections.

   An example of the Display TCP Connections table is shown in Figure 127.

```
        Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                        Production Switch

User: Manager                                  11:20:02 02-Jan-2006

                     Display TCP Connections

Total number of TCP Listening sockets : 2
Total number of TCP connections : 2
Index   Local Address       Foreign Address         State
-------------------------------------------------------------
0       0.0.0.0:80          0.0.0.0:0               LISTEN
1       0.0.0.0:23          0.0.0.0:0               LISTEN
4       169.254.37.1:23     169.254.37.138:1051     ESTABLISHED
24      169.254.37.1:80     169.254.37.101:1075     ESTABLISHED

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 127. Display TCP Connections Table

This table is for viewing purposes only. The columns in the table are defined here.

Total Number of TCP Listening sockets
The number of active listening sockets. There can be a maximum of three listening sockets. One is for the Telnet server, another for SSH, and the last for the web browser server. If a server is disabled, its listening socket does not appear in the table.

Total Number of TCP connections
The number of active Telnet, SSH, and web browser connections to the switch.

Index
The internal socket ID number assigned to the connection.

Local Address
The IP address of the switch, followed by the TCP port number used by the switch for the connection. The two values are divided by a colon, as illustrated in Figure 128. The port number indicates the type of TCP connection. A port number of 23 indicates a Telnet connection, 22 an SSH connection, and 80 or 443 a web browser HTTP or HTTPS connection, respectively.

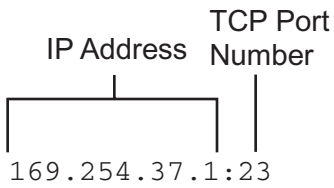

Figure 128. IP Address and TCP Port Number

Foreign Address
The IP address of the management workstation that initiated the connection, followed by the station's TCP port number.

State
The state of the TCP connection. A state of ESTABLISHED signals a successful TCP connection between the switch and the management workstation. For definitions of the TCP states, refer to RFC-793.

The entries for the listening sockets for the Telnet, SSH, and web browser servers are identified in the table with a TCP state of LISTEN. If you disable a server on the switch, its corresponding LISTEN entry is removed from the table. Disabling all the servers leaves the table empty. (The SSH server is disabled by default on the switch.)

The example in Figure 127 on page 367 shows that the Telnet and web browser servers are active on the switch. The table also includes two active TCP connections. Entry 4 is for a Telnet connection and entry 24 is for a web browser HTTP connection.

A web browser management session can have more than one TCP connection open at a time. The different connections carry different packets of the management session.

You cannot change any of the information in this table. The only operating parameter on the switch that affects management TCP connections that you can adjust, other than enabling or disabling the servers, is the TCP port used by the web browser server. The default values are port 80 for HTTP and 443 for HTTPS. For instructions on how to change this setting, refer to "Configuring the Web Server" on page 683. The management software does not allow you to change the default port number of 23 for Telnet connections or 22 for SSH connections.

Deleting a TCP Connection

This procedure explains how you can use the TCP connections table to end a remote Telnet, SSH or web browser management session on a switch. This procedure is useful if a manager forgot to log out after ending a session or if you suspect that an unauthorized person is accessing the switch's management software.

Before performing this procedure, display the TCP table by performing the procedure "Displaying the TCP Connections Table" on page 367 and write down on paper the index number of the connection you want to end. A web browser management session can consist of more than one TCP connection.

You cannot delete the entries for the listening sockets for the Telnet, SSH, and web browser servers. To remove a listening socket entry from the table, disable the corresponding server.

To delete a TCP connection so as to end the corresponding Telnet or web browser management session, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. From the System Utilities menu, type **6** to select Networking Stack.

   The Networking Stack menu is shown in Figure 124 on page 361.

4. From the Networking Stack menu, type **7** to select Delete TCP Connection. The following prompt is displayed:

   ```
   Enter the TCP Connection Index: [0 to 65535] ->
   ```

5. Enter the index number of the TCP connection you want to delete from the table. You can enter only one index number at a time. To display the index numbers, refer to "Displaying the TCP Connections Table" on page 367.

   Deleting a TCP connection immediately ends the associated Telnet or web browser management session.

6. To delete additional TCP connection, repeat steps 4 and 5.

7. Return to the Main Menu.

Displaying the TCP Global Information Table

The TCP Global Information table displays TCP status and statistics. To view the table, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **9** to select System Utilities.

   The System Utilities menu is shown in Figure 6 on page 57.

3. From the System Utilities menu, type **6** to select Networking Stack.

   The Networking Stack menu is shown in Figure 124 on page 361.

4. From the Networking Stack menu, type **6** to select Display TCP Global Information.

   The Display TCP Global Information table is shown in Figure 129.

```
          Allied Telesyn Ethernet Switch AT-8524POE - AT-S62
                          Production Switch

User: Manager                                    11:20:02 02-Jan-2006

                     Display TCP Global Information

TCP MIB parameters, counters
----------------------------

RTO min (ms):           1000      RTO max (ms):              240000
Max connections:        30
Active Opens:           0         Passive Opens:             0
Attempt Fails:          0         Established Resets:        0
Current Established:    0
In Segs:                0         In Segs Error:             0
Out Segs:               0         Out Segs Retran:           0
Out Segs with RST:      0

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 129. Display TCP Global Information Table

This table is for viewing purposes only. The fields are defined here.

RTO min (ms) and RTO max (min)
Retransmit time algorithm parameters.

Max connections
The maximum number of TCP connections allowed.

Active Opens
The number of active TCP opens. Active opens initiate connections.

Passive Opens
The number of TCP passive opens. Passive opens are issued to wait for a connection from another host.

Attempt Fails
The number of failed connection attempts.

Established Resets
The number of connections established but have not been reset.

Current Established
The number of current connections.

In Segs
The number of segments received.

In Segs Error
The number of segments received with an error.

Out Segs
The number of segments transmitted.

Out Segs Retran
The number of segments retransmitted.

Out Segs with RST
The number of segments transmitted with the RST bit set.

Section III

SNMPv3 Operations

This section contains the following chapter:

Chapter 21
SNMPv3

This chapter provides a description of the AT-S62 implementation of the SNMPv3 protocol. In addition, it provides procedures that allow you to create and modify SNMPv3 users. The following sections are provided:

Note
Several SNMPv3 parameters appear only in the AT-S62 version 1.1.1 software.

SNMPv3 Overview

The SNMPv3 protocol builds on the existing SNMPv1 and SNMPv2c protocol implementation which is described in Chapter 5: “SNMPv1 and SNMPv2c Configuration” on page 89. In the SNMPv3 protocol, User-based Security Model (USM) authentication is implemented along with encryption, allowing you to configure a secure SNMP environment.

The SNMP terminology changes in the SNMPv3 protocol. In the SNMPv1 and SNMPv2c protocols, there are two actors in an SNMP network—a manager and an agent. A *manager* is a server that runs SNMP management software. The manager is often called the Network Management System (NMS). An *agent* is the SNMP software that runs on a network device, such as the AT-8500 Series switch. An NMS is responsible for querying, or polling, agents in the network. In addition, the agent sends messages to the NMS indicating events. In the AT-S62 implementation of SNMPv3, the switch sends trap and inform messages.

In SNMPv3, managers and agents are both called *entities*. Each entity consists of an Engine Id and SNMP applications. Each AT-8500 Series switch has a unique Engine ID number. The roles of authoritative entity and non-authoritative entity can change depending on the type of message that is sent. Consider the following three cases:

- The NMS sends an inform message to the switch. Once a network device (either an NMS or the switch) sends an inform message, the network device expects a response to this type of message. When the switch receives an inform message, then the switch is considered an authoritative entity. In this case, the NMS is the non-authoritative entity.
- If the switch sends a trap message (a type of message that does not expect a response), then the switch is considered the authoritative entity. In this case, the NMS is the non-authoritative entity.
- If the switch sends an inform message, then the NMS is considered the authoritative entity. In this case, the switch is the non-authoritative entity

The concept of entities is important because they help define an internal architecture for the SNMPv3 protocol—as opposed to just defining a set of messages. This new architecture makes the protocol more secure. For more details about the architecture, consult the SNMPv3 RFCs. For the SNMP RFCs supported by this release of the AT-S62 software, see “SNMP Management Session” on page 34.

With the SNMPv3 protocol, you create users, determine the protocol used for message authentication as well as determine if data transmitted between an SNMP agent and an NMS is encrypted. In addition, you have the ability to restrict user privileges by determining the user’s view of the Management Information Bases (MIBs). In this way, you restrict which

MIBs the user can display and modify. In addition, you can restrict the types of messages the switch can send on behalf of a user.

After you have created a user, you define SNMPv3 message notification. This consists of determining where messages are sent and what types of messages can be sent. This configuration is similar to the SNMPv1 and SNMPv2c configuration because you configure IP addresses of trap receivers, or hosts. In addition, with the SNMPv3 implementation you decide what types of messages can be sent.

This section further describes the features of the SNMPv3 protocol. The following subsections are included:

SNMPv3 Authentication Protocols

The SNMPv3 protocol supports two authentication protocols—HMAC-MD5-96 (MD5) and HMAC-SHA-96 (SHA). Both MD5 and SHA use an algorithm to generate a message digest. Each authentication protocol authenticates a user by checking the message digest. In addition, both protocols use keys to perform authentication. The keys for both protocols are generated locally using the Engine ID, a unique identifier that is assigned to each switch automatically, and the user password. You modify a key only by modifying the user password.

In addition, you have the option of assigning no user authentication. In this case, no authentication is performed for this user. Allied Telesyn does not recommend this configuration for security reasons.

Note
The keys generated by the MD5 and SHA protocols are specific to the SNMPv3 protocol. They have no relation to the SSL and SSH keys for encryption.

SNMPv3 Privacy Protocol

After you have configured an authentication protocol, you have the option of assigning a privacy protocol if you have the encrypted version of the AT-S62 software. In SNMPv3 protocol terminology, privacy is equivalent to encryption. Currently, the DES protocol is the only encryption protocol supported. The DES privacy protocol requires the authentication protocol to be configured as either MD5 or SHA.

If you assign a DES privacy protocol to a user, then you are also required to assign a privacy password. If you choose to not assign the privacy to DES, then SNMPv3 messages are sent in plain text format.

Note
You are able to configure the Privacy Protocol only if you are using the encrypted version of the AT-S62 software.

SNMPv3 MIB Views

The SNMPv3 protocol allows you to configure MIB views for users and groups. The MIB tree is defined by RFC 1155 (Structure of Management Information). See Figure 130.

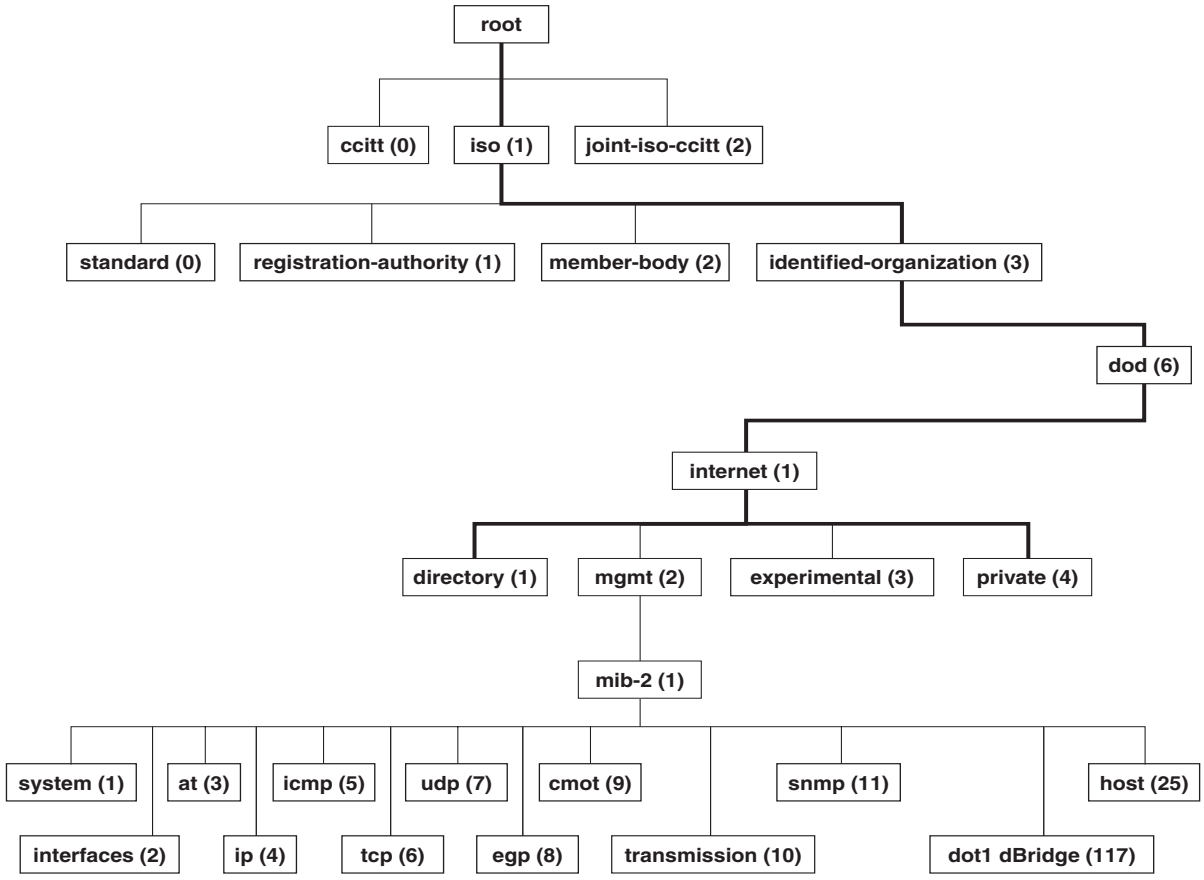

Figure 130. MIB Tree

The AT-S62 software supports the MIB tree, starting with the Internet MIBs, as defined by 1.3.6.1. There are two ways to specify a MIB view. You can enter the OID number of the MIB view or its equivalent text name. For example, to specify MIBs in the Internet view, you can enter the OID format “1.3.6.1” or the text name “internet.”

In addition, you can define a MIB view that the user can access or a MIB view that the user cannot access. When you want to permit a user to access a MIB view, you include a particular view. When you want to deny a user access to a MIB view, you exclude a particular view.

After you specify a MIB Subtree view you have the option of further restricting a view by defining a Subtree Mask. The relationship between a MIB Subtree View and a Subtree Mask is analogous to the relationship between an IP address and a subnet mask. The switch uses the subnet mask to determine which portion of an IP address represents the network address and which portion represents the node address. In a similar way, the Subtree Mask further refines the Subtree View and enables you to restrict a MIB view to a specific row of the OID MIB table. Naturally, you need a thorough understanding of the OID MIB table to define a Subtree Mask.

SNMPv3 Storage Types

Each SNMPv3 table entry has its own storage type. You can choose between NonVolatile storage which allows you to save the table entry or Volatile storage which does not allow you to save an entry. If you select the Volatile storage type, when you power off the switch your SNMPv3 configuration is lost and cannot be recovered.

At each SNMPv3 menu, you are prompted to configure a storage type. You do not have to configure the same storage type value for each table entry.

SNMPv3 Message Notification

When you generate an SNMPv3 message from the switch, there are three basic pieces of information included in the message:

- The type of message
- The destination of the message
- SNMP security information

To configure the type of message, you need to define if you are sending a Trap or Inform message. Basically, the switch expects the authoritative entity (or NMS) to respond to an Inform message. The switch does not expect the authoritative entity to respond to a Trap message. These two message types are defined in the SNMPv3 (RFC 2571-6).

To determine the destination of the message, you configure the IP address of the host. This configuration is similar to the SNMPv1 and SNMPv2c configuration.

The SNMP security information consists of information about the following:

- User
- View of the MIB Tree
- Security Level

- Security Model
- Authentication Level
- Privacy Protocol
- Group

To configure the SNMP security information, you associate a user and its related information—View, Security Level, Security Model, Authentication Level, Privacy Protocol and Group—with the type of message and the host IP address.

SNMPv3 Tables

The SNMPv3 configuration is neatly divided into configuring SNMPv3 user information and configuring the message notification. You must configure all seven tables to successfully configure the SNMPv3 protocol. You use the following tables for user configuration:

- Configure SNMPv3 User Table
- Configure SNMPv3 View Table
- Configure SNMPv3 Access Table
- Configure SNMPv3 SecurityToGroup Table

First, you create a user in the Configure SNMPv3 User Table. Then you define the MIB view this user has access to in the Configure SNMPv3 View Table. To configure a security group and associate a MIB view to a security group, you configure the Configure SNMPv3 Access Table. Finally, configure the Configure SNMPv3 SecurityToGroup Menu to associate a user to a security group. See Figure 131 for an illustration of how the user configuration tables are linked.

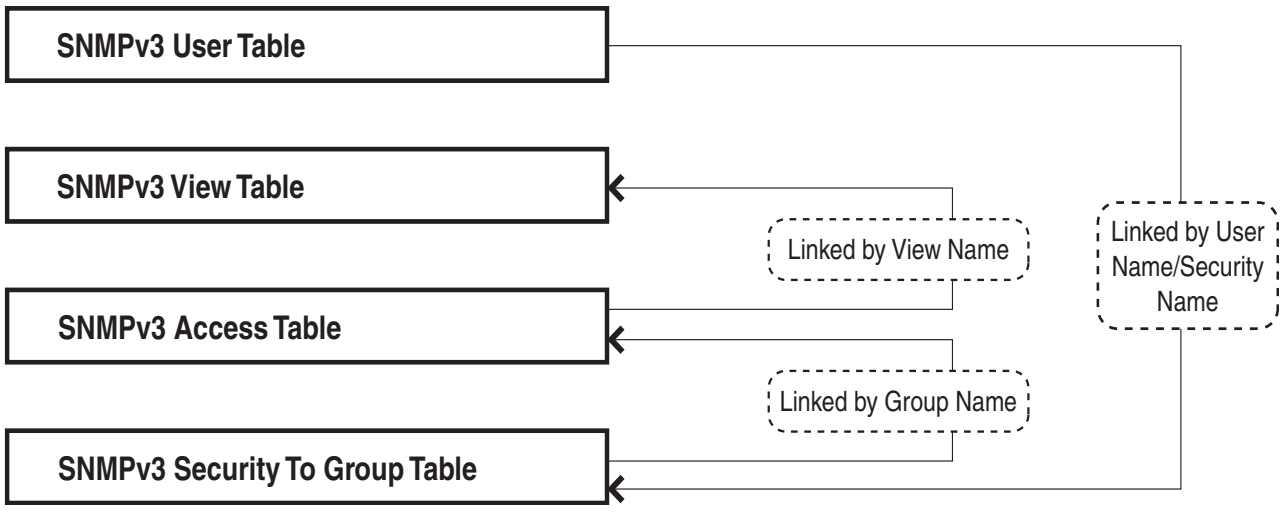

Figure 131. SNMPv3 User Configuration Process

In general, you focus on configuring security groups and then add and delete users from the groups as needed. For example, you may want to have two groups—one for manager privileges and a second one for operator privileges. See Appendix B, "SNMPv3" on page 375 for an example of manager and operator configurations.

After you configure an SNMPv3 user, you need to configure SNMPv3 message notification. This configuration is accomplished with the following tables:

- Configure SNMPv3 Notify Table
- Configure SNMPv3 Target Address Table
- Configure SNMPv3 Target Parameters Table

You start the message notification configuration by defining the type of message you want to send with the SNMPv3 Notify Table. Then you define a IP address that is used for notification in the Configure SNMPv3 Target Address Table. This is the IP address of the SNMPv3 manager. Finally, you associate the trap information with a user by configuring the Configure SNMPv3 Target Parameters Table.

See Figure 132 for an illustration of how the message notification tables are linked.

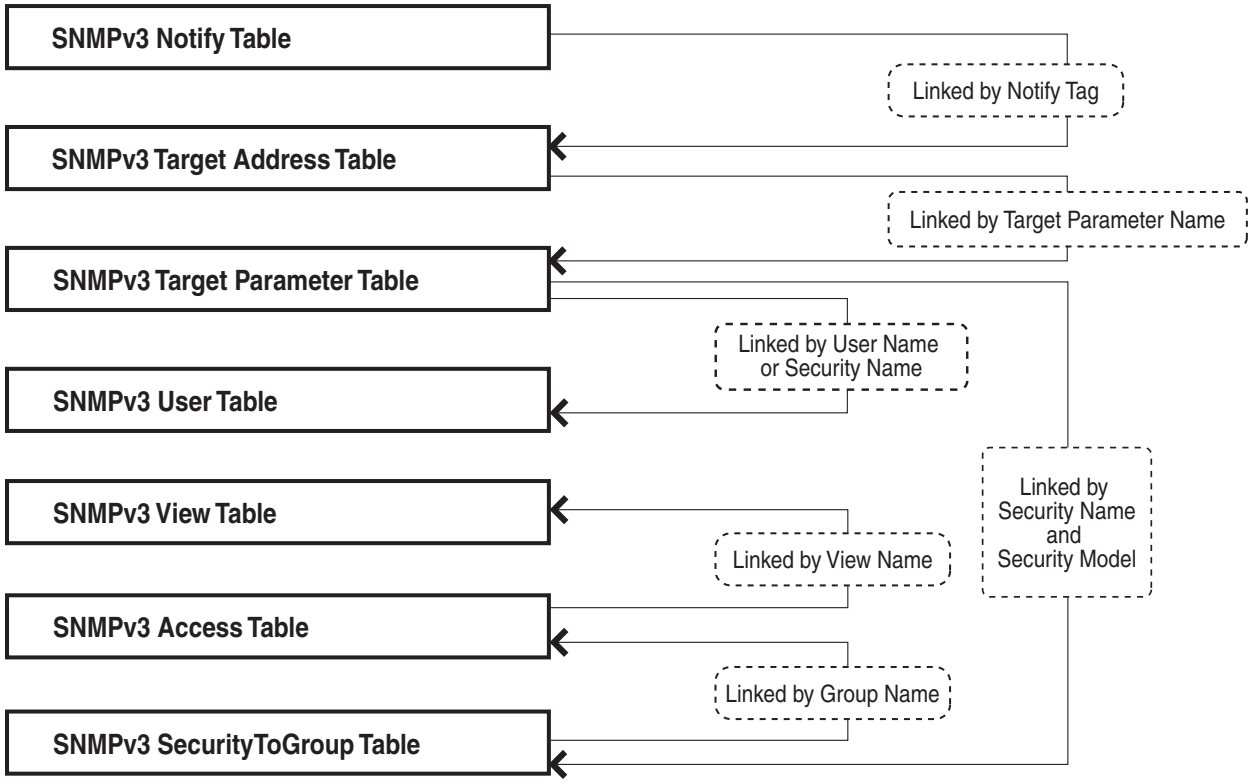

Figure 132. SNMPv3 Message Notification Process

For a more detailed description of the SNMPv3 Tables, see the following subsections:

SNMPv3 User Table

The Configure SNMPv3 User Table menu allows you to create an SNMPv3 user and provides the options of configuring authentication and privacy protocols. With an authentication protocol configured, users are authenticated when they send and receive messages. In addition, you can configure a privacy protocol and password so messages a user sends and receives are encrypted. The DES privacy algorithm uses the privacy password and the Engine ID to generate a key that is used for encryption. Lastly, you can configure a storage type for this table entry which allows you to save this user and its related configuration to flash memory.

SNMPv3 View Table

The Configure SNMPv3 View Table Menu allows you to create a view of the MIB OID Table. First, you configure a view of a subtree. Then you have the option of configuring a Subtree Mask that further refines the subtree view. For example, you can use a Subtree Mask to restrict a user’s view to one row of the MIB OID Table. In addition, you can chose to include or exclude a view. As a result, you can let a user see a particular view or prevent a user from seeing a particular view. Lastly, you can configure a storage type for this table entry which allows you to save this view to flash memory.

SNMPv3 Access Table

The Configure SNMPv3 Access Table Menu allows you to configure a security group. After you create a security group, you assign a set of users with the same access privileges to this group using the SNMPv3 SecurityToGroup Table. It is useful to consider the types of groups you want to create and the types of access privileges each group will have. In this way, it is easy to keep track of your users as belonging to one or two groups.

For each group, you can assign read, write, and notify views of the MIB table. The views you assign here have been previously defined in the

Configure SNMPv3 View Table Menu. For example, the Read View allows group members to view the specified portion of the OID MIB table. The Write View allows group members to write to, or modify, the MIBs in the specified MIB view. The Notify View allows group members to send trap messages defined by the MIB view. Lastly, you can configure a storage type for this table entry which allows you to save this view to flash memory.

SNMPv3 SecurityToGroup Table

The Configure SNMPv3 SecurityToGroup Table Menu allows you to associate a User Name with a security group called a Group Name. The User Name is previously configured with the Configure SNMPv3 User Table Menu. The security group is previously configured with the Configure SNMPv3 Access Table Menu. Lastly, you can configure a storage type for this table entry which allows you to save the entry to flash memory.

SNMPv3 Notify Table

The Configure SNMPv3 Notify Table Menu allows you to define the type of message that is sent from the switch (or non-authoritative entity) to the authoritative entity. You have the option of defining the message type as either an Inform or a Trap message. When a switch sends an Inform message, it expects a response from the authoritative entity. In comparison, when the switch sends a Trap message, it does not require a response from the authoritative entity.

In addition, you define a Notify Tag that links an SNMPv3 Notify Table entry to the host IP address defined in the Configure SNMPv3 Target Address Table Menu. Lastly, you can configure a storage type for this table entry which allows you to save the entry to flash memory.

SNMPv3 Target Address Table

The Configure SNMPv3 Target Address Table Menu allows you to configure the IP address of the host. Also, in an SNMPv3 Target Address Table entry, you configure the values of the Tag List parameter with the previously defined Notify Tag parameter values. The Notify Tag parameter is configured in the Configure SNMPv3 Notify Table. In this way, the Notify and Target Address tables are linked. Lastly, you can configure a storage type for this table entry which allows you to save the entry to flash memory.

SNMPv3 Target Parameters Table

The Configure SNMPv3 Target Parameters Table Menu allows you to define which user can send messages to the host IP address defined in the Configure SNMPv3 Target Address Table. The user and its associated information is previously configured in the Configure SNMPv3 User Table,

SNMPv3 View Table, SNMPv3 Access Table, and SNMPv3 SecurityToGroup Table. Lastly, you can configure a storage type for this table entry which allows you to save the entry to flash memory.

SNMPv3 Community Table

The Configure SNMPv3 Community Table Menu allows you to configure SNMPv1 and SNMPv2c communities. If you are going to use the SNMPv3 Tables to configure SNMPv1 and SNMPv2c communities, start with the SNMPv3 Community Table. See “Configuring the SNMPv3 Community Table” on page 462.

Note
Allied Telesyn recommends that you use the procedures described in Chapter 5: “SNMPv1 and SNMPv2c Configuration” on page 89 to configure the SNMPv1 and SNMPv2c protocols.

SNMPv3 Configuration Example

You may want to have two classes of SNMPv3 users—Managers and Operators. In this scenario, you would configure one group, called Managers, with full access privileges. Then you would configure a second group, called Operators, with monitoring privileges only. For a detailed example of this configuration, see Appendix B, “SNMPv3 Configuration Examples” on page 799.

Configuring the SNMPv3 Protocol

This section describes how to configure the SNMPv3 protocol using the SNMPv3 Tables. To successfully configure this protocol, you must perform the procedures in the order given. For overview information about SNMPv3, see the "SNMPv3 Overview" on page 376.

In order to allow an NMS to access the switch, you need to enable SNMP access. In addition, to allow the switch to send a trap when it receives a request message, you need to enable authentication failure traps. See "Enabling or Disabling SNMP Management" on page 93.

The following SNMPv3 tables are described in this chapter:

- "Configuring the SNMPv3 User Table" on page 386
- "Configuring the SNMPv3 View Table" on page 396
- "Configuring the SNMPv3 Access Table" on page 405
- "Configuring the SNMPv3 SecurityToGroup Table" on page 421
- "Configuring the SNMPv3 Notify Table" on page 429
- "Configuring the SNMPv3 Target Address Table" on page 436
- "Configuring the SNMPv3 Target Parameters Table" on page 449
- "Configuring the SNMPv3 Community Table" on page 462

The SNMPv3 User, View, Access, and SecurityToGroup tables are concerned with setting up a user, determining authentication and privacy, and associating a user to a security group. The SNMPv3 Notify, Target Address, and Target Parameters tables are concerned with message notification. You use the SNMPv3 Community Table to configure SNMPv1 and SNMPv2 communities.

Due to the complexity of the SNMPv3 configuration, Allied Telesyn recommends that you configure the SNMPv3 protocol with the procedures listed above, in the order they are listed. However, it is possible to configure the SNMPv3 protocol using the above procedures in any order.

Note
New entries to the SNMPv3 tables are added alphabetically.

Configuring the SNMPv3 User Table

This section contains a description of the SNMPv3 User Table and how to create, delete, and modify table entries. Configure the SNMPv3 User Table first. Creating this table, allows you to create an entry in an SNMPv3 User Table for a User Name. In addition, this table allows you to associate a User Name with the following parameters:

- Authentication Protocol
- Authentication Password
- Privacy Protocol
- Privacy Password

> **Note**
> You are prompted to configure the Privacy Protocol only if you are using the encrypted version of the AT-S62 software.

There are three functions you can perform with the SNMPv3 User Table.

- "Creating an SNMPv3 User Table Entry" on page 386
- "Deleting an SNMPv3 User Table Entry" on page 390
- "Modifying an SNMPv3 User Table Entry" on page 391

Creating an SNMPv3 User Table Entry

To create an entry in the SNMPv3 User Table, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **5** to select SNMP Configuration.

   The SNMP Configuration menu is shown in Figure 17 on page 93.

3. From the SNMP Configuration menu, type **5** to select Configure SNMPv3 Table.

The Configure SNMPv3 Table Menu is shown in Figure 133.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006
                       Configure SNMPv3 Table
1 - SNMP Engine..............80:00:00:CF:31:00:30:84:FD:57:DA
2 - Configure SNMPv3 User Table
3 - Configure SNMPv3 View Table
4 - Configure SNMPv3 Access Table
5 - Configure SNMPv3 SecurityToGroup Table
6 - Configure SNMPv3 Notify Table
7 - Configure SNMPv3 Target Address Table
8 - Configure SNMPv3 Target Parameters Table
9 - Configure SNMPv3 Community Table

R - Return to Previous Menu

Enter your selection?
```

Figure 133. Configure SNMPv3 Table Menu

Note
The SNMP Engine field is a read-only field. You cannot change the setting. The field displays the SNMP engine identifier that is assigned automatically to the switch.

4. From the Configure SNMPv3 Table Menu, type **2** to select Configure SNMPv3 User Table. The Configure SNMPv3 User Table Menu is shown in Figure 134.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  00:14:33 15-Jan-2006
                    Configure SNMPv3 User Table
Engine ID ................. 80:00:00:CF:03:00:30:84:FD:57:DA
User Name ................. jenny
Authentication Protocol ... MD5
Privacy Protocol .......... DES
Storage Type .............. NonVolatile
Row Status ................ Active

1 - Create SNMPv3 Table Entry
2 - Delete SNMPv3 Table Entry
3 - Modify SNMPv3 Table Entry

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 134. Configure SNMPv3 User Table Menu

5. To create a new user table, type **1** to select Create SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter User (Security) Name:
   ```

6. Enter a descriptive name of the user.

   You can enter a name that consists of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter Authentication Protocol [M-MD5, S-SHA,
   N-None]:
   ```

7. Enter one of the following:

   M-MD5
   This value represents the MD5 authentication protocol. With this selection, users are authenticated with the MD5 authentication protocol after a message is received. This algorithm generates the message digest. The user is authenticated when the authentication protocol checks the message digest. With the MD5 selection, you can configure a Privacy Protocol.

   S-SHA
   This value represents the SHA authentication protocol. With this selection, users are authenticated with the SHA authentication protocol after a message is received. This algorithm generates the message digest. The user is authenticated when the authentication protocol checks the message digest. With the SHA selection, you can configure a Privacy Protocol.

   N-None
   This value represents no authentication protocol. When messages are received, users are not authenticated. With the None selection, you cannot configure a Privacy Protocol.

   If you select NONE, you are prompted for the Storage Type. Go to Step 13.

   If you select MD5 or SHA, the following prompt is displayed:

   ```
   Enter Authentication Password:
   ```

8. Enter an authentication password of up to 32-alphanumeric characters and press Return.

   You are prompted to re-enter the password.

The following prompt is displayed:

```
Enter Privacy Protocol [D-DES, N-None]:
```

Note
If you have the non encrypted version of the AT-S62 software, then the Privacy Protocol field is read-only.

Note
You can only configure the Privacy Protocol if you have configured the Authentication Protocol with the MD5 or SHA values.

9. Select one of the following options:

   D -DES
   Select this value to make the DES privacy (or encryption) protocol the privacy protocol for this User Table entry. With this selection, messages transmitted between the host and the switch are encrypted with the DES protocol.

   N -None
   Select this value if you do not want a privacy protocol for this User Table entry. With this selection, messages transmitted between the host and the switch are not encrypted.

   If you select NONE, you are prompted for the Storage Type. Go to Step 13.

   If you select DES, the following prompt is displayed:

   ```
   Enter Privacy Password:
   ```

10. Enter a privacy password of up to 32-alphanumeric characters.

    You are prompted to re-enter the password.

    The following prompt is displayed:

    ```
    Enter Storage Type [V-Volatile, N-NonVolatile]:
    ```

11. Select one of the following storage types for this table entry:

    V - Volatile
    Select this storage type if you do not want the ability to save an entry in the SNMPv3 User Table to nonvolatile memory. After making changes to an SNMPv3 User Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

    N-NonVolatile
    Select this storage type if you want the ability to save an entry in the SNMPv3 User Table to nonvolatile memory. After making changes to

an SNMPv3 User Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

Note
The Row Status parameter is a read-only field in the Telnet and Local interfaces. The Active value indicates the SNMPv3 User Table entry takes effect immediately.

12. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an SNMPv3 User Table Entry

You may want to delete an entry from the SNMPv3 User Table. When you delete an entry in the SNMPv3 User Table, there is no way to undelete, or recover it.

To delete an entry in the SNMPv3 User Table, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133.

2. From the Configure SNMPv3 Table Menu, type **2** to select Configure SNMPv3 User Table.

   The SNMPv3 User Table is shown in Figure 134.

3. From the SNMPv3 User Table, type **2** to select Delete SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter User (Security) Name:
   ```

4. Enter the User Name of the User Table entry you want to delete.

   The following prompt is displayed:

   ```
   Do you want to delete this table entry? (Y/N): [Yes/No]->
   ```

5. Enter **Y** to delete the user or **N** to save the user.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an SNMPv3 User Table Entry

This section describes how to modify parameters in an SNMPv3 Notify Table entry. See the following procedures:

- "Modifying the Authentication Protocol and Password" on page 391
- "Modifying the Privacy Protocol and Password" on page 393
- "Modifying the Storage Type" on page 394

Modifying the Authentication Protocol and Password

To modify the Authentication Protocol and Password in an SNMPv3 User Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133.

2. From the Configure SNMPv3 Table Menu, type **2** to select Configure SNMPv3 User Table.

   The SNMPv3 User Table is shown in Figure 134.

3. From the SNMPv3 User Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 User Table is shown in Figure 135.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  00:14:33 15-Jan-2006

                       Modify SNMPv3 User Table

Engine ID ................. 80:00:00:CF:03:00:30:84:FD:57:DA
User Name ................. wilson
Authentication Protocol ... SHA
Privacy Protocol .......... DES
Storage Type .............. NonVolatile
Row Status ................ Active

1 - Set Authentication Protocol & Password
2 - Set Privacy Protocol & Password
3 - Set Storage Type

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 135. Modify SNMPv3 User Table Menu

4. To change the authentication protocol and password, type **1** to select Set Authentication Protocol & Password.

   The following prompt is displayed:

   ```
   Enter User Name:
   ```

5. Enter the User Name of the User Table you want to modify.

   The following prompt is displayed:

   ```
   Enter Authentication Protocol [M-MD5, S-SHA,
   N-None]:
   ```

6. Enter one of the following:

   M-MD5
   This value represents the MD5 authentication protocol. With this selection, users are authenticated with the MD5 authentication protocol after a message is received. This algorithm generates the message digest. The user is authenticated when the authentication protocol checks the message digest. With the MD5 selection, you can configure a Privacy Protocol.

   S-SHA
   This value represents the SHA authentication protocol. With this selection, users are authenticated with the SHA authentication protocol after a message is received. This algorithm generates the message digest. The user is authenticated when the authentication protocol checks the message digest. With the SHA selection, you can configure a Privacy Protocol.

   N-None
   This value represents no authentication protocol. When messages are received, users are not authenticated. With the None selection, you cannot configure a Privacy Protocol.

   If you select None, go to step 9.

   If you select MD5 or SHA, the following prompt is displayed:

   ```
   Enter Authentication Password:
   ```

7. Enter an authentication password of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Re-enter Authentication password:
   ```

8. Re-enter the password.

   The following message is displayed:

   ```
   Authentication protocol algorithm has been changed.
   ```

The following prompt is displayed:

```
Please enter privacy password to regenerate privacy key.
```

9. Enter the Privacy Password for this User Name.

   The following prompt is displayed:

   ```
   Re-enter Privacy password:
   ```

10. Re-enter the password.

11. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Privacy Protocol and Password

To modify the Privacy Protocol and Password in an SNMPv3 User Table entry, perform the following procedure.

> **Note**
> You can only configure the Privacy Protocol if you have configured the Authentication Protocol with the MD5 or SHA values.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133.

2. From the Configure SNMPv3 Table Menu, type **2** to select Configure SNMPv3 User Table.

   The SNMPv3 User Table is shown in Figure 134 on page 387.

3. From the SNMPv3 User Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Table Menu is shown in Figure 135 on page 391.

4. Type **2** to select Privacy Protocol & Password.

   The following prompt is displayed:

   ```
   Enter User (Security) Name:
   ```

5. Enter the User Name.

   The following prompt is displayed:

   ```
   Enter Privacy Protocol [D-DES, N-None]:
   ```

6. Choose one of the following Privacy Protocols:

D -DES
Select this value to make the DES privacy (or encryption) protocol the privacy protocol for this User Table entry. With this selection, messages transmitted between the host and the switch are encrypted with the DES protocol.

N -None
Select this value if you do not want a privacy protocol for this User Table entry. With this selection, messages transmitted between the host and the switch are not encrypted.

If you select None, proceed to step 9.

If you select DES, the following prompt is displayed:

```
Enter Privacy Password:
```

7. Enter a privacy password of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Re-enter Authentication password:
   ```

8. Re-enter the password.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Storage Type

To modify the Storage Type in an SNMPv3 User Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **2** to select Configure SNMPv3 User Table.

   The SNMPv3 User Table is shown in Figure 134 on page 387.

3. From the SNMPv3 User Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Table Menu is shown in Figure 135 on page 391.

4. To change the storage type, type **3** to select Set Storage Type.

   The following prompt is displayed:

```
Enter User (Security) Name:
```

5. Enter the User Name.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-NonVolatile]:
   ```

6. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 User Table to nonvolatile memory. After making changes to an SNMPv3 User Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 User Table to nonvolatile memory. After making changes to an SNMPv3 User Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the SNMPv3 View Table

This section contains a description of the SNMPv3 View Table and how to create, delete, and modify table entries. Creating this table, allows you to specify a view using the following parameters:

- Subtree OID
- Subtree Mask
- MIB OID Table View

To configure the SNMPv3 View Table, you need to be very familiar with the MIB tree. You can be very specific about the view a user can or cannot access—down to a column or row of the tree. AT-S62 supports the Internet subtree of the MIB tree. See RFC 2575 for detailed information about defining a view.

There are three functions you can perform with the SNMPv3 User Table.

- "Creating an SNMPv3 View Table Entry" on page 396
- "Deleting an SNMPv3 View Table Entry" on page 399
- "Modifying an SNMPv3 View Table Entry" on page 400

Creating an SNMPv3 View Table Entry

To create an entry in the SNMPv3 View Table, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **3** to select Configure SNMPv3 View Table.

The Configure SNMPv3 View Table Menu is shown in Figure 136.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                 00:14:33 15-Jan-2006

                    Configure SNMPv3 View Table

View Name ................. internet
Subtree OID ............... 1.3.6.1
Subtree Mask ..............
View Type ................. Included
Storage Type .............. NonVolatile
Row Status ................ Active

1 - Create SNMPv3 Table Entry
2 - Delete SNMPv3 Table Entry
3 - Modify SNMPv3 Table Entry

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 136. Configure SNMPv3 View Table Menu

3. From the Configure SNMPv3 View Table Menu, type **1** to select Create SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter View Name:
   ```

4. Enter a descriptive name of this View.

   Enter a unique name of up to 32-alphanumeric characters.

   Note
   The "defaultViewAll" value is the default entry for the SNMPv1 and SNMPv2c configuration. You cannot use the default value for an SNMPv3 View Table entry.

   The following prompt is displayed:

   ```
   Enter View Subtree (OID format/Text Name):
   ```

5. Enter subtree that this view will or will not be permitted to display.

   You can enter either a numeric value in hex format or the equivalent text name. For example, the OID hex format for TCP/IP is:

   ```
   1.3.6.1.2.1.6
   ```

The text format is for TCP/IP is:

```
tcp
```

The following prompt is displayed:

```
Enter Subtree Mask (Hex format):
```

6. Enter a subtree mask.

   This is an optional parameter that is used to further refine the value in the View Subtree parameter. This parameter is in binary format.

   The View Subtree parameter defines a MIB View and the Subtree Mask further restricts a user's view, for example, to a specific row of the MIB tree. The value of the Subnet Mask parameter is dependent on the subtree you select. See RFC 2575 for detailed information about defining a subnet mask.

   The following prompt is displayed:

   ```
   Enter View Type [I-Included, E-Excluded]:
   ```

7. Enter one of the following view types:

   I - Included
   Enter this value to permit the View Name to see the subtree specified above.

   E - Excluded
   Enter this value to not permit the View Name to see the subtree specified above.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-NonVolatile]:
   ```

8. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 View Table to the configuration file. After making changes to an SNMPv3 View Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 View Table to the configuration file. After making changes to an SNMPv3 View Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

Note
The Row Status parameter is a read-only field in the Telnet and Local interfaces. The Active value indicates the SNMPv3 View Table entry takes effect immediately.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an SNMPv3 View Table Entry

You may want to delete an entry from the SNMPv3 View Table. After you delete an SNMPv3 View Table entry, there is no way to undelete, or recover it.

To delete an entry in the SNMPv3 View Table, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **3** to select Configure SNMPv3 View Table.

   The SNMPv3 View Table is shown in Figure 136 on page 397.

3. From the SNMPv3 View Table, type **2** to select Delete SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter View Name:
   ```

4. Enter the View Name of the View Table entry you want to delete.

   The following prompt is displayed:

   ```
   Enter View Subtree (OID format/Text Name):
   ```

5. Enter the subtree for this view.

   ```
   Do you want to delete this table entry? (Y/N): [Yes/No]->
   ```

6. Enter **Y** to delete the view or **N** to save the view.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an SNMPv3 View Table Entry

This section describes how to modify parameters in an SNMPv3 Notify Table entry. See the following procedures:

- “Modifying a Subtree Mask” on page 400
- “Modifying a View Type” on page 401
- “Modifying a Storage Type” on page 403

Modifying a Subtree Mask

To modify the Subtree Mask parameter in an SNMPv3 View Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **3** to select Configure SNMPv3 View Table.

   The Configure SNMPv3 View Table Menu is shown in Figure 136 on page 397.

3. From the Configure SNMPv3 View Table Menu, type **3** to select Modify SNMPv3 Table Entry. The Modify SNMPv3 View Table Menu is shown in Figure 137.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                  11:20:02 02-Jan-2006

                     Modify SNMPv3 View Table

View Name ................. tcp
Subtree OID ............... 1.3.6.1.2.1.6
Subtree Mask .............. ff:ff
View Type ................. Included
Storage Type .............. NonVolatile
Row Status ................ Active

1 - Set Subtree Mask
2 - Set View Type
3 - Set Storage Type

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 137. Modify SNMPv3 View Table Menu

4. To modify the Subtree Mask for this view, type **1** to select Set Subtree Mask.

   The following prompt is displayed:

   ```
   Enter View Name:
   ```

5. Enter an existing View Name.

   The following prompt is displayed:

   ```
   Enter View Subtree (OID format/Text Name):
   ```

6. Enter Subtree that this view will or will not be permitted to display.

   You can enter either a numeric value in hex format or the equivalent text name. For example, the OID hex format for TCP/IP is:

   ```
   1.3.6.1.2.1.6
   ```

   The text format is for TCP/IP is:

   ```
   tcp
   ```

   The following prompt is displayed:

   ```
   Enter Subtree Mask (Hex format):
   ```

7. Enter a Subtree Mask.

   This is an optional parameter that is used to further refine the value in the View Subtree parameter. This parameter is in binary format.

   The View Subtree parameter defines a MIB View and the Subtree Mask further restricts a user's view, for example, to a specific row of the MIB tree. The value of the Subnet Mask parameter is dependent on the subtree you select. See RFC 2575 for detailed information about defining a subnet mask.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a View Type

To modify the View Type parameter in an SNMPv3 View Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **3** to select Configure SNMPv3 View Table.

   The Configure SNMPv3 View Table Menu is shown in Figure 136 on page 397.

3. From the Configure SNMPv3 View Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Table Menu is shown in Figure 137 on page 400.

4. To modify the View Type, type **2** to select Set View Type.

   The following prompt is displayed:

   ```
   Enter View Name:
   ```

5. Enter a View Name that was previously configured.

   The following prompt is displayed:

   ```
   Enter View Subtree (OID format/Text Name):
   ```

6. Enter the View Subtree value for this View Name.

   You can enter either a numeric value in hex format or the equivalent text name. For example, the OID hex format for TCP/IP is:

   ```
   1.3.6.1.2.1.6
   ```

   The text format is for TCP/IP is:

   ```
   tcp
   ```

   The following prompt is displayed:

   ```
   Enter View Type [I-Included, E-Excluded]:
   ```

7. Choose one of the following view types:

   I - Included
   Enter this value to permit the View Name to see the subtree specified above.

   E - Excluded
   Enter this value to not permit the View Name to see the subtree specified above.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Storage Type

To modify the Storage Type parameter in an SNMPv3 View Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **3** to select Configure SNMPv3 View Table.

   The Configure SNMPv3 View Table Menu is shown in Figure 136 on page 397.

3. From the Configure SNMPv3 View Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Table Menu is shown in Figure 137 on page 400.

4. To modify the storage type, type **3** to select Set Storage Type.

   The following prompt is displayed:

   ```
   Enter View Name:
   ```

5. Enter the View Name you want to modify.

   The following prompt is displayed:

   ```
   Enter View Subtree (OID format/Text Name):
   ```

6. Enter the View Subtree for this View Name.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-Nonvolatile]:
   ```

7. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 View Table to the configuration file. After making changes to an SNMPv3 View Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 View Table to the configuration file. After making changes to an SNMPv3 View Table entry with a NonVolatile storage type, the **S** -

Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the SNMPv3 Access Table

This section contains a description of the SNMPv3 Access Table and how to create, delete, and modify table entries. The SNMPv3 Access Table allows you to configure a security group. Each user must belong to a security group. After you have configured a security group, use the SecurityToGroup Table to assign users to security groups. See "Creating an SNMPv3 SecurityToGroup Table Entry" on page 421.

For each security group, you can assign the following attributes:

- a Security Model (SNMPv1, SNMPv2c, SNMPv3)
- Read, write, and notify views
- A security level
- A storage type

Before you begin this procedure, you will need to configure entries in the View Table. These values are used to configure the Read, Write, and Notify View parameters in this procedure. See "Configuring the SNMPv3 View Table" on page 396.

There are three functions you can perform with the SNMPv3 Access Table.

- "Creating an SNMPv3 Access Table Entry" on page 405
- "Deleting an SNMPv3 Access Table Entry" on page 409
- "Modifying an SNMPv3 Access Table Entry" on page 411

Creating an SNMPv3 Access Table Entry

To create an entry in the SNMPv3 Access Table, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **4** to select Configure SNMPv3 Access Table.

The Configure SNMPv3 Access Table Menu is shown in Figure 138.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                 00:14:33 15-Jan-2006

                      Configure SNMPv3 Access Table

Group Name .... softwareengineering        Security Model . v3
Context Prefix.                            Security Level . AuthPriv
Read View...... internet                   Context Match .. Exact
Write View .... tcp                        Storage Type ... NonVolatile
Notify View ... tcp                        Row Status ..... Active

1 - Create SNMPv3 Table Entry
2 - Delete SNMPv3 Table Entry
3 - Modify SNMPv3 Table Entry

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 138. Configure SNMPv3 Access Table Menu

3. To create a group in the SNMPv3 Access Table, type **1** to select Create SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Group Name:
   ```

4. Enter a descriptive name of the group. The Group Name can consist of up to 32-alphanumeric characters.

   The Group Name can consist of up to 32-alphanumeric characters.

   You are not required to enter a unique value here because the SNMPv3 Access Table entry is index with the Group Name, Security Model, and Security Level parameter values. However, unique group names makes it easier to tell the groups apart.

   There are four default values for this field:

   - defaultV1GroupReadOnly
   - defaultV1GroupReadWrite
   - defaultV2cGroupReadOnly
   - defaultV2cGroupReadWrite

   These values are reserved for SNMPv1 and SNMPv2c implementations.

Note
The Context Prefix and the Context Match fields are a read only fields. The Context Prefix field is always set to null. The Context Match field is always set to exact.

The following prompt is displayed:

```
Enter Security Model [1-v1, 2-v2c, 3-v3]:
```

5. Select one of the following SNMP protocols as the Security Model for this Group Name.

   1-v1
   Select this value to associate the Group Name with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Group Name with the SNMPv2c protocol.

   3-v3
   Select this value to associate the Group Name with the SNMPv3 protocol. The SNMPv3 protocol allows you to configure the group to authenticate SNMPv3 users and encrypt messages.

   The following prompt is displayed:

   ```
   Enter Security Level [N-NoAuthNoPriv, A-AuthNoPriv,
   P-AuthPriv]:
   ```

6. Select one of the following security levels:

   N-NoAuthNoPriv
   This option represents no authentication and no privacy protocol. Select this security level if you do not want to authenticate SNMP users and you do not want to encrypt messages using a privacy protocol. This security level provides the least security.

   Note
   If you have selected SNMPv1 or SNMPv2c, N-NoAuthNoPriv is the only security level you can select.

   A-AuthNoPriv
   This option represents authentication, but no privacy protocol. Select this security level if you want to authenticate SNMP users, but you do not want to encrypt messages using a privacy protocol.You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   P-AuthPriv
   This option represents authentication and the privacy protocol. Select

this security level to encrypt messages using a privacy protocol and authenticate SNMP users. This level provides the greatest level of security. You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

The following prompt is displayed:

```
Enter Read View Name:
```

7. Enter a value that you configured with the View Name parameter in the SNMPv3 View Table.

   A Read View Name allows the users assigned to this Group Name to view the information specified by the View Table entry. This value does not need to be unique.

   The following prompt is displayed:

   ```
   Enter Write View Name:
   ```

8. Enter a value that you configured with the View Name parameter in the SNMPv3 View Table.

   A Write View Name allows the users assigned to this Security Group to write, or modify, the information in the specified View Table. This value does not need to be unique.

   The following prompt is displayed:

   ```
   Enter Notify View Name:
   ```

9. Enter a value that you configured with the View Name parameter in the SNMPv3 View Table.

   A Notify View Name allows the users assigned to this Group Name to send traps permitted in the specified View. This value does not need to be unique.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-NonVolatile]:
   ```

10. Select one of the following storage types for this table entry:

    V - Volatile
    Select this storage type if you do not want the ability to save an entry in the SNMPv3 Access Table to the configuration file. After making changes to an SNMPv3 Access Table entry with a Volatile storage type, the
    S - Save Configuration Changes option does not appear on the Main Menu.

    N-NonVolatile
    Select this storage type if you want the ability to save an entry in the

SNMPv3 Access Table to the configuration file. After making changes to an SNMPv3 Access Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

Note
The Row Status parameter is a read-only field in the Telnet and Local interfaces. The Active value indicates the SNMPv3 Access Table entry will take effect immediately.

11. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an SNMPv3 Access Table Entry

You may want to delete an entry from the SNMPv3 Access Table. After you delete an SNMPv3 Access Table, there is no way to undelete, or recover, it.

To delete an entry in the SNMPv3 Access Table, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **4** to select Configure SNMPv3 Access Table.

   The SNMPv3 Access Table is shown in Figure 138 on page 406.

   Note
   To display a particular Group Name and its associated parameters from the Configure SNMPv3 Access Table Menu, type **N** to display the Next Page and **P** to display the previous page.

3. From the SNMPv3 Access Table, type **2** to select Delete SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Group Name:
   ```

4. Enter the Group Name that you want to delete.

   The following prompt is displayed:

   ```
   Enter Security Model [1-v1, 2-v2c, 3-v3]:
   ```

5. Enter the Security Model of this Group Name.

   Select one of the following security levels:

   1-v1
   Select this value to associate the Group Name with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Group Name with the SNMPv2c protocol.

   3-v3
   Select this value to associate the Group Name with the SNMPv3 protocol. The following prompt is displayed:

   ```
   Enter the Security Level [N-NoAuthNoPriv,
   A-AuthNoPriv, P-AuthPriv]:
   ```

6. Enter the Security Level of this Group Name.

   Select one of the following Security Levels:

   N-NoAuthNoPriv
   This option represents no authentication and no privacy protocol. Select this security level if you do not want to authenticate SNMP users and you do not want to encrypt messages using a privacy protocol. This security level provides the least security.

   Note
   If you have selected SNMPv1 or SNMPv2c, N-NoAuthNoPriv is the only security level you can select.

   A-AuthNoPriv
   This option represents authentication, but no privacy protocol. Select this security level if you want to authenticate SNMP users, but you do not want to encrypt messages using a privacy protocol.You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   P-AuthPriv
   This option represents authentication and the privacy protocol. Select this security level to encrypt messages using a privacy protocol and authenticate SNMP users. This level provides the greatest level of security. You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   The following prompt is displayed:

   ```
   Do you want to delete this table entry? (Y/N): [Yes/No]->
   ```

7. Enter **Y** to delete the view or **N** to save the view.

The following prompt is displayed:

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an SNMPv3 Access Table Entry

This section describes how to modify parameters in an SNMPv3 Access Table entry. For each entry in the SNMPv3 Access Table, you can modify the following parameters:

- Read View Name
- Write View Name
- Notify View Name
- Storage Type

Configure the values of the Read View Name, Write View Name, and Notify View Name parameters with values previously configured with the View Name parameter in the SNMPv3 View Table. This is the only way to associate a Group Name with these Views. See "Creating an SNMPv3 View Table Entry" on page 396.

See the following procedures:

- "Modifying the Read View Name" on page 411
- "Modifying the Write View Name" on page 414
- "Modifying the Notify View Name" on page 416
- "Modifying the Storage Type" on page 418

Modifying the Read View Name

To modify the Read View Name parameter in an SNMPv3 Access Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **4** to select Configure SNMPv3 Access Table.

   The Configure SNMPv3 Access Table is shown in Figure 138 on page 406.

3. From the Configure SNMPv3 Access Table, type **3** to select Modify SNMPv3 Table Entry.

The Modify SNMPv3 Access Table is shown in Figure 139.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                    11:20:02 02-Jan-2006

                      Modify SNMPv3 Access Table

Group Name .... sales                      Security Model . v3
Context Prefix.                            Security Level . AuthNoPriv
Read View...... systemmanagers             Context Match .. Exact
Write View .... salespeople                Storage Type ... Volatile
Notify View ... salespeople                Row Status ..... Active

1 - Set Read View Name
2 - Set Write View Name
3 - Set Notify View Name
4 - Set Storage Type

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 139. Modify SNMPv3 Access Table Menu

4. To modify the Read View Name parameter, type **1** to select Set Read View Name.

   The following prompt is displayed:

   ```
   Enter Group Name:
   ```

5. Enter a Group Name that was previously configured.

   The following prompt is displayed:

   ```
   Enter Security Model [1-v1, 2-v2c, 3-v3]:
   ```

6. Enter the Security Model configured for this Group Name. You cannot change the value of the Security Model parameter.

   Select one of the following SNMP protocols:

   1-v1
   Select this value to associate the Group Name with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Group Name with the SNMPv2c protocol.

   3-v3
   Select this value to associate the Group Name with the SNMPv3 protocol.

The following prompt is displayed:

```
Enter Security Level [N-NoAuthNoPriv, A-AuthNoPriv,
P-AuthPriv]:
```

7. Select one of the following security levels:

   N-NoAuthNoPriv
   This option represents no authentication and no privacy protocol. Select this security level if you do not want to authenticate SNMP users and you do not want to encrypt messages using a privacy protocol. This security level provides the least security.

   Note
   If you have selected SNMPv1 or SNMPv2c, N-NoAuthNoPriv is the only security level you can select.

   A-AuthNoPriv
   This option represents authentication, but no privacy protocol. Select this security level if you want to authenticate SNMP users, but you do not want to encrypt messages using a privacy protocol.You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   P-AuthPriv
   This option represents authentication and the privacy protocol. Select this security level to encrypt messages using a privacy protocol and authenticate SNMP users. This level provides the greatest level of security. You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   The following prompt is displayed:

   ```
   Enter Read View Name:
   ```

8. Enter a value that you configured with the View Name parameter in the SNMPv3 View Table. See "Creating an SNMPv3 View Table Entry" on page 396.

   A Read View Name allows the users assigned to this Security Group to view the information specified in the View Table. This value does not need to be unique.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Write View Name

To modify the Write View Name parameter in an SNMPv3 Access Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **4** to select Configure SNMPv3 Access Table.

   The Configure SNMPv3 Access Table is shown in Figure 138 on page 406.

3. From the Configure SNMPv3 Access Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Table Menu is shown in Figure 139 on page 412.

4. To modify the Write View Name parameter, type **2** to select Set Write View Name.

   The following prompt is displayed:

   ```
   Enter Group Name:
   ```

5. Enter a Group Name that was previously configured.

   The following prompt is displayed:

   ```
   Enter Security Model[1-v1, 2-v2c, 3-v3]:
   ```

6. Enter the Security Model configured for this Group Name. You cannot change the value of the Security Model parameter.

   Select one of the following SNMP protocols:

   1-v1
   Select this value to associate the Group Name with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Group Name with the SNMPv2c protocol.

   3-v3
   Select this value to associate the Group Name with the SNMPv3 protocol.

The following prompt is displayed:

```
Enter Security Level [N-NoAuthNoPriv, A-AuthNoPriv,
P-AuthPriv]:
```

7. Enter the Security Level configured for this Group Name. You cannot change the value of the Security Level parameter.

   Select one of the following security levels:

   N-NoAuthNoPriv
   This option represents no authentication and no privacy protocol. Select this security level if you do not want to authenticate SNMP users and you do not want to encrypt messages using a privacy protocol. This security level provides the least security.

   Note
   If you have selected SNMPv1 or SNMPv2c, N-NoAuthNoPriv is the only security level you can select.

   A-AuthNoPriv
   This option represents authentication, but no privacy protocol. Select this security level if you want to authenticate SNMP users, but you do not want to encrypt messages using a privacy protocol.You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   P-AuthPriv
   This option represents authentication and the privacy protocol. Select this security level to encrypt messages using a privacy protocol and authenticate SNMP users. This level provides the greatest level of security. You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   The following prompt is displayed:

   ```
   Enter Write View Name:
   ```

8. Enter a value that you configured with the View Name parameter in the SNMPv3 View Table.

   A Write View Name allows the people assigned to this Security Group to write, or modify, to the information in the specified View Table. This value does not need to be unique.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Notify View Name

To modify the Notify View Name parameter in an SNMPv3 Access Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **4** to select Configure SNMPv3 Access Table.

   The Configure SNMPv3 Access Table is shown in Figure 138 on page 406.

3. From the Configure SNMPv3 Access Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Table Menu is shown in Figure 139 on page 412.

4. To modify the Notify View Name parameter, type **3** to select Set Notify View Name.

   The following prompt is displayed:

   ```
   Enter Group Name:
   ```

5. Enter a Group Name that was previously configured.

   The following prompt is displayed:

   ```
   Enter Security Model[1-v1, 2-v2c, 3-v3]:
   ```

6. Enter the Security Model configured for this Group Name. You cannot change the value of the Security Model parameter.

   Select one of the following SNMP protocols:

   1-v1
   Select this value to associate the Group Name with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Group Name with the SNMPv2c protocol.

   3-v3
   Select this value to associate the Group Name with the SNMPv3 protocol.

The following prompt is displayed:

```
Enter Security Level [N-NoAuthNoPriv, A-AuthNoPriv,
P-AuthPriv]:
```

7. Enter the Security Level configured for this Group Name. You cannot change the value of the Security Level parameter.

   Select one of the following security levels:

   N-NoAuthNoPriv
   This option represents no authentication and no privacy protocol. Select this security level if you do not want to authenticate SNMP users and you do not want to encrypt messages using a privacy protocol. This security level provides the least security.

   Note
   If you have selected SNMPv1 or SNMPv2c, N-NoAuthNoPriv is the only security level you can select.

   A-AuthNoPriv
   This option represents authentication, but no privacy protocol. Select this security level if you want to authenticate SNMP users, but you do not want to encrypt messages using a privacy protocol.You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   P-AuthPriv
   This option represents authentication and the privacy protocol. Select this security level to encrypt messages using a privacy protocol and authenticate SNMP users. This level provides the greatest level of security. You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   The following prompt is displayed:

   ```
   Enter Notify View Name:
   ```

8. Enter a value that you configured with the View Name parameter in the SNMPv3 View Table.

   A Notify View Name permits the users assigned to this Security Group to send traps specified in this view of the MIB tree. This value does not need to be unique.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Storage Type

To modify the Storage Type parameter in an SNMPv3 Access Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **4** to select Configure SNMPv3 Access Table.

   The Configure SNMPv3 Access Table is shown in Figure 138 on page 406.

3. From the Configure SNMPv3 Access Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Table Menu is shown in Figure 139 on page 412.

4. To modify the Storage Type parameter, type **4** to select Set Storage Type.

   The following prompt is displayed:

   ```
   Enter Group Name:
   ```

5. Enter a Group Name that was previously configured.

   The following prompt is displayed:

   ```
   Enter Security Model[1-v1, 2-v2c, 3-v3]:
   ```

6. Enter the Security Model configured for this Group Name. You cannot change the value of the Security Model parameter.

   Select one of the following SNMP protocols:

   1-v1
   Select this value to associate the Group Name with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Group Name with the SNMPv2c protocol.

   3-v3
   Select this value to associate the Group Name with the SNMPv3 protocol.

The following prompt is displayed:

```
Enter Security Level [N-NoAuthNoPriv, A-AuthNoPriv,
P-AuthPriv]:
```

7. Enter the Security Level configured for this Group Name. You cannot change the value of the Security Level parameter.

   Select one of the following security levels:

   N-NoAuthNoPriv
   This option represents no authentication and no privacy protocol. Select this security level if you do not want to authenticate SNMP users and you do not want to encrypt messages using a privacy protocol. This security level provides the least security.

   Note
   If you have selected SNMPv1 or SNMPv2c, N-NoAuthNoPriv is the only security level you can select.

   A-AuthNoPriv
   This option represents authentication, but no privacy protocol. Select this security level if you want to authenticate SNMP users, but you do not want to encrypt messages using a privacy protocol.You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   P-AuthPriv
   This option represents authentication and the privacy protocol. Select this security level to encrypt messages using a privacy protocol and authenticate SNMP users. This level provides the greatest level of security. You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-NonVolatile]:
   ```

8. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 Access Table to the configuration file. After making changes to an SNMPv3 Access Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 Access Table to the configuration file. After making changes to an SNMPv3 Access Table entry with a NonVolatile storage type, the

S - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the SNMPv3 SecurityToGroup Table

This section contains a description of the SNMPv3 SecurityToGroup Table and how to create, delete, and modify table entries. The SNMPv3 SecurityToGroup Table allows you to associate a User Name with a Group Name. The User Name is configured in the Configure SNMPv3 User Table Menu while the Group Name is configured in the Configure SNMPv3 Access Table Menu. In addition, the configuration in the Configure SNMPv3 Access Table Menu defines which MIB views this User can read, write (modify), and send traps from. For each User Name, you can assign:

- A Security Model (SNMPv1, SNMPv2c, SNMPv3)
- A Group Name
- A Storage Type

There are three functions you can perform with the SNMPv3 Access Table.

- "Creating an SNMPv3 SecurityToGroup Table Entry" on page 421
- "Deleting an SNMPv3 SecurityToGroup Table Entry" on page 424
- "Modifying an SNMPv3 SecurityToGroup Table Entry" on page 425

Creating an SNMPv3 SecurityToGroup Table Entry

To create an entry in the SecurityToGroup Table, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **5** to select Configure SNMPv3 SecurityToGroup Table.

The Configure SNMPv3 SecurityToGroup Table Menu is shown in Figure 140.

```
         Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                  00:14:33 15-Jan-2006
               Configure SNMPv3 SecurityToGroup Table
Security Model.................. v3
Security Name ................. spike
Group Name .................... marketing
Storage Type .................. NonVolatile
Row Status .................... Active

1 - Create SNMPv3 Table Entry
2 - Delete SNMPv3 Table Entry
3 - Modify SNMPv3 Table Entry

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 140. Configure SNMPv3 SecurityToGroup Table Menu

3. To configure a group in the SNMPv3 SecurityToGroup Table, type **1** to select Create SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter User (Security) Name:
   ```

4. Enter the User Name that you want to associate with a group.

   Enter a User Name that you configured in “Creating an SNMPv3 User Table Entry” on page 386.

   The following prompt is displayed:

   ```
   Enter Security Model [1-v1, 2-v2c, 3-v3]:
   ```

5. Select the SNMP protocol that was configured for this User Name.

   Choose from the following:

   1-v1
   Select this value to associate the Group Name with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Group Name with the SNMPv2c protocol.

3-v3
Select this value to associate the Group Name with the SNMPv3 protocol.

The following prompt is displayed:

```
Enter Group Name:
```

6. Enter a Group Name that you configured in the SNMPv3 Access Table. See. "Creating an SNMPv3 Access Table Entry" on page 405.

   There are four default values for this field:

   - defaultV1GroupReadOnly
   - defaultV1GroupReadWrite
   - defaultV2cGroupReadOnly
   - defaultV2cGroupReadWrite

   These values are reserved for SNMPv1 and SNMPv2c implementations.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-NonVolatile]:
   ```

7. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 SecurityToGroup Table to the configuration file. After making changes to an SNMPv3 SecurityToGroup Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 SecurityToGroup Table to the configuration file. After making changes to an SNMPv3 SecurityToGroup Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

   Note
   The Row Status parameter is a read-only field in the Telnet and Local interfaces. The Active value indicates the SNMPv3 SecurityToGroup Table entry will take effect immediately.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an SNMPv3 SecurityToGroup Table Entry

You may want to delete an entry from the SNMPv3 SecurityToGroup Table. When you delete an SNMPv3 SecurityToGroup Table entry, there is no way to undelete, or recover, it.

To delete an entry in the SNMPv3 SecurityToGroup Table, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **5** to select Configure SNMPv3 SecurityToGroup Table.

   The SNMPv3 SecurityToGroup Table is shown in Figure 140 on page 422.

   Note
   To display a Group Name and its associated parameters from the Configure SNMPv3 SecurityToGroup Table Menu, type **N** to display the Next Page and **P** to display the previous page.

3. From the SNMPv3 SecurityToGroup Table, type **2** to select Delete SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter User (Security) Name:
   ```

4. Enter a User Name.

   The following prompt is displayed:

   ```
   Enter Security Model [1-v1, 2-v2c, 3-v3]:
   ```

5. Enter the Security Model of this User Name. Choose from the following:

   1-v1
   Select this value to associate the Group Name with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Group Name with the SNMPv2c protocol.

   3-v3
   Select this value to associate the Group Name with the SNMPv3 protocol.

The following prompt is displayed:

```
Do you want to delete this table entry? (Y/N): [Yes/No]->
```

6. Enter **Y** to delete this SecurityToGroup entry or **N** to save it.
7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an SNMPv3 SecurityToGroup Table Entry

This section describes how to modify parameters in an SNMPv3 SecurityToGroup Table entry. See the following procedures:

- "Modifying the Group Name" on page 425
- "Modifying the Storage Type" on page 427

Modifying the Group Name

To modify the Group Name in an SNMPv3 SecurityToGroup Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **5** to select Configure SNMPv3 SecurityToGroup Table.

   The Configure SNMPv3 SecurityToGroup Table is shown in Figure 138 on page 406.

3. From the Configure SNMPv3 SecurityToGroup Table, type **3** to select Modify SNMPv3 Table Entry.

The Modify SecurityToGroup Table is displayed as shown Figure 140.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                  11:20:02 02-Jan-2006

                  Modify SNMPv3 SecurityToGroup Table

Security Model................. v3
Security Name ................. cleo72
Group Name .................... engineering
Storage Type .................. Volatile
Row Status .................... Active

1 - Set Group Name
2 - Set Storage Type

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 141. Modify SNMPv3 SecurityToGroup Table Menu

4. To modify the Group Name, type **1** to select Set Group Name.

   The following prompt is displayed:

   ```
   Enter User (Security) Name:
   ```

5. Enter a User Name.

   The User Name must be previously configured in the Configure SNMPv3 User Table Menu. See “Creating an SNMPv3 User Table Entry” on page 386.

   The following prompt is displayed:

   ```
   Enter Security Model [1-v1, 2-v2c, 3-v3]:
   ```

6. Enter the Security Model configured for this User Name. You cannot change the value of the Security Model parameter.

   Select one of the following SNMP protocols:

   1-v1
   Select this value if this User Name is configured with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the User Name with the SNMPv2c protocol.

3-v3
Select this value to associate the User Name with the SNMPv3 protocol.

The following prompt is displayed:

```
Enter Group Name:
```

7. Enter the new Group Name.

   This value must match a value configured in the Group Name parameter in the Configure SNMPv3 Access Table. See "Creating an SNMPv3 Access Table Entry" on page 405.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Storage Type

To modify the Storage Type in an SNMPv3 SecurityToGroup Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **5** to select Configure SNMPv3 SecurityToGroup Table.

   The Configure SNMPv3 SecurityToGroup Table is shown in Figure 138 on page 406.

3. From the Configure SNMPv3 SecurityToGroup Table, type **3** to select Modify SNMPv3 Table Entry.

4. To modify the storage type, type **2** to select Set Storage Type.

   The following prompt is displayed:

   ```
   Enter User (Security) Name:
   ```

5. Enter a User Name.

   The User Name must be previously configured in the Configure SNMPv3 User Table Menu. See "Creating an SNMPv3 User Table Entry" on page 386.

   The following prompt is displayed:

   ```
   Enter Security Model [1-v1, 2-v2c, 3-v3]:
   ```

6. Enter the Security Model configured for this User Name. You cannot change the value of the Security Model parameter.

   Select one of the following SNMP protocols:

   1-v1
   Select this value if this User Name is configured with the SNMPv1 protocol.

   2-v2c
   Select this value if this User Name is configured with the SNMPv2c protocol.

   3-v3
   Select this value if this User Name is configured with the SNMPv3 protocol.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-NonVolatile]:
   ```

7. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 SecurityToGroup Table to the configuration file. After making changes to an SNMPv3 SecurityToGroup Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 SecurityToGroup Table to the configuration file. After making changes to an SNMPv3 SecurityToGroup Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the SNMPv3 Notify Table

This section contains a description of the SNMPv3 Notify Table Menu and how to create, delete, and modify table entries. The Configure SNMPv3 Notify Table Menu allows you to define a name for sending traps. In each Notify Table entry, you define if the switch sends a trap or an inform message. The two message types, trap and inform, have different packet formats.

For each Notify group, you can configure:

- Notify Name
- Notify Tag
- Notify Type
- Storage Type

The value of the Notify Tag is linked with the Tag List parameter in the Configure SNMPv3 Target Address Table Menu. After you configure a value for the Notify Tag parameter, you use the same value in the Target List parameter that is located on the Target Address Table Menu. As a result of this connection between the two tables, the Notify Tag parameter assigns a Target IP address to the Notify Table internally.

There are three functions you can perform with the Configure SNMPv3 Notify Table Menu.

- "Creating an SNMPv3 Notify Table Entry" on page 429
- "Deleting an SNMPv3 Notify Table Entry" on page 431
- "Modifying an SNMPv3 Notify Table Entry" on page 432

Creating an SNMPv3 Notify Table Entry

To create an entry in the SNMPv3 Notify Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **6** to select Configure SNMPv3 Notify Table.

The Configure SNMPv3 Notify Table Menu is shown in Figure 142.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  00:14:33 15-Jan-2006

                     Configure SNMPv3 Notify Table

Notify Name ..................... hardwareengineeringTrap
Notify Tag ...................... hardwareengineeringtag
Notify Type ..................... Trap
Storage Type .................... NonVolatile
Row Status ...................... Active

1 - Create SNMPv3 Table Entry
2 - Delete SNMPv3 Table Entry
3 - Modify SNMPv3 Table Entry

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 142. Configure SNMPv3 Notify Table Menu

3. To create an entry in the table, type **1** to select Create SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Notify Name:
   ```

4. Enter the name associated with this trap message.

   Enter a name of up to 32-alphanumeric characters. For example, you might want to define a trap message for hardware engineering and enter a value of “hardwareengineeringtrap” for the Notify Name.

   The following prompt is displayed:

   ```
   Enter Notify Tag:
   ```

5. Enter the name of the Notify Tag.

   Enter a name of up to 32 alphanumeric characters. The following prompt is displayed:

   ```
   Enter Notify Type [T-Trap, I-Inform]:
   ```

6. Enter one of the following message types:

   T-Trap
   Indicates this notify table is used to send traps. With this message type, the switch does not expects a response from the authoritative entity.

I-Inform
Indicates this notify table is used to send inform messages. With this message type, the switch expects a response from the authoritative entity.

The following prompt is displayed:

```
Enter Storage Type [V-Volatile, N-NonVolatile]:
```

7. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 Notify Table to the configuration file. After making changes to an SNMPv3 Notify Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 Notify Table to the configuration file. After making changes to an SNMPv3 Notify Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

   Note
   The Row Status parameter is a read-only field in the Telnet and Local interfaces. The Active value indicates the SNMPv3 Notify Table entry takes effect immediately.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an SNMPv3 Notify Table Entry

You may want to delete an entry from the Configure SNMPv3 Notify Table Menu. When you delete a Configure SNMPv3 Notify Table entry, there is no way to undelete, or recover, it.

To delete an entry in the Configure SNMPv3 Notify Table Menu, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **6** to select Configure SNMPv3 Notify Table.

The Configure SNMPv3 Notify Table Menu is shown in Figure 142 on page 430.

Note
To display a Group Name and its associated parameters from the Configure SNMPv3 SecurityToGroup Table Menu, type **N** to display the Next Page and **P** to display the previous page.

3. To delete an SNMPv3 Notify Table entry, type **2** to select Delete SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Notify Name:
   ```

4. Enter a Notify Name.

   The following prompt is displayed:

   ```
   Do you want to delete this table entry? (Y/N): [Yes/No]->
   ```

5. Enter **Y** to delete the SNMPv3 Notify Table entry or **N** to save it.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an SNMPv3 Notify Table Entry

This section describes how to modify parameters in an SNMPv3 Notify Table entry. See the following procedures:

- “Modifying a Notify Tag” on page 432
- “Modifying a Notify Type” on page 434
- “Modifying a Storage Type” on page 435

Modifying a Notify Tag

To modify the Notify Tag parameter in an SNMPv3 Notify Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **6** to select Configure SNMPv3 Notify Table.

   The Configure SNMPv3 Notify Table Menu is shown in Figure 142 on page 430.

3. From the Configure SNMPv3 Notify Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Notify Table Menu is displayed as shown in Figure 143.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                 11:20:02 02-Jan-2006
                     Modify SNMPv3 Notify Table
Notify Name ..................... softwareeengineering
Notify Tag...................... softwareeengineeringtag
Notify Type..................... Inform
Storage Type .................... NonVolatile
Row Status ...................... Active

1 - Set Notify Tag
2 - Set Notify Type
3 - Set Storage Type

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 143. Modify SNMPv3 Notify Table Menu

Note
To display a Group Name and its associated parameters from the Configure SNMPv3 SecurityToGroup Table Menu, type **N** to display the Next Page and **P** to display the previous page.

4. To modify the Notify Tag, type **1** to select Set Notify Tag.

   The following prompt is displayed:

   ```
   Enter Notify Name:
   ```

5. Enter a Notify Name.

   The following prompt is displayed:

   ```
   Enter Notify Tag:
   ```

6. Enter the new Notify Tag.

   Enter an alphanumeric value of up to 32 characters.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Notify Type

To modify the Notify Type parameter in an SNMPv3 Notify Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **6** to select Configure SNMPv3 Notify Table.

   The Configure SNMPv3 Notify Table Menu is shown in Figure 142 on page 430.

3. From the Configure SNMPv3 Notify Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Notify Table is shown in Figure 143 on page 433.

4. To modify the Notify Type, type **2** to select Set Notify Type.

   The following prompt is displayed:

   ```
   Enter Notify Name:
   ```

5. Enter a Notify Name.

   The following prompt is displayed:

   ```
   Enter Notify Type [T-Trap, I-Inform]:
   ```

6. Enter one of the following message types:

   T-Trap
   Indicates this notify table is used to send traps. With this message type, the switch does not expect a response from the host.

   I-Inform
   Indicates this notify table is used to send inform messages. With this message type, the switch expects a response from the host.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Storage Type

To modify the Storage Type parameter in an SNMPv3 Notify Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **6** to select Configure SNMPv3 Notify Table.

   The Configure SNMPv3 Notify Table Menu is shown in Figure 142 on page 430.

3. From the Configure SNMPv3 Notify Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Notify Table is shown in Figure 143 on page 433.

4. To modify the Storage Type, type **3** to select Set Storage Type.

   The following prompt is displayed:

   ```
   Enter Notify Name:
   ```

5. Enter a Notify Name.

   The following prompt is displayed:

   ```
   Enter Storage type [V-Volatile, N-NonVolatile]:
   ```

6. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 Notify Table to the configuration file. After making changes to an SNMPv3 Notify Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 Notify Table to the configuration file. After making changes to an SNMPv3 Notify Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the SNMPv3 Target Address Table

This section contains a description of the SNMPv3 Target Address Table Menu and how to create, delete, and modify table entries. You use the SNMPv3 Target Address Table Menu to assign the IP address of a host that is used for generating notifications. The Configure SNMPv3 Target Address Table Menu is linked internally to the Configure SNMPv3 Notify Table through the Tag List parameter. The Configure SNMPv3 Notify Table Menu receives the host IP address through the configuration of the SNMPv3 Target Address Table Menu.

For each Target Address Table entry, you can configure the following parameters:

- Target Address Name
- Target IP Address
- UDP Port
- Timeout Value
- Number of Retries
- Tag List
- Target Parameters
- Storage Type

You must configure the Tag List parameter with values previously configured in the Notify Tag parameter. The Notify Tag parameter is located on the Notify Table Menu. See “Creating an SNMPv3 Notify Table Entry” on page 429.

There are three functions you can perform with the Configure SNMPv3 Target Address Table Menu.

- “Creating an SNMPv3 Target Address Table Entry” on page 437
- “Deleting an SNMPv3 Target Address Table Entry” on page 439
- “Modifying an SNMPv3 Target Address Table Entry” on page 440

Creating an SNMPv3 Target Address Table Entry

To create an entry in the Configure SNMPv3 Target Address Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Address Table Menu is shown in Figure 144.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                 00:14:33 15-Jan-2006
                Configure SNMPv3 Target Address Table
Target Addr Name ... host451                  Timeout ..... 1500
Target Parameters .. SNMPmanagerPC            Retries ..... 3
IP Address ......... 198.35.11.1              UDP Port# ... 162
Storage Type ....... NonVolatile              Row Status .. Active
Tag List ........... hwengTrap hwengInform swengTrap swengInform
1 - Create SNMPv3 Table Entry
2 - Delete SNMPv3 Table Entry
3 - Modify SNMPv3 Table Entry

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 144. Configure SNMPv3 Target Address Table Menu

3. To create an entry in the SNMPv3 Target Address Table, type **1** to select Create SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Target Address Name:
   ```

4. Enter the name of the SNMP manager, or host, that manages the SNMP activity on your switch.

   You can enter a name of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter IP Address:
   ```

5. Enter the IP address of the host.

   Use the following format for an IP address:
   XXX.XXX.XXX.XXX

   The following prompt is displayed:

   ```
   Enter UDP Port#: [0 to 65535]-> 162
   ```

6. Enter a UDP port.

   You can enter a UDP port in the range of 0 to 65,535. The default UDP port is 162.

   The following prompt is displayed:

   ```
   Enter Timeout (10mS): [0 to 2147483647]-> 1500
   ```

7. Enter a timeout value in milliseconds.

   When an Inform message is generated, it requires a response from the switch. The timeout value determines how long the switch considers the Inform message an active message. This parameter applies to Inform messages only. The range is from 0 to 2,147,483,647 milliseconds. The default value is 1500 milliseconds.

   The following prompt is displayed:

   ```
   Enter Retries:[0 to 255]-> 3
   ```

8. Enter the number of times the switch will retry, or resend, an Inform message.

   When an Inform message is generated, it requires a response from the switch. This parameter determines how many times the switch resends an Inform message. The Retries parameter applies to Inform messages only. The range is 0 to 255 retries. The default is 3 retries.

   The following prompt is displayed:

   ```
   Enter Tag List:
   ```

9. Enter a Tag List.

   This list consists of a tag or list of tags you configured in a Configure SNMPv3 Notify Table entry with the Notify Tag parameter. See “Creating an SNMPv3 Notify Table Entry” on page 429. Enter a Tag List of up to 256 alphanumeric characters. Use a space to separate entries, for example:

   ```
   hwengtag swengtag testengtag
   ```

   The following prompt is displayed:

   ```
   Enter Target Parameters:
   ```

10. Enter a Target Parameters name.

    This name can consist of up to 32-alphanumeric characters. The value configured here must match the value configured with the Target Parameters Name parameter in the Configure SNMPv3 Target Parameters Table.

    The following prompt is displayed:

    ```
    Enter Storage Type [V-Volatile, N-NonVolatile]:
    ```

11. Select one of the following storage types for this table entry:

    V - Volatile
    Select this storage type if you do not want the ability to save an entry in the SNMPv3 Target Address Table to the configuration file. After making changes to an SNMPv3 Target Address Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

    N-NonVolatile
    Select this storage type if you want the ability to save an entry in the SNMPv3 Target Address Table to the configuration file. After making changes to an SNMPv3 Target Address entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

    Note
    The Row Status parameter is a read-only field in the Telnet and Local interfaces. The Active value indicates the SNMPv3 Target Address Table entry will take effect immediately.

12. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an SNMPv3 Target Address Table Entry

You may want to delete an entry from the SNMPv3 Target Address Table. After you delete an SNMPv3 Target Address Table entry, there is no way to undelete, or recover, it.

To delete an entry in the SNMPv3 Target Address Table, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

The Configure SNMPv3 Target Address Table Menu is shown in Figure 146 on page 450.

Note
To display a Group Name and its associated parameters from the Configure SNMPv3 SecurityToGroup Table Menu, type **N** to display the Next Page and **P** to display the previous page.

3. To delete an SNMPv3 Target Address Table entry, type **2** to select Delete SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Target Address Name:
   ```

4. Enter a Target Address Name.

   The following prompt is displayed:

   ```
   Do you want to delete this table entry? (Y/N): [Yes/No]->
   ```

5. Enter **Y** to delete the SNMPv3 Target Address Table entry or **N** to save it.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an SNMPv3 Target Address Table Entry

This section describes how to modify parameters in an SNMPv3 Target Address Table entry. See the following procedures:

- “Modifying a Target IP Address” on page 441
- “Modifying the Target Address UDP Port” on page 442
- “Modifying the Target Address Timeout” on page 443
- “Modifying the Target Address Retries” on page 444
- “Modifying the Target Address Tag List” on page 445
- “Modifying the Target Parameters Field” on page 446
- “Modifying the Storage Type” on page 447

Note
You cannot modify the Target Address Name parameter.

Modifying a Target IP Address

To modify the target IP address in an SNMPv3 Target Address Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Address Table Menu is shown in Figure 144 on page 437.

3. From the Configure SNMPv3 Target Address Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Address Table Menu is shown in Figure 145.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                 11:20:02 02-Jan-2006

                  Modify SNMPv3 Target Address Table

Target Addr Name ... host451                  Timeout ..... 1500
Target Parameters .. SNMPmanagerPC            Retries ..... 3
IP Address ......... 198.35.11.1              UDP Port# ... 162
Storage Type ....... NonVolatile              Row Status .. Active
Tag List ........... hwengTrap hwengInform swengTrap swengInform

1 - Set Target IP Address
2 - Set Target Address UDP Port
3 - Set Target Address Timeout
4 - Set Target Address Retries
5 - Set Target Address TagList
6 - Set Target Parameters
7 - Set Storage Type

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 145. Modify SNMPv3 Target Address Table Menu

4. To change the Target IP Address, type **1** to select Set Target IP Address.

The following prompt is displayed:

```
Enter Target Address Name:
```

5. Enter a previously configured Target Address Name.

   This is the name of the SNMP manager, or host, that manages the SNMP activity on your switch. You can enter a name of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter IP Address:
   ```

6. Enter the IP address of the host.

   Use the following format for an IP address:
   XXX.XXX.XXX.XXX

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Target Address UDP Port

To modify the Target Address UDP Port parameter in an SNMPv3 Target Address Table entry, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Address Table Menu is shown in Figure 144 on page 437.

3. From the Configure SNMPv3 Target Address Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Address Table Menu is shown in Figure 145 on page 441.

4. To change the Target Address UDP Port, type **2** to select Set Target Address UDP Port.

   The following prompt is displayed:

   ```
   Enter Target Address Name:
   ```

5. Enter a previously configured Target Address Name.

   This is the name of the SNMP manager, or host, that manages the SNMP activity on your switch. You can enter a name of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter UDP Port#: [0 to 65535]-> 162
   ```

6. Enter a UDP port.

   You can enter a UDP port in the range of 0 to 65,535. The default UDP port is 162.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Target Address Timeout

The Target Address Timeout parameter only applies when the message type is an Inform message. To modify the Target Address Timeout parameter in an SNMPv3 Target Address Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Address Table Menu is shown in Figure 144 on page 437.

3. From the Configure SNMPv3 Target Address Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Address Table Menu is shown in Figure 145 on page 441.

4. To modify the Target Address Timeout, type **3** to select Set Target Address Timeout.

   The following prompt is displayed:

   ```
   Enter Target Address Name:
   ```

5. Enter a previously configured Target Address Name.

This is the name of the SNMP manager, or host, that manages the SNMP activity on your switch. You can enter a name of up to 32-alphanumeric characters.

The following prompt is displayed:

```
Enter Timeout (10mS): [0 to 2147483647]-> 1500
```

6. Enter a timeout value in milliseconds.

   When an Inform message is generated, it requires a response from the switch. The timeout value determines how long the switch considers the Inform message an active message. This parameter applies to Inform messages only. The range is from 0 to 2,147,483,647 milliseconds. The default value is 1500 milliseconds.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Target Address Retries

The Target Address Retries parameter only applies when the message type is an Inform message. To modify the Target Address Retries parameter in an SNMPv3 Target Address Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Address Table Menu is shown in Figure 144 on page 437.

3. From the Configure SNMPv3 Target Address Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Address Table Menu is shown in Figure 145 on page 441.

4. To modify the Target Address Retries, type **4** to select Set Target Address Retries.

   The following prompt is displayed:

   ```
   Enter Target Address Name:
   ```

5. Enter a previously configured Target Address Name.

This is the name of the SNMP manager, or host, that manages the SNMP activity on your switch. You can enter a name of up to 32-alphanumeric characters.

The following prompt is displayed:

```
Enter Retries:[0 to 255]-> 3
```

6. Enter the number of times the switch will retry, or resend, the Inform message.

   The range is 0 to 255 retries. The default is 3 retries.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Target Address Tag List

To modify the Target Address Tag List parameter in an SNMPv3 Target Address Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Address Table Menu is shown in Figure 144 on page 437.

3. From the Configure SNMPv3 Target Address Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Address Table Menu is shown in Figure 145 on page 441.

4. To modify the Target Address Tag List, type **5** to select Set Target Address TagList.

   The following prompt is displayed:

   ```
   Enter Target Address Name:
   ```

5. Enter a previously configured Target Address Name.

   This is the name of the SNMP manager, or host, that manages the SNMP activity on your switch. You can enter a name of up to 32-alphanumeric characters.

The following prompt is displayed:

```
Enter Tag List:
```

Enter a Tag List of up to 256 alphanumeric characters. Use a space to separate entries. This list consists of a tag or list of tags you configured in a Configure SNMPv3 Notify Table entry with the Notify Tag parameter. See “Creating an SNMPv3 Notify Table Entry” on page 429.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Target Parameters Field

To modify the Target Parameters field in an SNMPv3 Target Address Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Address Table Menu is shown in Figure 144 on page 437.

3. From the Configure SNMPv3 Target Address Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Address Table Menu is shown in Figure 145 on page 441.

4. To modify the Target Parameters field, type **6** to select Set Target Parameters.

   The following prompt is displayed:

   ```
   Enter Target Address Name:
   ```

5. Enter a previously configured Target Address Name.

   This is the name of the SNMP manager, or host, that manages the SNMP activity on your switch. You can enter a name of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter Target Parameters:
   ```

6. Enter a Target Parameters Name.

   The value configured here must match the value configured with the Target Parameters Name parameter in the Configure SNMPv3 Target Parameters Table. This name can consist of up to 32-alphanumeric characters.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Storage Type

To modify the Storage Type parameter in an SNMPv3 Target Address Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **7** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Address Table Menu is shown in Figure 144 on page 437.

3. From the Configure SNMPv3 Target Address Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Address Table Menu is shown in Figure 145 on page 441.

4. To modify the Storage Type, type **7** to select Set Storage Type.

   The following prompt is displayed:

   ```
   Enter Target Address Name:
   ```

5. Enter a previously configured Target Address Name.

   This is the name of the SNMP manager, or host, that manages the SNMP activity on your switch. You can enter a name of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-NonVolatile]:
   ```

6. Select one of the following storage types for this table entry:

V - Volatile
Select this storage type if you do not want the ability to save an entry in the SNMPv3 Target Address Table to the configuration file. After making changes to an SNMPv3 Target Address Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

N-NonVolatile
Select this storage type if you want the ability to save an entry in the SNMPv3 Target Address Table to the configuration file. After making changes to an SNMPv3 Target Address entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the SNMPv3 Target Parameters Table

This section contains a description of the SNMPv3 Target Parameters Table and how to create, delete, and modify table entries. The SNMPv3 Target Parameters Table links the user security information with the message notification information configured in the Configure SNMPv3 Notify Table Menu and Configure SNMPv3 Target Address Table Menu.

In the SNMPv3 Target Parameters Table, you specify the SNMP parameters that are used when a message is generated to a target, or host, IP address. The SNMPv3 Target Parameters Table also links a User Name and its related security information, called *user security information*, with a host. The user security information consists of the following parameters listed in the SNMPv3 tables where they are configured:

- User Name parameter configured in the SNMPv3 User Table Menu
- View Name parameter configured in the SNMPv3 View Table Menu
- Group Name, Security Model, and Security Level parameters configured in the SNMPv3 Access Table
- User Name, Security Model, and Group Name configured in the SNMPv3 SecurityToGroup Table

When you enter user security information in an SNMPv3 Target Parameters Table entry, it must match the configuration in the SNMPv3 tables listed above. If the user security information in the SNMPv3 Target Parameters Table entry does not match the configuration in the tables listed above, messages are not sent on behalf of the user.

Note
In the SNMPv3 Target Parameters Table, the Security Name parameter is the equivalent to the User Name parameter in the SNMPv3 User Table.

For each Target Address Table entry, you can configure:

- Target Parameters Name
- Security Name (User Name)
- Security Model
- Security Level
- Storage Type

There are three functions you can perform with the Configure SNMPv3 Target Parameters Table Menu.

- "Creating an SNMPv3 Target Parameters Table Entry" on page 450

Creating an SNMPv3 Target Parameters Table Entry

To create an entry in the Configure SNMPv3 Target Parameters Table, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **8** to select Configure SNMPv3 Target Parameters Table Menu.

   The Configure SNMPv3 Target Parameters Table Menu is shown in Figure 146.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                  00:14:33 15-Jan-2006

              Configure SNMPv3 Target Parameters Table

Target Parameters Name ... host125parm
Message Processing Model . v3
Security Model............ v3
Security Name ............ murthy
Security Level ........... AuthPriv
Storage Type ............. NonVolatile
Row Status ............... Active

1 - Create SNMPv3 Table Entry
2 - Delete SNMPv3 Table Entry
3 - Modify SNMPv3 Table Entry

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 146. Configure SNMPv3 Target Parameters Table Menu

3. To create an SNMPv3 Target Parameters Table, type **1** to select Create SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Target Parameters Name:
   ```

4. Enter a name of the Target Parameters.

Enter a value of up to 32-alphanumeric characters.

Note
You are prompted to enter a value for the Message Processing Model parameter only if you select SNMPv1 or SNMPv2c as the Security Model. If you select the SNMPv3 protocol as the Security Model, then the Message Processing Model is automatically assigned to SNMPv3.

The following prompt is displayed:

```
Enter User (Security) Name:
```

5. Enter a User Name.

   The value of this parameter is previously configured with the Configure SNMPv3 User Table. See "Creating an SNMPv3 User Table Entry" on page 386.

   The following prompt is displayed:

   ```
   Enter Security Model [1-v1, 2-v2c, 3-v3]:
   ```

6. Select one of the following SNMP protocols as the Security Model for this Security Name, or User Name.

   1-v1
   Select this value to associate the Security Name, or User Name, with the SNMPv1 protocol.

   2-v2c
   Select this value to associate the Security Name, or User Name, with the SNMPv2c protocol.

   3-v3
   Select this value to associate the Security Name, or User Name, with the SNMPv3 protocol. The SNMPv3 protocol allows you to configure the group to authenticate SNMPv3 users and to encrypt messages.

   The following prompt is displayed:

   ```
   Enter Security Level [N-NoAuthNoPriv, A-AuthNoPriv,
   P-AuthPriv]:
   ```

7. Select one of the following Security Levels:

   Note
   The value you configure for the Security Level must match the value configured for the User Name in the Configure SNMPv3 User Table Menu. See "Creating an SNMPv3 User Table Entry" on page 386.

N-NoAuthNoPriv
This option represents no authentication and no privacy protocol. Select this security level if you do not want to authenticate SNMP users and you do not want to encrypt messages using a privacy protocol. This security level provides the least security.

Note
If you have selected SNMPv1 or SNMPv2c, N-NoAuthNoPriv is the only security level you can select.

A-AuthNoPriv
This option represents authentication, but no privacy protocol. Select this security level if you want to authenticate SNMP users, but you do not want to encrypt messages using a privacy protocol.You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

P-AuthPriv
This option represents authentication and the privacy protocol. Select this security level to encrypt messages using a privacy protocol and authenticate SNMP users. This level provides the greatest level of security. You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

The following prompt is displayed:

```
Enter Storage Type [V-Volatile, N-NonVolatile]:
```

8. Select one of the following storage types for this table entry:

V - Volatile
Select this storage type if you do not want the ability to save an entry in the SNMPv3 Target Parameters Table to the configuration file. After making changes to an SNMPv3 Target Parameters Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

N-NonVolatile
Select this storage type if you want the ability to save an entry in the SNMPv3 Target Parameters Table to the configuration file. After making changes to an SNMPv3 Target Parameters Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

Note
The Row Status parameter is a read-only field in the Telnet and Local interfaces. The Active value indicates the SNMPv3 Target Parameters Table entry will take effect immediately.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an SNMPv3 Target Parameters Table Entry

You may want to delete an entry from the SNMPv3 Target Parameters Table. When you delete an SNMPv3 Target Parameters Table entry, there is no way to undelete, or recover, it.

To delete an entry in the SNMPv3 Target Parameters Table, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **8** to select Configure SNMPv3 Target Parameters Table.

   The Configure SNMPv3 Parameters Table Menu is shown in Figure 146 on page 450.

   Note
   To display a Group Name and its associated parameters from the Configure SNMPv3 SecurityToGroup Table Menu, type **N** to display the Next Page and **P** to display the previous page.

3. To delete an SNMPv3 Target Parameters Table entry, type **2** to select Delete SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Target Parameters Name:
   ```

4. Enter a Target Parameters Name.

   The following prompt is displayed:

   ```
   Do you want to delete this table entry? (Y/N): [Yes/No]->
   ```

5. Enter **Y** to delete the SNMPv3 Target Address Table entry or **N** to save it.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an SNMPv3 Target Parameters Table Entry

This section provides procedures for modifying parameters in an SNMPv3 Target Parameters Table entry. The parameter values configured in the Target Parameters Table must match those configured in the other tables. For a more detailed explanation, see “Creating an SNMPv3 Target Parameters Table Entry” on page 450.

In an SNMPv3 Target Parameters Table entry, the Security Name parameter is linked to the User Name parameter on the SNMPv3 User Table. In an SNMPv3 User Table entry, the User Name parameter is used as an index for the entry. Because the User Name and Security Name parameters are linked, the information you configure that relates to a User Table entry must match the information you configure in the SNMPv3 Target Parameters Table entry. In addition, the values configured for the following parameters in an SNMPv3 Target Parameters Table entry must match those configured in the corresponding table entry:

- User Name parameter in the SNMPv3 User Table
- View Name parameter in the SNMPv3 View Table
- Group Name, Security Model, and Security Level parameters in the SNMPv3 Access Table
- User Name, Security Model, Group Name parameters in the SNMPv3 SecurityToGroup Table

See the following procedures:

- “Modifying the Security Name (User Name)” on page 454
- “Modifying the Security Model” on page 456
- “Modifying the Security Level” on page 457
- “Modifying the Message Process Model” on page 459
- “Modifying the Storage Type” on page 460

Note
You cannot modify the Target Params Name parameter.

Note
You cannot modify an entry in the SNMPv3 Target Parameter Table that contains a value of “default” in the Target Parameters Name field.

Modifying the Security Name (User Name)

In the AT-S62 implementation of the SNMPv3 protocol, the Security Name and the User Name parameters are equivalent. In the SNMPv3 Target Parameters Table Menu, the Security Name and the User Name parameters are used interchangeably.

When you modify the Security Name parameter, you must use a value that you configured with the User Name parameter in the Configure SNMPv3 User Table Menu. If you do not use a value configured with the User Name parameter, messages are not sent on behalf of this User Name. See "Creating an SNMPv3 User Table Entry" on page 386.

To modify the Security Name parameter in an SNMPv3 Target Parameter Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **8** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Parameters Table Menu is shown in Figure 146 on page 450.

3. From the Configure SNMPv3 Target Parameters Table Menu, type **3** to select Modify SNMPv3 Table Entry. The Modify SNMPv3 Target Parameters Table Menu is shown in Figure 147.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006

                Modify SNMPv3 Target Parameters Table

Target Parameters Name ... host27
Message Processing Model . v3
Security Model..... v3
Security Name ..... hoa
Security Level ..... AuthNoPriv
Storage Type ..... NonVolatile
Row Status ..... Active

1 - Set Security Name
2 - Set Security Model
3 - Set Security Level
4 - Set Message Processing Model
5 - Set Storage Type

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 147. Modify SNMPv3 Target Parameters Table Menu

4. To change the Security Name parameter, type **1** to select Set Security Name.

   The following prompt is displayed:

   ```
   Enter Target Parameters Name:
   ```

5. Enter a previously configured Target Parameters Name.

   Enter a value of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter User (Security) Name:
   ```

6. Enter a User Name.

   Enter a value that you previously configured with the Configure SNMPv3 User Table Menu. You can enter a value of up to 32-alphanumeric characters.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Security Model

For the Security or User Name you have selected, the value of the Security Model parameter in an SNMPv3 Target Parameter Table entry must match the value of the Security Model parameter in the SNMPv3 Access Table entry.

Caution
If the values of the Security Model parameter in the SNMPv3 User Table and the SNMPv3 Target Parameter Table entry do not match, notification messages are not generated on behalf of this User (Security) Name.

To modify the Security Model parameter in an SNMPv3 Target Parameter Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **8** to select Configure SNMPv3 Target Address Table.

The Configure SNMPv3 Target Parameters Table Menu is shown in Figure 146.

3. From the Configure SNMPv3 Target Parameters Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Parameters Table Menu is shown in Figure 147 on page 455.

4. To change the Security Model, type **2** to select Security Model.

   The following prompt is displayed:

   ```
   Enter Target Parameters Name:
   ```

5. Enter a previously configured Target Parameters Name.

   Enter a value of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter Security Model [1-v1, 2-v2c, 3-v3]:
   ```

6. Select one of the following SNMP protocols that was previously configured as the Security Model for this Security Name, or User Name.

   1-v1
   Select this value if this User Name is associated with the SNMPv1 protocol.

   2-v2c
   Select this value if this User Name is associated with the SNMPv2c protocol.

   3-v3
   Select this value if this User Name is associated with the SNMPv3 protocol.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Security Level

For the Security or User Name you have selected, the value of the Security Level parameter in an SNMPv3 Target Parameter Table entry must match the value of the Security Level parameter in the SNMPv3 User Table entry.

To modify the Security Level parameter in an SNMPv3 Target Parameter Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type

5->**5**->**5**.

The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **8** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Parameters Table Menu is shown in Figure 146.

3. From the Configure SNMPv3 Target Parameters Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Parameters Table Menu is shown in Figure 147 on page 455.

4. To modify the Security Level, type **3** to select Set Security Level.

   The following prompt is displayed:

   ```
   Enter Target Parameters Name:
   ```

5. Enter a previously configured Target Parameters Name.

   Enter a value of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter Security Level [N-NoAuthNoPriv, A-AuthNoPriv,
   P-AuthPriv]:
   ```

6. Enter the Security Level.

   Select one of the following Security Levels:

   Note
   The value you configure for the Security Level must match the value configured for the User Name in the Configure SNMPv3 User Table Menu. See “Creating an SNMPv3 User Table Entry” on page 386.

   N-NoAuthNoPriv
   This option represents no authentication and no privacy protocol. Select this security level if you do not want to authenticate SNMP users and you do not want to encrypt messages using a privacy protocol. This security level provides the least security.

   Note
   If you have selected SNMPv1 or SNMPv2c, N-NoAuthNoPriv is the only security level you can select.

A-AuthNoPriv
This option represents authentication, but no privacy protocol. Select this security level if you want to authenticate SNMP users, but you do not want to encrypt messages using a privacy protocol.You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

P-AuthPriv
This option represents authentication and the privacy protocol. Select this security level to encrypt messages using a privacy protocol and authenticate SNMP users. This level provides the greatest level of security. You can select this value if you configured the Security Model parameter with the SNMPv3 protocol.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Message Process Model

You can modify the Message Process Model for SNMPv1 and SNMPv2c protocol configurations only. When you configure the SNMPv3 protocol, the Message Process Model is automatically assigned to the SNMPv3 protocol.

To modify the Message Process Model parameter in an SNMPv3 Target Parameter Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **8** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Parameters Table Menu is shown in Figure 146.

3. From the Configure SNMPv3 Target Parameters Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Parameters Table Menu is shown in Figure 147 on page 455.

4. To modify the Message Process Model, type **4** to select Set Message Processing Model.

   The following prompt is displayed:

   ```
   Enter Target Parameters Name:
   ```

5. Enter a previously configured Target Parameters Name.

   Enter a value of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter Message Processing Model[1-v1,2-v2c,3-v3]:
   ```

6. Select one of the following SNMP protocols that is used to process, or send messages:

   1-v1
   Select this value to process messages with the SNMPv1 protocol.

   2-v2c
   Select this value to process messages with the Security Name, or User Name, with the SNMPv2c protocol.

   3-v3
   Select this value to process messages with the SNMPv3 protocol.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Storage Type

To modify the Storage Type parameter in an SNMPv3 Target Parameter Table entry, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **8** to select Configure SNMPv3 Target Address Table.

   The Configure SNMPv3 Target Parameters Table Menu is shown in Figure 146.

3. From the Configure SNMPv3 Target Parameters Table Menu, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Target Parameters Table Menu is shown in Figure 147 on page 455.

4. To modify the Storage Type, type **5** to select Storage Type.

   The following prompt is displayed:

   ```
   Enter Target Parameters Name:
   ```

5. Enter a previously configured Target Parameters Name.

   Enter a value of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter Storage Type [V-Volatile, N-NonVolatile]:
   ```

6. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 Target Parameters Table to the configuration file. After making changes to an SNMPv3 Target Parameters Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 Target Parameters Table to the configuration file. After making changes to an SNMPv3 Target Parameters Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the SNMPv3 Community Table

This section contains a description of the SNMPv3 Community Table and how to create, delete, and modify table entries. The SNMPv3 Community Table allows you to create SNMPv1 and SNMPv2c Communities using the SNMPv3 Tables.

Allied Telesyn does not recommend that you use the menu described in this section to configure SNMPv1 and SNMPv2c communities. Instead, use the procedures described in Chapter 5: “SNMPv1 and SNMPv2c Configuration” on page 89.

However, if you want to configure SNMPv1 and SNMPv2c with the SNMPv3 Tables you need to start your configuration with the SNMPv3 Community Table and then create entries in the following tables:

- SNMPv3 View Table—See “Creating an SNMPv3 View Table Entry” on page 396.
- SNMPv3 Access Table—See “Creating an SNMPv3 Access Table Entry” on page 405.
- SNMPv3 SecurityToGroup Table—See “Creating an SNMPv3 SecurityToGroup Table Entry” on page 421.
- SNMPv3 Notify Table—See “Configuring the SNMPv3 Notify Table” on page 429.
- SNMPv3 Target Address Table—See “Creating an SNMPv3 Target Address Table Entry” on page 437.
- SNMPv3 Target Parameters Table—See “Creating an SNMPv3 Target Parameters Table Entry” on page 450.

It is important to note that you do not create an entry in the SNMPv3 User Table when you are configuring SNMPv1 and SNMPv2c with the SNMPv3 Tables. When you configure the SNMPv3 protocol, the various tables are linked with the User Name parameter and its related information. With the SNMPv1 and SNMPv2c configuration, the Security Name parameter and its related information (configured in the SNMPv3 Community Table Menu) links an SNMPv3 Community Table entry to the other SNMPv3 Table entries.

Note
In the SNMPv3 Community Table entry, the Security Name parameter is not related to the User Name parameter.

For each SNMPv3 Community Table entry, you can configure the following parameters:

- Community Index
- Community Name
- Security Name
- Transport Tag
- Storage Type

In addition, you can display the entries configured with the Configure SNMPv1 & SNMPv2c Community Menu in the Configure SNMPv3 Community Table Menu. However, you cannot modify an SNMPv1 & SNMPv2c Community Table entry with the Configure SNMPv3 Community Table Menu.

There are three functions you can perform with the Configure SNMPv3 Target Parameters Table Menu.

- "Creating an SNMPv3 Community Table Entry" on page 463
- "Deleting an SNMPv3 Community Table Entry" on page 466
- "Modifying an SNMPv3 Community Table Entry" on page 467

Creating an SNMPv3 Community Table Entry

To create an entry in the Configure SNMPv3 Community Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **9** to select Configure SNMPv3 Community Table.

The Configure SNMPv3 Community Table Menu is shown in Figure 148.

```
         Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                 00:14:33 15-Jan-2006
                  Configure SNMPv3 Community Table
Community Index ............... ATIIndex1
Community Name ................ 451engineering75
Security Name ................. debashi48
Transport Tag ................. sampletag
Storage Type .................. NonVolatile
Row Status .................... Active

1 - Create SNMPv3 Table Entry
2 - Delete SNMPv3 Table Entry
3 - Modify SNMPv3 Table Entry

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 148. Configure SNMPv3 Community Table Menu

3. To create an entry in the SNMPv3 Community Table, type **1** to select Create SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Community Index:
   ```

4. Enter the name of this Community Index.

   This parameter describes the name of this community. It is used to index the other parameters in an SNMPv3 Community Table entry. Enter a value of up to 32-alphanumeric characters.

   The following prompt is displayed:

   ```
   Enter Community Name:
   ```

5. Enter a Community Name of up to 64 alphanumeric characters. The value of the Community Name parameter acts as a password for the SNMPv3 Community Table entry. This parameter is case sensitive.

Note
Allied Telesyn recommends that you select SNMP Community Names carefully to ensure these names are known only to authorized personnel.

The following prompt is displayed:

```
Enter Security Name:
```

6. Enter the name of an SNMPv1 and SNMPv2c user.

   This name must be unique. Enter a value of up to 32-alphanumeric characters.

   Note
   Do not use a value configured with the User Name parameter in the SNMPv3 User Table.

   The following prompt is displayed:

   ```
   Enter Transport Tag:
   ```

7. Enter a name of up to 32-alphanumeric characters for the Transport Tag.

   The Transport Tag parameter is similar to the Notify Tag parameter in the SNMPv3 Notify Table. Add the value you configure for the Transport Tag parameter to the Tag List parameter in the Target Address Table. In this way, the Transport Tag parameter links an SNMPv3 Community Table entry with an entry in the SNMPv3 Target Address Table. See "SNMPv3 Target Address Table" on page 383.

   The following prompt is displayed:

   ```
   Enter Storage type [V-volatile, N-NonVolatile]:
   ```

8. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to save an entry in the SNMPv3 Community Table to the configuration file. After making changes to an SNMPv3 Community Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 Community Table to the configuration file. After making changes to an SNMPv3 Community Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

   Note
   The Row Status parameter is a read-only field in the Telnet and Local interfaces. The Active value indicates the SNMPv3 Community Table entry takes effect immediately.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an SNMPv3 Community Table Entry

You may want to delete an entry from the SNMPv3 Community Table. When you delete an entry in the SNMPv3 Community Table, there is no way to undelete or recover it.

To delete an entry in the SNMPv3 Community Table, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **9** to select Configure SNMPv3 Community Table.

   The Configure SNMPv3 Community Table Menu is shown in Figure 148 on page 464.

3. To delete an entry in the SNMPv3 Community Table, type **2** to select Delete SNMPv3 Table Entry.

   The following prompt is displayed:

   ```
   Enter Community Index:
   ```

4. Enter the Community Index that you want to delete.

   The following prompt is displayed:

   ```
   Do you want to delete this table entry? (Y/N): [Yes/No]->
   ```

5. Choose one of the following:

   Y
   Type Y to delete an SNMPv3 Community table entry.

   N
   Type N to retain the SNMPv3 Community table entry.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an SNMPv3 Community Table Entry

For each entry in the SNMPv3 Community Table, you can modify the following parameters:

- Community Name
- Security Name
- Transport Tag
- Storage Type

However, you cannot modify the Community Index parameter.

Although you can display the SNMPv1 and SNMPv2c configuration created with the procedures described in Chapter 5: "SNMPv1 and SNMPv2c Configuration" on page 89, you cannot modify these Community Table entries with the SNMPv3 Tables.

See the following procedures:

- "Modifying the Community Name" on page 467
- "Modifying the Security Name" on page 469
- "Modifying the Transport Tag" on page 469
- "Modifying the Storage Type" on page 470

Modifying the Community Name

To modify the Community Name parameter in an SNMPv3 Community Table entry, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **9** to select Configure SNMPv3 Community Table.

   The Configure SNMPv3 Community Table Menu is shown in Figure 148 on page 464.

3. From the Configure SNMPv3 Community Table, type **3** to select Modify SNMPv3 Table Entry.

The Modify SNMPv3 Community Table Menu is shown in Figure 149.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                 00:14:33 15-Jan-2006

                    Modify SNMPv3 Community Table

Community Index ............... alliedtelesynindex
Community Name ................ 789bothel23wa
Security Name ................. buster
Transport Tag ................. 72
Storage Type .................. Volatile
Row Status .................... Active

1 - Set Community Name
2 - Set Security Name
3 - Set Transport Tag
4 - Set Storage Type

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 149. Modify SNMPv3 Community Table Menu

4. To change the Community Name, type **1** to select Set Community Name.

   The following prompt is displayed:

   ```
   Enter Community Index:
   ```

5. Enter the Community Index that you want to modify.

   The following prompt is displayed:

   ```
   Enter Community Name:
   ```

6. Enter the new Community Name.

   The value of the Community Name parameter acts as a password for the SNMPv3 Community Table entry. This parameter is case sensitive. Enter a value of up to 64 alphanumeric characters.

   Note
   Allied Telesyn recommends that you select SNMP Community Names carefully to ensure these names are known only to authorized personnel.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Security Name

To modify the Security Name parameter in an SNMPv3 Community Table entry, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **9** to select Configure SNMPv3 Community Table.

   The Configure SNMPv3 Community Table Menu is shown in Figure 148 on page 464.

3. From the Configure SNMPv3 Community Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Community Table Menu is shown in Figure 149 on page 468.

4. To change the Security Name, type **2** to select Set Security Name.

   The following prompt is displayed:

   ```
   Enter Community Index:
   ```

5. Enter the Community Index of the Security Name you want to change.

   The following prompt is displayed:

   ```
   Enter Security Name:
   ```

6. Enter the new Security Name.

   Enter a value of up to 32-alphanumeric characters.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Transport Tag

To modify the Transport Tag parameter in an SNMPv3 Community Table entry, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in "Creating an SNMPv3 User Table Entry" on page 386. Or, from the Main Menu type **5**->**5**->**5**.

The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **9** to select Configure SNMPv3 Community Table.

   The Configure SNMPv3 Community Table Menu is shown in Figure 148 on page 464.

3. From the Configure SNMPv3 Community Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Community Table Menu is shown in Figure 149 on page 468.

4. To change the Transport Tag, type **3** to select Set Transport Tag.

   The following prompt is displayed:

   ```
   Enter Community Index:
   ```

5. Enter the Community Index of the Transport Tag you want to change.

   The following prompt is displayed:

   ```
   Enter Transport Tag:
   ```

6. Enter the new value for the Transport Tag.

   Enter a name of up to 32-alphanumeric characters.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying the Storage Type

To modify the Storage Type parameter in an SNMPv3 Community Table entry, perform the following procedure:

1. Follow steps 1 through 3 in the procedure described in “Creating an SNMPv3 User Table Entry” on page 386. Or, from the Main Menu type **5**->**5**->**5**.

   The Configure SNMPv3 Table Menu is shown in Figure 133 on page 387.

2. From the Configure SNMPv3 Table Menu, type **9** to select Configure SNMPv3 Community Table.

   The Configure SNMPv3 Community Table Menu is shown in Figure 148 on page 464.

3. From the Configure SNMPv3 Community Table, type **3** to select Modify SNMPv3 Table Entry.

   The Modify SNMPv3 Community Table Menu is shown in Figure 149 on page 468.

4. To change the Storage Type, type **4** to select Set Storage Type.

   The following prompt is displayed:

   ```
   Enter Community Index:
   ```

5. Enter the Community Index of the Storage Type you want to change.

   The following prompt is displayed:

   ```
   Enter Storage type [V-volatile, N-NonVolatile]:
   ```

6. Select one of the following storage types for this table entry:

   V - Volatile
   Select this storage type if you do not want the ability to an entry in the SNMPv3 Community Table to the configuration file. After making changes to an SNMP Community Table entry with a Volatile storage type, the **S** - Save Configuration Changes option does not appear on the Main Menu.

   N-NonVolatile
   Select this storage type if you want the ability to save an entry in the SNMPv3 Community Table to the configuration file. After making changes to an SNMPv3 Community Table entry with a NonVolatile storage type, the **S** - Save Configuration Changes option appears on the Main Menu, allowing you to save your changes.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying SNMPv3 Table Menus

The procedures in this section describe how to display the SNMPv3 Tables. The following procedures are provided:

Displaying the Display SNMPv3 User Table Menu

This section describes how to display the Display SNMPv3 User Table Menu. For information about the SNMPv3 User Table, see “Creating an SNMPv3 User Table Entry” on page 386.

To display the Display SNMPv3 User Table Menu, perform the following procedure.

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **5** to select SNMP Configuration.

   The SNMP Configuration menu is shown in Figure 17 on page 93.

3. From the SNMP Configuration menu, type **5** to select Configure SNMPv3 Table.

   The Configure SNMP Menu is shown in Figure 133 on page 387.

4. From the Configure SNMP Menu, type **6** to select Display SNMPv3 Table.

The Display SNMPv3 Table Menu is shown in Figure 150.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  00:14:33 15-Jan-2006

                        Display SNMPv3 Table

1 - Display SNMPv3 User Table
2 - Display SNMPv3 View Table
3 - Display SNMPv3 Access Table
4 - Display SNMPv3 SecurityToGroup Table
5 - Display SNMPv3 Notify Table
6 - Display SNMPv3 Target Address Table
7 - Display SNMPv3 Target Parameters Table
8 - Display SNMPv3 Community Table

R - Return to Previous Menu

Enter your selection?
```

Figure 150. Display SNMPv3 Table Menu

5. From the Display SNMPv3 Table Menu, type **1** to select Display SNMPv3 User Table.

   The Display SNMPv3 User Table is shown in Figure 151.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  00:14:33 15-Jan-2006

                     Display SNMPv3 User Table

Engine Id ................. 80:00:00:CF:31:00:30:84:FD:57:DA
User Name ................. spike
Authentication Protocol ... MD5
Privacy Protocol .......... DES
Storage Type .............. NonVolatile
Row Status ................ Active

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 151. Display SNMPv3 User Table Menu

Displaying the Display SNMPv3 View Table Menu

This section describes how to display the Display SNMPv3 View Table Menu. For information about the SNMPv3 View Table parameters, see “Creating an SNMPv3 View Table Entry” on page 396.

To display the Display SNMPv3 View Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Displaying the Display SNMPv3 User Table Menu” on page 472. Or, from the Main Menu type **5**->**5**->**6**.

2. From the Display SNMPv3 Table Menu, type **2** to select Display SNMPv3 View Table.

   The Display SNMPv3 View Table Menu is shown in Figure 152.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                 00:14:33 15-Jan-2006

                     Display SNMPv3 View Table

View Name ..................... tcp
Subtree OID ................... 1.3.6.1
Subtree Mask ..................
View Type ..................... Included
Storage Type .................. NonVolatile
Row Status .................... Active

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 152. Display SNMPv3 View Table Menu

Displaying the Display SNMPv3 Access Table Menu

This section describes how to display the Display SNMPv3 Access Table Menu. For information about the SNMPv3 Access Table parameters, see "Creating an SNMPv3 Access Table Entry" on page 405.

To display the Display SNMPv3 Access Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Displaying the Display SNMPv3 User Table Menu" on page 472. Or, from the Main Menu type **5**->**5**->**6**.

2. From the Display SNMPv3 Table Menu, type **3** to select Display SNMPv3 Access Table.

   The Display SNMPv3 Access Table Menu is shown in Figure 153.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                           Production Switch
User: Manager                                      11:20:02 02-Jan-2006

                       Display SNMPv3 Access Table

Group Name .... technicalsales             Security Model . v3
Context Prefix.                            Security Level . AuthPriv
Read View...... internet                   Context Match .. Exact
Write View ....                            Storage Type ... NonVolatile
Notify View ...                            Row Status ..... Active

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 153. Display SNMPv3 Access Table Menu

Displaying the Display SNMPv3 SecurityToGroup Table Menu

This section describes how to display the Display SNMPv3 SecurityToGroup Table Menu. For more information about the parameters in the SNMPv3 SecurityToGroup Table Menu, see “Creating an SNMPv3 SecurityToGroup Table Entry” on page 421.

To display the Display SNMPv3 SecurityToGroup Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Displaying the Display SNMPv3 User Table Menu” on page 472. Or, from the Main Menu type **5**->**5**->**6**.

2. From the Display SNMPv3 Table Menu, type **4** to select Display SNMPv3 SecurityToGroup Table.

   The Display SNMPv3 SecurityToGroup Table Menu is shown in Figure 154.

```
         Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                  11:20:02 02-Jan-2006
               Display SNMPv3 SecurityToGroup Table
Security Model................. v3
Security Name ................. praveen
Group Name .................... hardwareengineering
Storage Type .................. NonVolatile
Row Status .................... Active

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 154. Display SNMPv3 SecurityToGroup Table Menu

Displaying the Display SNMPv3 Notify Table Menu

This section describes how to display the Display SNMPv3 Notify Table Menu. For information about the SNMPv3 Notify Table parameters, see "Creating an SNMPv3 Notify Table Entry" on page 429.

To display the Display SNMPv3 Notify Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Displaying the Display SNMPv3 User Table Menu" on page 472. Or, from the Main Menu type **5**->**5**->**6**.

2. From the Display SNMPv3 Table Menu, type **5** to select Display SNMPv3 Notify Table.

   The Display SNMPv3 Notify Table Menu is shown in Figure 154.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006

                     Display SNMPv3 Notify Table

Notify Name ..................... testengineeringTrap
Notify Tag ...................... testengineeringtag
Notify Type ..................... Inform
Storage Type .................... NonVolatile
Row Status ...................... Active


N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 155. Display SNMPv3 Notify Table Menu

Displaying the Display SNMPv3 Target Address Table Menu

This section describes how to display the Display SNMPv3 Target Address Table Menu. For information about the SNMPv3 Target Address Table parameters, see “Creating an SNMPv3 Target Address Table Entry” on page 437.

To display the Display SNMPv3 Target Address Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Displaying the Display SNMPv3 User Table Menu” on page 472. Or, from the Main Menu type **5**->**5**->**6**.

2. From the Display SNMPv3 Table Menu, type **6** to select Display SNMPv3 Target Address Table.

   The Display SNMPv3 Target Address Table Menu is shown in Figure 154.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006
                  Display SNMPv3 Target Address Table
Target Addr Name ... host99                    Timeout ..... 1500
Target Parameters .. SNMPmanagerPC             Retries ..... 5
IP Address ......... 198.35.11.1               UDP Port# ... 162
Storage Type ....... NonVolatile               Row Status .. Active
Tag List ........... engTrap engInform

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 156. Display SNMPv3 Target Address Table Menu

Displaying the Display SNMPv3 Target Parameters Table Menu

This section describes how to display the Display SNMPv3 Target Parameters Table Menu. For information about the SNMPv3 Target Parameters Table parameters, see "Creating an SNMPv3 Target Parameters Table Entry" on page 450.

To display the Display SNMPv3 Target Parameters Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in "Displaying the Display SNMPv3 User Table Menu" on page 472. Or, from the Main Menu type **5**->**5**->**6**.

2. From the Display SNMPv3 Table Menu, type **7** to select Display SNMPv3 Target Parameters Table.

   The Display SNMPv3 Target Parameters Table Menu is shown in Figure 154.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                         Production Switch
User: Manager                                  11:20:02 02-Jan-2006
               Display SNMPv3 Target Parameters Table
Target Parameters Name ... TargetIndex21
Message Processing Model . v3
Security Model ........... v3
Security Name ............ wilson
Security Level ........... AuthPriv
Storage Type ............. NonVolatile
Row Status ............... Active

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 157. Display SNMPv3 Target Parameters Table Menu

Displaying the Display SNMPv3 Community Table Menu

This section describes how to display the Display SNMPv3 Community Table Menu. For information about the SNMPv3 Community Table parameters, see “Creating an SNMPv3 Community Table Entry” on page 463.

To display the Display SNMPv3 Community Table Menu, perform the following procedure.

1. Follow steps 1 through 3 in the procedure described in “Displaying the Display SNMPv3 User Table Menu” on page 472. Or, from the Main Menu type **5**->**5**->**6**.

2. From the Display SNMPv3 Table Menu, type **8** to select Display SNMPv3 Community Table.

   The Display SNMPv3 Community Table Menu is shown in Figure 154.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006

                    Display SNMPv3 Community Table

Community Index ........ atiindex14
Community Name ......... sunnyvale
Security Name .......... hoa
Transport Tag........... sampletag14
Storage Type ........... NonVolatile
Row Status ............. Active

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 158. Display SNMPv3 Community Table Menu

Section IV

Spanning Tree Protocols

The chapters in this section explain the spanning tree protocols. The chapters include:

Chapter 22

Spanning Tree and Rapid Spanning Tree Protocols

This chapter provides background information on the Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP). The chapter also contains procedures on how to adjust the STP and RSTP bridge and port parameters. The sections in this chapter include:

Note
For detailed information on the Spanning Tree Protocol, refer to IEEE Std 802.1D. For detailed information on the Rapid Spanning Tree Protocol, refer to IEEE Std 802.1w.

The switch also supports the Multiple Spanning Tree Protocol. For background information on this spanning tree protocol and how to configure the parameter settings, refer to Chapter 23, “Multiple Spanning Tree Protocol” on page 507.

STP and RSTP Overview

The performance of a Ethernet network can be severely impaired by the existence of a physical loop in the network topology. A loop exists when two or more nodes on a network can transmit data to each other over more than one traffic path. The problem that loops pose is that Ethernet 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 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.

STP and RSTP can also activate a redundant path if the main path goes down. They maintain network connectivity by activating a backup redundant path in the event a main link fails or is taken off-line.

The principal different between the two protocols 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 physical 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 network traffic.

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.

The AT-S62 management software features both spanning tree protocols. Only one spanning tree protocol can be active on a switch at a time. The default active spanning tree is RSTP.

The STP implementation on the AT-S62 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 running spanning tree perform 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.

By changing the bridge priority number in the AT-S62 software, you can designate which switch on your network you want as the root bridge by giving it the lowest bridge priority number. You might 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. There are sixteen increments. You specify the increment that represents the desired bridge priority value. The increments are shown in Table 11.

Table 11. 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 must determine if the network contains redundant paths. If one is found, they must 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 will be the primary, active path, and which path(s) will be 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 redundant paths are placed into blocking state.

Path cost is determined through an evaluation of *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 an AT-8500 Series switch is adjustable through the management software. For STP, the range is 0 to 65,535. For RSTP, the range is 0 to 20,000,000.

Port cost also has an Auto-Detect feature. This feature allows spanning tree to automatically set the port cost according to the speed of the port, assigning a lower value for higher speeds. Auto-Detect is the default setting. Table 12 lists the STP port costs with Auto-Detect.

Table 12. STP Auto-Detect Port Costs

Port Speed	Port Cost
10 Mbps	100
100 Mbps	10
1000 Mbps	4

Table 13 lists the STP port costs with Auto-Detect when a port is part of a port trunk.

Table 13. STP Auto-Detect Port Trunk Costs

Port Speed	Port Cost
10 Mbps	4
100 Mbps	4
1000 Mbps	2

Table 14 lists the RSTP port costs with Auto-Detect.

Table 14. RSTP Auto-Detect Port Costs

Port Speed	Port Cost
10 Mbps	2,000,000
100 Mbps	200,000
1000 Mbps	20,000

Table 15 lists the RSTP port costs with Auto-Detect when the port is part of a port trunk.

Table 15. RSTP Auto-Detect Port Trunk Costs

Port Speed	Port Cost
10 Mbps	20,000
100 Mbps	20,000
1000 Mbps	2,000

You can override Auto-Detect and set the port cost manually.

Port Priority

If two paths have the same cost, the bridges must select a preferred path. In some instances this can involve the use of the *port priority* parameter. This parameter is used as a tie-breaker when two paths have the same cost. The lower the value, the higher the priority given to the port.

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 16 on page 488 lists the values and increments. The default value is 128, which is increment 8.

Table 16. Port Priority Value Increments

Increment	Port Priority	Increment	Port Priority
0	0	8	128
1	16	9	144
2	32	10	160
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 might take time for the root bridge to notify all bridges that a topology change has occurred, especially if it is a large network. If a topology change is made before all bridges have been notified, a temporary data loop could occur, 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-S62 management software. The appropriate value for this parameter depends on a number of variables, the size of your network being 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 network traffic.

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 function as 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-S62 software. The interval is measured in seconds and the default is two seconds. Consequently, if an AT-8500 Series switch is selected as the root bridge of a spanning tree domain, it transmits a BPDU every two seconds.

Point-to-Point Ports and Edge Ports

Note
This section applies only to RSTP and MSTP.

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 159 illustrates two AT-8524M switches that have been connected with one data link. With the link operating in full-duplex, the ports are point-to-point ports.

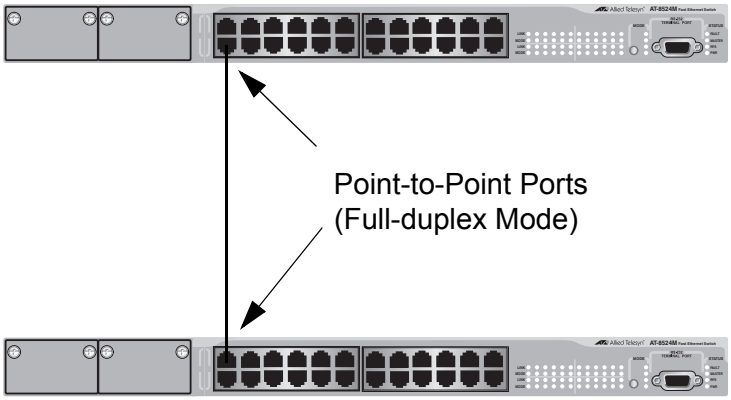

Figure 159. Point-to-Point Ports

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 160 illustrates an edge port on an AT-8524M 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.

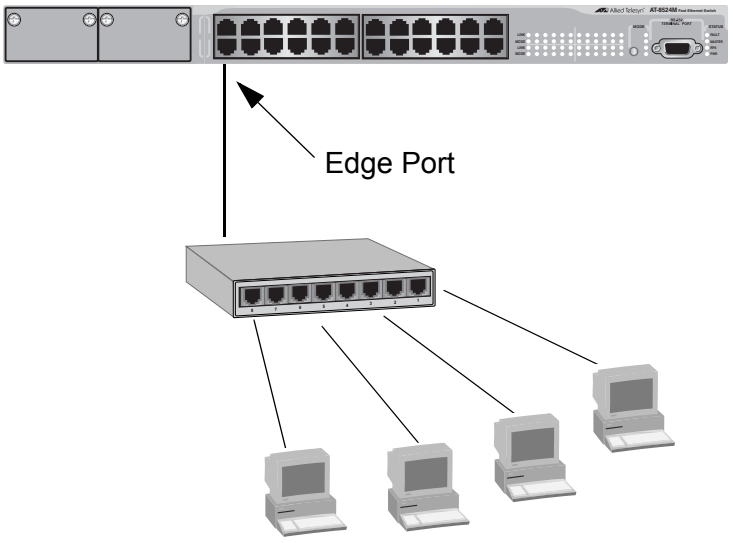

Figure 160. Edge Port

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 161 illustrates a port functioning as both a point-to-point and edge port.

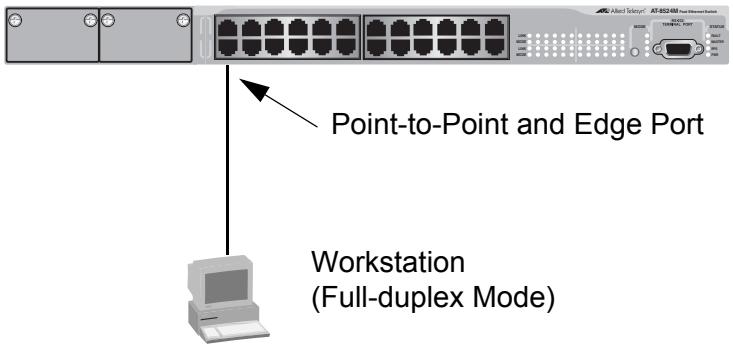

Figure 161. Point-to-Point and Edge Port

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 concept. In most cases, the default values work well.

Mixed STP and RSTP Network

RSTP IEEE 802.1w is 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.

There is no reason not to activate RSTP on an AT-8500 Series switch even when all other switches are running STP. The switch can combine its RSTP with the STP of the other switches. The 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 STP and RSTP implementations in the AT-S62 software are single-instance spanning trees. The protocols support just one spanning tree.

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 or RSTP might block a data link because it detects a data loop. This can cause fragmentation of your VLANs.

This issue is illustrated in Figure 162. Two VLANs, Sales and Production, span two AT-8524M 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.

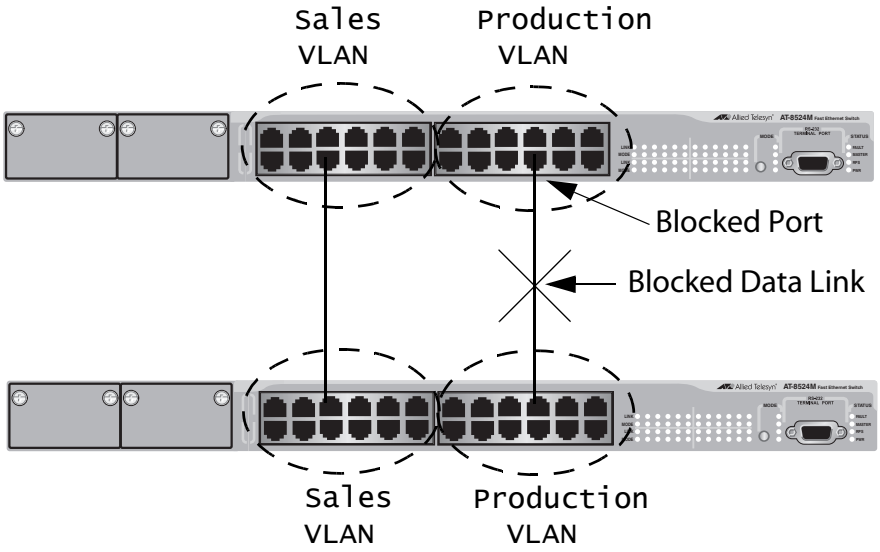

Figure 162. VLAN Fragmentation

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 24, “Port-based and Tagged Virtual LANs” on page 545.) Another approach is to use the Multiple Spanning Tree Protocol, explained in Chapter 23 on page 507, to create multiple spanning tree domains within a network.

Enabling or Disabling a Spanning Tree Protocol

The AT-S62 software supports STP, RSTP, and MSTP. (MSTP is explained in Chapter 23 on page 507.) Only one spanning tree protocol can be active on the switch at a time. Before you can configure or enable a spanning tree protocol, you must first select it as the active spanning tree protocol on the switch. After you have selected it as the active protocol, you can then configure it and enable or disable it.

To select and activate a spanning tree protocol, or to disable spanning tree, perform the following procedure:

1. From the Main Menu, type **3** to select Spanning Tree Configuration.

   The Spanning Tree Configuration menu is shown in Figure 163.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
              Spanning Tree Configuration
1 - Spanning Tree Status ...... Disabled
2 - Active Protocol Version ... RSTP
3 - Configure Active Protocol

R - Return to Previous Menu

Enter your selection?
```

Figure 163. Spanning Tree Configuration Menu

> **Note**
> Do not enable spanning tree on the switch until after you have selected an activate spanning tree protocol and configured the settings. To disable spanning tree, go to Step 5.

2. To change the active version of spanning tree on the switch, type **2** to select Active Protocol Version.

   The following prompt is displayed:

```
Enter new value (S-STP, R-RSTP, M-MSTP):
```

3. Type **S** to select STP, **R** to select RSTP, or **M** to select MSTP.

> **Note**
> A change to the active spanning tree is automatically saved on the switch.

4. If you selected STP as the active spanning tree protocol, go to "Configuring STP" on page 495 for further instructions. If you selected RSTP, go to "Configuring RSTP" on page 501. If you selected MSTP, go to Chapter 23, "Multiple Spanning Tree Protocol" on page 507.

   Note
   Once you have configured the spanning tree parameters, perform Steps 5 through 7 to enable spanning tree.

5. To enable or disable spanning tree, type **1** to select Spanning Tree Status.

   The following prompt is displayed:

   ```
   Enter new value (E-Enable, D-Disable):
   ```

6. Type **E** to enable spanning tree or **D** to disable it. The default is disabled.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring STP

This section contains the following procedures:

- "Configuring STP Bridge Settings", next
- "Configuring STP Port Settings" on page 497
- "Displaying STP Port Settings" on page 499

Configuring STP Bridge Settings

This section contains the procedure for configuring a bridge's STP settings.

> **Caution**
> The default STP parameters are adequate for most networks. Changing them without prior experience and an understanding of how STP works might have a negative effect on your network. You should consult the IEEE 802.1d standard before changing any of the STP parameters.

To configure the bridge settings, do the following:

1. From the Spanning Tree Configuration menu, type **3** to select Configure Active Protocol.

   The STP Menu is shown in Figure 164.

```
       Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                              11:20:02 02-Jan-2006
                                STP Menu
1 - Bridge Priority ........ 32768
2 - Bridge Hello Time ...... 2/2 (Configured/Actual)
3 - Bridge Forwarding ...... 15/15 (Configured/Actual)
4 - Bridge Max Age ......... 20/20 (Configured/Actual)
5 - Bridge Identifier ...... 32768/00:30:84:00:00:02
6 - Root Bridge ............ 00:30:84:00:00:02
7 - Root Priority .......... 32768
8 - Root Path Cost ......... 0

P - STP Port Parameters

R - Return to Previous Menu
Enter your selection?
```

Figure 164. STP Menu

The bridge hello time, bridge forwarding, and bridge max age parameters will have two values if STP is enabled on the switch (for example, Bridge Forwarding .. 15/15). The first number is the configured value on the switch for the parameter and the second is the value the switch obtained from the root bridge and is actually using for the parameter. The switch displays only the configured values for these parameters if spanning tree is not activated on the switch.

2. Adjust the bridge STP settings as needed. The parameters are described below.

1 - Bridge Priority
The priority number for the bridge. 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, the bridge with the numerically lowest MAC address becomes the root bridge. When a root bridge goes offline, the bridge with the next priority number automatically takes over as the root bridge. This parameter can be from 0 (zero) to 61,440 in increments of 4096, with 0 being the highest priority. For a list of the increments, refer to Table 11, Bridge Priority Value Increments on page 485.

2 - Bridge Hello Time
The time interval between generating and sending configuration messages by the bridge. This parameter can be from 1 to 10 seconds. The default is 2 seconds.

3 - Bridge Forwarding
The waiting period in seconds before a bridge changes to a new state, for example, becomes the new root bridge after the topology changes. If the bridge transitions too soon, not all links may have yet adapted to the change, resulting in network loops. The range is 4 to 30 seconds. The default is 15 seconds.

4 - Bridge Max Age
The length of time after which stored bridge protocol data units (BPDUs) are deleted by the bridge. All bridges in a bridged LAN use this aging time to test the age of stored configuration messages called bridge protocol data units (BPDUs). For example, if you use the default value 20, all bridges delete current configuration messages after 20 seconds. This parameter can be from 6 to 40 seconds.

When you select a value for maximum age, observe the following rules:

MaxAge must be greater than (2 x (HelloTime + 1))

MaxAge must be less than (2 x (ForwardingDelay - 1))

Note
The aging time for BPDUs is different from the aging time used by the MAC address table.

5 - Bridge Identifier
The MAC address of the switch. This value cannot be changed.

6 - Root Bridge
The MAC address of the root bridge of the spanning tree domain. This value cannot be changed and is only displayed when spanning tree is activated on the switch.

7 - Root Priority
The priority value on the root bridge of the spanning tree domain. This parameter is only displayed when spanning tree is enabled on the switch. To change the priority value on the root bridge, you must start a management session on the switch that is functioning as the root bridge and change its bridge priority value.

8 - Root Path Cost
The cost of the path from the current switch to the root switch of the spanning tree domain. If the current switch is the root switch, root path cost will be "0". This value cannot be changed and is only displayed when STP is activated on the switch.

3. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

4. To change STP port settings, go to the next procedure.

Configuring STP Port Settings

To adjust STP port parameters, perform the following procedure:

1. From the Spanning Tree Configuration menu, type **3** to select STP Configuration.

   The STP Menu is shown in Figure 164 on page 495.

2. From the STP Menu, type **P** to select STP Port Parameters.

The STP Port Parameters menu is shown in Figure 165.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
User: Manager                       11:20:02 02-Jan-2006
                   STP Port Parameters
1 - Configure STP Port Settings
2 - Display STP Port Configuration

R - Return to Previous Menu

Enter your selection?
```

Figure 165. STP Port Parameters Menu

3. Type **1** to select Configure STP Port Settings.

   The following prompt is displayed:

   ```
   Start Port to Configure [1 to 26] ->
   ```

4. Enter the number of the port you want to configure. To configure a range of ports, enter the first port of the range.

   The following prompt is displayed:

   ```
   End Port to Configure [1 to 24] ->
   ```

5. To configure just one port, enter the same port number here as you entered in the previous step. To configure a range of ports, enter the last port of the range.

   The Configure STP Port Settings menu is shown in Figure 166.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
              Configure STP Port Settings
Configuring Ports 4-4
1 - Port Priority ..... 128
2 - Port Cost ......... Automatic-Update

R - Return to Previous Menu

Enter your selection?
```

Figure 166. Configure STP Port Settings Menu

6. Adjust the settings as needed. The parameters are described below.

1 - Port Priority
This parameter is used as a tie breaker when two or more ports are determined to have equal costs to the root bridge. The range is 0 to 240 in increments of 16. The default value is 8 (priority value 128). For a list of the increments, refer to Table 16, "Port Priority Value Increments" on page 488.

2 - Port Cost
The spanning tree algorithm uses the cost parameter to decide which port provides the lowest cost path to the root bridge for that LAN. The range is 0 to 65,535. The default setting is Automatic Update, which sets port cost depending on the speed of the port. For the default values used by Automatic Update, refer Table 12 on page 486 and Table 13 on page 486.

All changes are immediately activated on the switch.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying STP Port Settings

To display STP port settings, perform the following procedure:

1. From the Spanning Tree Configuration menu, type **3** to select STP Configuration.

   The STP Menu is shown in Figure 164 on page 495.

2. From the STP Menu, type **P** to select STP Port Parameters.

   The STP Port Parameters menu is shown in Figure 165 on page 498.

3. From the STP Port Parameters menu, type **2** to select Display STP Port Configuration.

The Display STP Port Configuration menu is shown in Figure 167.

```
    Allied Telesyn AT-8400 Series AT-8524M - AT-S60
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
            Display STP Port Configuration
Port   State       Cost            Priority
-----------------------------------------------
1      Enabled     Auto-Update     128
2      Enabled     Auto-Update     128
3      Enabled     Auto-Update     128
4      Enabled     Auto-Update     128
5      Enabled     Auto-Update     128
6      Enabled     Auto-Update     128
7      Enabled     Auto-Update     128
8      Enabled     Auto-Update     128

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 167. Display STP Port Configuration Menu

The information in the menu is as follows:

Port - The port number.

State - Current state of the port. The possible states are Enabled or Disabled.

Cost - Port cost of the port. The default is Auto-Update.

Priority - The number used as a tie-breaker when two or more ports have equal costs to the root bridge.

Configuring RSTP

This section contains the following procedures:

- "Configuring RSTP Bridge Settings", next
- "Configuring RSTP Port Settings" on page 503
- "Displaying Port RSTP Status" on page 505

Configuring RSTP Bridge Settings

This section contains the procedure for configuring a bridge's RSTP settings.

Caution
The default RSTP parameters are adequate for most networks. Changing them without prior experience and an understanding of how RSTP works might have a negative effect on your network. You should consult the IEEE 802.1w standard before changing any of the RSTP parameters.

To configure the RSTP bridge settings, do the following

1. From the Spanning Tree Configuration menu, type **3** to select Configure Active Protocol.

   The RSTP Menu is shown in Figure 168.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                    11:20:02 02-Jan-2006
                                RSTP Menu
1 - Force Version .......... RSTP
2 - Bridge Priority ........ 32768 (In multiples of 4096: 8)
3 - Bridge Hello Time ...... 2/2 (Configured/Actual)
4 - Bridge Forwarding ...... 15/15 (Configured/Actual)
5 - Bridge Max Age ......... 20/20 (Configured/Actual)
6 - Bridge Identifier ...... 32768/00:30:84:00:00:02
7 - Root Bridge ............ 00:30:84:00:00:02
8 - Root Priority .......... 32768
9 - Root Path Cost ......... 0

P - RSTP Port Parameters

R - Return to Previous Menu

Enter your selection?
```

Figure 168. RSTP Menu

The bridge hello time, bridge forwarding, and bridge max age parameters will have two values if RSTP is enabled on the switch (for example, Bridge Forwarding..15/15). The first number is the configured value on the switch for the parameter and the second is the value the switch obtained from the root bridge and is currently using for the parameter. The switch displays only the configured values for these parameters if spanning tree is not activated on the switch.

2. Adjust the parameters as needed. The parameters are defined below.

1 - Force Version
This selection determines whether the bridge will operate with RSTP or in an STP-compatible mode. If you select RSTP, the bridge will operate all ports in RSTP, except for those ports that receive STP BPDU packets. If you select Force STP Compatible, the bridge will operate in RSTP, using the RSTP parameter settings, but it will send only STP BPDU packets out the ports.

2 - Bridge Priority
The priority number for the bridge. This number is used in determining 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, the bridge with the numerically lowest MAC address becomes the root bridge. When a root bridge goes off-line, the bridge with the next priority number automatically takes over as the root bridge. This parameter can be from 0 (zero) to 61,440 in increments of 4096, with 0 being the highest priority. For a list of the increments, refer to Table 11, “Bridge Priority Value Increments” on page 485.

3 - Bridge Hello Time
The time interval between generating and sending configuration messages by the bridge. This parameter can be from 1 to 10 seconds. The default is 2 seconds.

4 - Bridge Forwarding
The waiting period before a bridge changes to a new state, for example, becomes the new root bridge after the topology changes. If the bridge transitions too soon, not all links may have yet adapted to the change, possibly resulting in a network loop. The range is 4 to 30 seconds. The default is 15 seconds. This setting applies only to ports running in the STP-compatible mode.

5 - Bridge Max Age
The length of time after which stored bridge protocol data units (BPDUs) are deleted by the bridge. All bridges in a bridged LAN use this aging time to test the age of stored configuration messages called bridge protocol data units (BPDUs). For example, if you use the default 20, all bridges delete current configuration messages after 20 seconds. This parameter can be from 6 to 40 seconds. The default is 20 seconds.

When you select a value for maximum age, observe the following rules:

MaxAge must be greater than (2 x (HelloTime + 1))

MaxAge must be less than (2 x (ForwardingDelay - 1))

6 - Bridge Identifier
The bridge identifier of the switch. The identifier consists of the switch's bridge priority value and MAC address. The values are separated by a slash (/). To change the switch's priority value, use option 2, Bridge Priority. The MAC address of the switch cannot be changed.

7 - Root Bridge
The MAC address of the root bridge of the spanning tree domain. This value cannot be changed and is only displayed when RSTP is activated on the switch.

8 - Root Priority
The priority value on the root bridge of the spanning tree domain. This parameter is only displayed when RSTP is enabled on the switch. To change the priority value on the root bridge, you must start a management session on the switch functioning as the root bridge and change its bridge priority value.

9 - Root Path Cost
The cost of the path from the current switch to the root switch of the spanning tree domain. If the current switch is the root switch, root path cost will be "0". This value cannot be changed and is only displayed when RSTP is activated on the switch.

3. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring RSTP Port Settings

To adjust RSTP port parameters, perform the following procedure:

1. From the Spanning Tree Configuration menu, type **3** to select STP Configuration.

   The STP Menu is shown in Figure 164 on page 495.

2. From the STP Menu, type **P** to select RSTP Port Parameters.

The RSTP Port Parameters menu is shown in Figure 169.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
                 RSTP Port Parameters
1 - Configure RSTP Port Settings
2 - Display RSTP Port Configuration
3 - Display RSTP Port State

R - Return to Previous Menu

Enter your selection?
```

Figure 169. RSTP Port Parameters Menu

3. Type **1** to select Configure RSTP Port Settings.

   The following prompt is displayed:

   ```
   Starting Port to Configure [1 to 24] ->
   ```

4. Enter the number of the port you want to configure. To configure a range of ports, enter the first port of the range.

   The following prompt is displayed:

   ```
   Ending Port to Configure [1 to 24] ->
   ```

5. To configure just one port, enter the same port number here as you entered in the previous step. To configure a range of ports, enter the last port of the range.

   The Configure RSTP Port Settings menu is shown in Figure 170.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                    11:20:02 02-Jan-2006
             Configure RSTP Port Settings
Configuring Ports 4-4
1 - Port Priority ...... 128
2 - Port Cost .......... Automatic Update
3 - Point-to-Point ..... Auto Detect
4 - Edge Port .......... Yes

R - Return to Previous Menu

Enter your selection?
```

Figure 170. Configure RSTP Port Settings Menu

6. Adjust the settings as needed. The parameters are explained below.

   1 - Port Priority
   This parameter functions as a tie breaker when two or more ports have equal costs to the root bridge. The range is 0 to 240 in increments of 16. The default value is 8 (priority value 128). For a list of the increments, refer to Table 16, "Port Priority Value Increments" on page 488.

   2 - Port Cost
   The spanning tree algorithm uses the cost parameter in deciding which port provides the lowest cost path to the root bridge for that LAN. The range is 0 to 20,000,000. The default setting is Automatic Update, which sets port cost depending on the speed of the port. For the default values used by Automatic Update, refer Table 14 on page 487 and Table 15 on page 487.

   3 - Point-to-Point
   This parameter defines whether the port is functioning as a point-to-point port. For an explanation of this parameter, refer to "Point-to-Point Ports and Edge Ports" on page 489.

   4 - Edge Port
   This parameter defines whether the port is functioning as an edge port. For an explanation of this parameter, refer to "Point-to-Point Ports and Edge Ports" on page 489.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying Port RSTP Status

The RSTP Port Parameters menu has two selections for displaying a variety of RSTP port information. The two menu selections are discussed below.

2 - Display RSTP Port Configuration
This selection displays a menu that contains the current port settings for the following RSTP parameters:

Port - The port number.

Edge-Port - Whether or not the port is operating as an edge port. The possible settings are Yes and No.

Point-to-Point - Whether or not the port is functioning as a point-to-point port.

Cost - Port cost of the port. The default is Auto-Update.

Priority - The number used as a tie-breaker when two or more ports have equal costs to the root bridge.

3 - Display RSTP Port State
This selection displays a menu that contains the following RSTP operating status for a port:

Port - The port number.

State - Identifies the RSTP state of the port. Possible states are: discarding, learning, and forwarding. A state of disabled means the port has not established a link with its end node.

Role - Indicates the RSTP role of the port. Possible roles are: root, alternate, backup, and designated.

P2P - Whether or not the port is functioning as a point-to-point port.

Version - Indicates whether the port is operating in RSTP mode or STP-compatible mode.

Port Cost - Indicates the port cost of the port.

Chapter 23

Multiple Spanning Tree Protocol

This chapter provides background information on the Multiple Spanning Tree Protocol (MSTP). The chapter also explains how to adjust multiple spanning tree bridge and port parameters. The sections in this chapter include:

Note
For detailed information on the Multiple Spanning Tree Protocol, refer to IEEE Std 802.1s.

Note
You cannot configure MSTP parameters until you have selected the protocol as the active spanning tree protocol on the switch. For instructions, refer to “Selecting MSTP as the Active Spanning Tree Protocol” on page 522.

MSTP Overview

As explained in the previous chapter, STP and RSTP are single-instance spanning trees that search for physical loops across all VLANs in a bridged network. When loops are detected, the protocols stop the loops by placing one or more bridge ports in a blocking state.

As explained in “Spanning Tree and VLANs” on page 491, activating STP or RSTP can result in VLAN fragmentation when VLANs that span multiple bridges are interconnected with untagged ports. The untagged ports creating the links can represent a physical loop in the network, which will be blocked by spanning tree. The result can be a loss of communications between different parts of the same VLAN.

One way to resolve this, other than by not activating spanning tree on your network, is to link the switches using tagged ports, which can handle traffic from multiple VLANs simultaneously. The drawback is that the link formed by the tagged ports can create a bottleneck to your Ethernet traffic, resulting in reduced network performance.

Another approach is to use the Multiple Spanning Tree Protocol (MSTP). This spanning tree shares many of the same characteristics as RSTP. It features rapid convergence and has many of the same parameters. But the main difference is that while RSTP, just like STP, supports only a single-instance spanning tree, MSTP supports multiple spanning trees within a network.

The following sections describe the terms and concepts of MSTP. If you are not familiar with spanning tree or RSTP, you should first review the section “STP and RSTP Overview” on page 484.

Note
Do not activate MSTP on an AT-8500 Series switch without first familiarizing yourself with the following concepts and guidelines. Unlike STP and RSTP, you cannot activate this spanning tree protocol on a switch without first configuring the protocol parameters.

Note
The AT-S62 implementation of MSTP complies with the IEEE 802.1s standard and is compatible with versions from other vendors that conform to the standard.

Multiple Spanning Tree Instance (MSTI)

The individual spanning trees in MSTP are referred to as Multiple Spanning Tree Instances (MSTIs). A MSTI can span any number of switches. An AT-8500 Series switch can support up to 16 MSTIs at a time.

To create a MSTI, you first assign it a number, referred to as the MSTI ID. The range is 1 to 15. (The switch comes with a default MSTI with an MSTI ID of 0. This default spanning tree instance is discussed later in "Common and Internal Spanning Tree (CIST)" on page 516.)

Once you have selected an MSTI ID, you need to define its scope by assigning one or more VLANs to it. An instance can contain any number of VLANs, but a VLAN can belong to only one MSTI at a time.

Here are a couple of examples. Figure 171 illustrates two AT-8524M switches, each containing the two VLANs Sales and Production. The two parts of each VLAN are connected with a direct link using untagged ports on both switches.

If the switches were running STP or RSTP, one of the links would be blocked because the links constitute a physical loop. Which link would be blocked would depend on the STP or RSTP bridge settings. In the example, the link between the two parts of the Production VLAN is blocked, resulting in a loss of communications between the two parts of the Production VLAN.

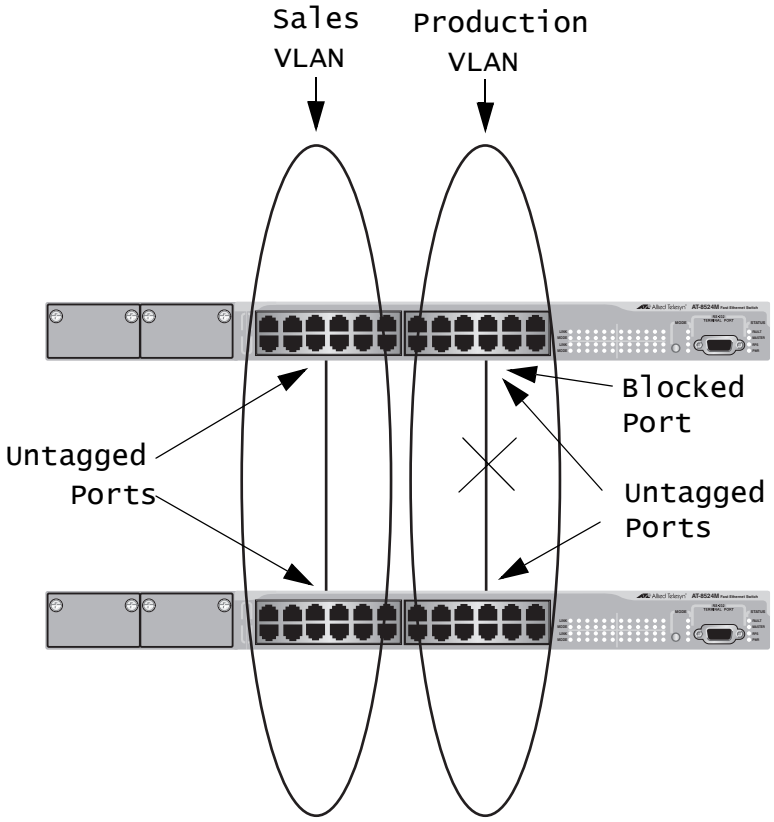

Figure 171. VLAN Fragmentation with STP or RSTP

Figure 172 illustrates the same two AT-8524M switches and the same two virtual LANs. But in this example, the two switches are running MSTP and the two VLANs have been assigned to different spanning tree instances. Both links remain active now that they reside in different MSTIs, enabling the VLANs to forward traffic over their respective direct link.

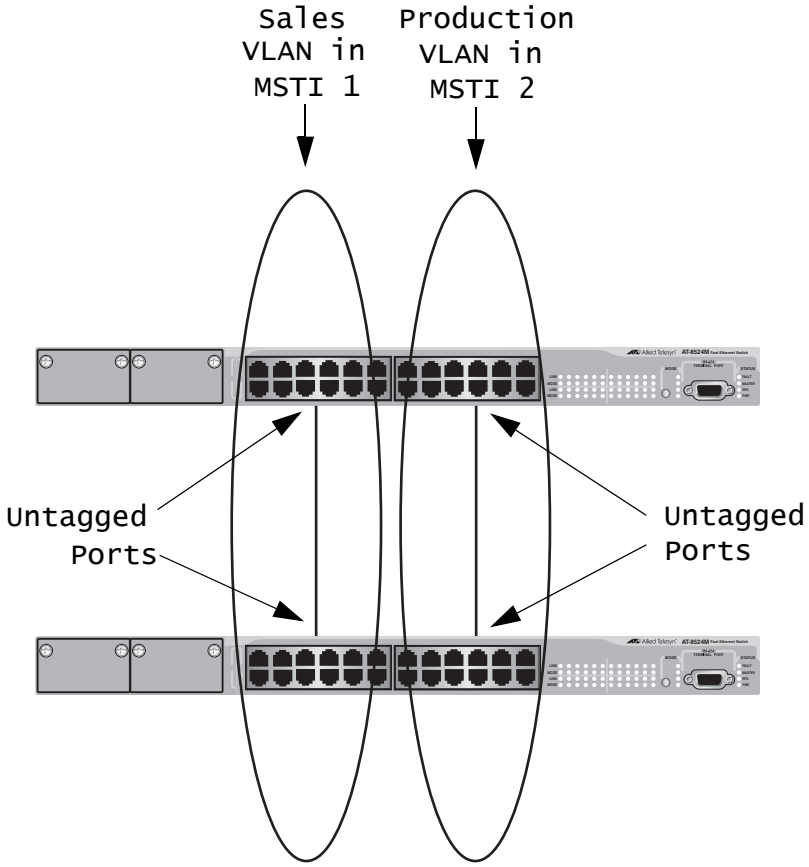

Figure 172. MSTP Example of Two Spanning Tree Instances

A MSTI can contain more than one VLAN. This is illustrated in Figure 173 where there are two AT-8524M switches with four VLANs. There are two MSTIs, each containing two VLANs. MSTI 1 contains the Sales and Presales VLANs and MSTI 2 contains the Design and Engineering VLANs.

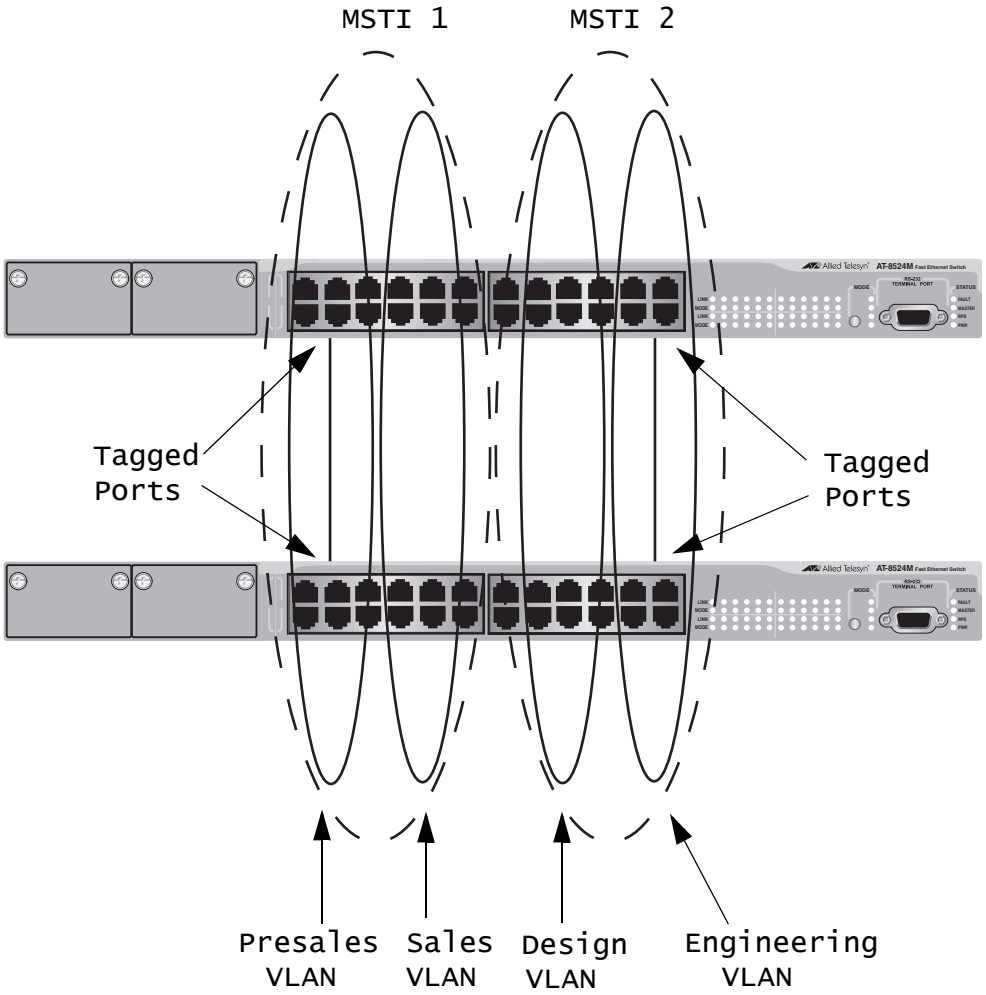

Figure 173. Multiple VLANs in a MSTI

You should note in this example that since an MSTI contains more than one VLAN, the links between the VLAN parts is made with tagged, not untagged, ports so that they can carry traffic from more than one virtual LAN. Referring again to Figure 173, the tagged link in MSTI 1 is carrying traffic for both the Presales and Sales VLANs while the tagged link in MSTI 2 is carrying traffic for the Design and Engineering VLANs.

MSTI Guidelines

Here are several guidelines to keep in mind about MSTIs:

- An AT-8500 Series can support up to 16 spanning tree instances, including the CIST, at a time.
- A MSTI can contain any number of VLANs.
- A VLAN can belong to only one MSTI at a time.
- A switch port can belong to more than one spanning tree instance at a time. This allows you to assign a port as an untagged and tagged member of VLANs that belong to different MSTIs. What makes this possible is a port's ability to be in different MSTP states for different MSTIs. For example, a port can be in the MSTP blocking state for one MSTI and the forwarding state for another spanning tree instance, simultaneously.
- A router or Layer 3 network device is required to forward traffic between different VLANs.

VLAN and MSTI Associations

Part of the task to configuring MSTP involves assigning VLANs to spanning tree instances. The mapping of VLANs to MSTIs is called *associations*. A VLAN, either port-based or tagged, can belong to only one instance at a time, but an instance can contain any number of VLANs.

Ports in Multiple MSTIs

The AT-8500 Series switch allows a port to be a member of more than one MSTI at a time. This can happen if a port is a tagged member of one or more VLANs and the VLANs are assigned to different MSTIs. If this occurs, it is possible that a port might be required to operate in different spanning tree states simultaneously, depending on the requirements of the MSTIs. For example, a port that is a member of two VLANs assigned to two different MSTIs might be operating in the forwarding state for one MSTI and in the blocking state for the other.

When you configure a port's MSTI parameter settings you will notice that the parameters are divided into two groups. The first group is referred to as generic parameters. These are set just once on a port, regardless of the number of MSTIs where a port happens to be a member. One of these parameters is the external path cost, which sets the operating cost of the port in situations where it is connected to a device that is outside its region. A port can have only one external path cost even if it belongs to multiple MSTIs. Other generic parameters are used to designate a port as an edge port or a point-to-point port.

The second group can be applied independently on a port for each MSTI where the port is a member. One of the parameters is the internal path cost. This parameter specifies the port's operating cost if it is connected to a bridge that is a part of the same MSTP region. You can give a port a different internal path cost for each MSTI where it is a member. This group also has a parameter for setting port priority, used as a tie breaker when

two or more ports have equal costs to a regional root bridge. Again, as with the internal path cost, you can assign a port a different priority value for each of its MSTIs.

Multiple Spanning Tree Regions

Another important concept of MSTP is *regions*. A MSTP region is defined as a group of bridges that share exactly the same MSTI characteristics. Those characteristics are:

- Configuration name
- Revision level
- VLANs
- VLAN to MSTI ID associations

A *configuration name* is a name you assign to a region to help you identify it. You must assign each bridge in a region exactly the same name, even the same upper and lowercase lettering. Identifying the regions in your network is easier if you choose names that are characteristic of the functions of the nodes and bridges of the region. Examples are Sales Region and Engineering Region.

The *revision level* is an arbitrary number you can assign to a region. You can use the number to keep track of the revision level of a region's configuration. For example, you might use this value to maintain the number of times you revise a particular MSTP region. It is not important that you maintain this number, only that each bridge in a region have the same number.

The bridges of a particular region must also have the same VLANs. The names of the VLANs and the VIDs must be same on all bridges of a region.

Finally, the VLANs in the bridges must be associated to the same MSTIs.

If any of the above information is different on two bridges, MSTP considers the bridges as residing in different regions.

Figure 174 illustrates the concept of regions. It shows one MSTP region consisting of two AT-8524M switches. Each switch in the region has the same configuration name and revision level. The switches also have the same five VLANs and the VLANs are associated with the same MSTIs.

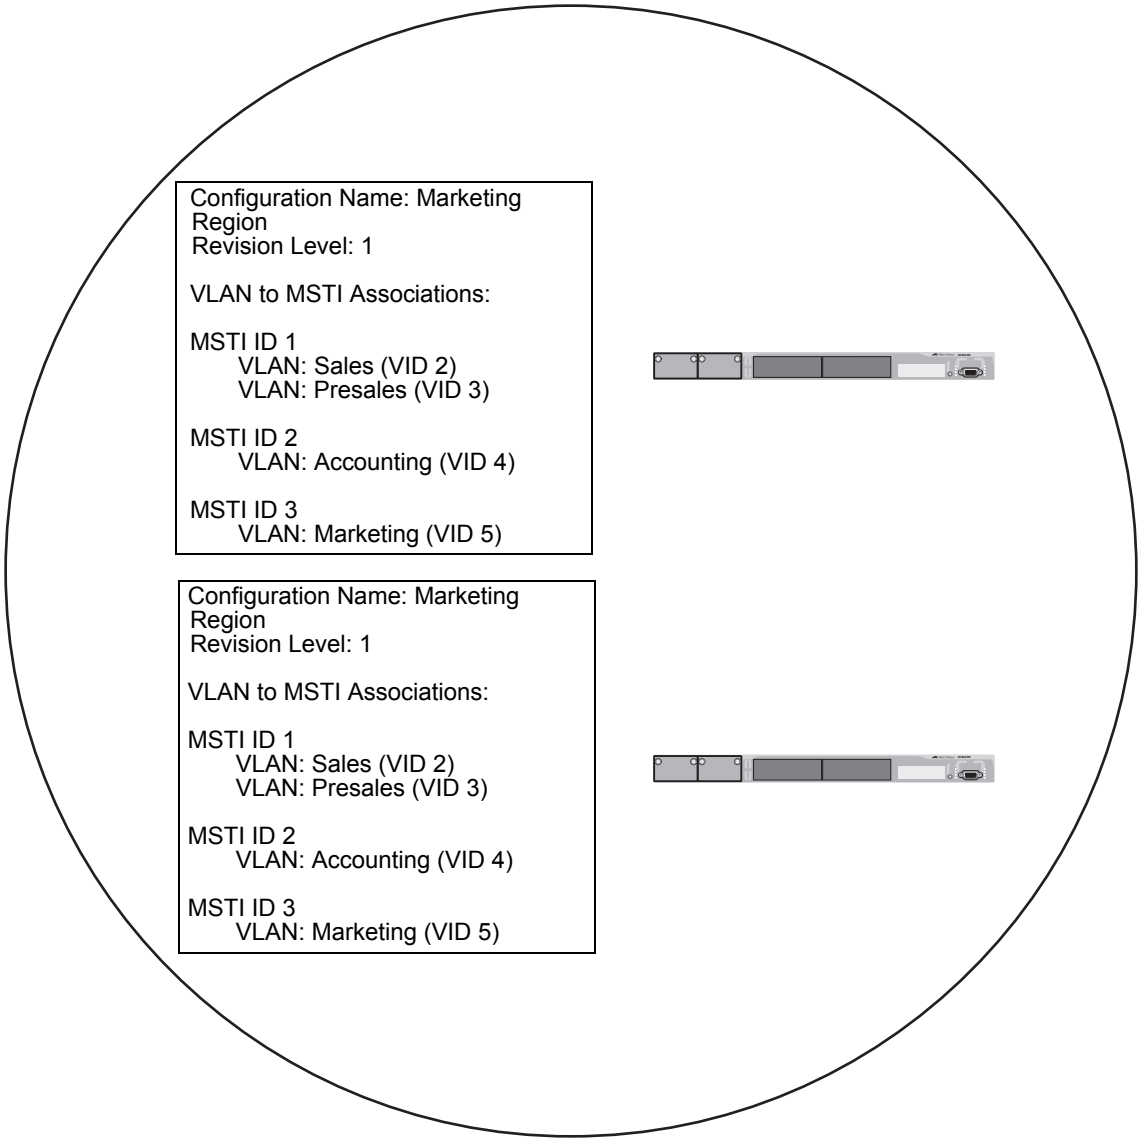

Figure 174. Multiple Spanning Tree Region

The AT-8500 Series switch determines regional boundaries by examining the MSTP BPDUs received on the ports. A port that receives a MSTP BPDU from another bridge with regional information different from its own is considered to be a boundary port and the bridge connected to the port as belonging to another region.

The same is true for any ports connected to bridges running the single-instance spanning tree STP or RSTP. Those ports are also considered as part of another region.

Each MSTI functions as an independent spanning tree within a region. Consequently, each MSTI must have a root bridge to locate physical loops within the spanning tree instance. An MSTI's root bridge is called a *regional root*. The MSTIs within a region may share the same regional root or they can have different regional roots.

A regional root for an MSTI must be within the region where the MSTI is located. An MSTI cannot have a regional root that is outside its region.

A regional root is selected by a combination of the *MSTI priority* value and the bridge's MAC address. The MSTI priority is analogous to the RSTP bridge priority value. Where they differ is that while the RSTP bridge priority is used to determine the root bridge for an entire bridged network, MSTI priority is used only to determine the regional root for a particular MSTI.

The range for this parameter is the same as the RSTP bridge priority; from 0 to 61,440 in sixteen increments of 4,096. To set the parameter, you specify the increment that represents the desired MSTI priority value. Table 11 on page 485 lists the increments.

Region Guidelines

Here are several points to remember about regions.

- A network can contain any number of regions and a region can contain any number of switches that support MSTP.
- An AT-8500 Series switch can belong to only one region at a time.
- A region can contain any number of VLANs.
- All of the bridges in a region must have the same configuration name, revision level, VLANs, and VLAN to MSTI associations.
- An MSTI cannot span multiple regions.
- Each MSTI must have a regional root for locating loops in the instance. MSTIs can share the same regional root or have different roots. A regional root is determined by the MSTI priority value and a bridge's MAC address.
- The regional root of a MSTI must be in the same region as the MSTI.

Common and Internal Spanning Tree (CIST)

MSTP has a default spanning tree instance called the Common and Internal Spanning Tree (CIST). This instance has an MSTI ID of 0.

This instance has unique features and functions that make it different from the MSTIs that you create yourself. First, you cannot delete this instance and you cannot change its MSTI ID.

Second, when you create a new port-based or tagged VLAN, it is by default associated with the CIST and is automatically given an MSTI ID of 0. The Default_VLAN is also associated by default with CIST.

Another critical difference is that when you assign a VLAN to another MSTI, it still partially remains a member of CIST. This is because CIST is used by MSTP to communicate with other MSTP regions and with any RSTP and STP single-instance spanning trees in the network. MSTP uses CIST to participate in the creation of a spanning tree between different regions and between regions and single-instance spanning tree, to form one spanning tree for the entire bridged network.

The reason MSTP uses CIST to form the spanning tree of an entire bridged network is because CIST can cross regional boundaries, while a MSTI cannot. If a port is a boundary port, that is, if it is connected to another region, that port automatically belongs solely to CIST, even if it was assigned to an MSTI, because only CIST is active outside of a region.

As mentioned earlier, every MSTI must have a root bridge, referred to as a regional root, in order to locate loops within the instance. CIST must also have a regional root. However, the CIST regional root communicates with the other MSTP regions and single-instance spanning trees in the bridged network.

The CIST regional root is set with the *CIST Priority* parameter. This parameter, which functions similar to the RSTP bridge priority value, is used to select the root bridge for the entire bridged network. If an AT-8500 Series switch has the lowest CIST Priority value among all the spanning tree bridges, it functions as the root bridge for all the MSTP regions and STP and RSTP single-instance spanning trees in the network.

MSTP with STP and RSTP

MSTP is fully compatible with STP and RSTP. If a port on an AT-8500 Series switch running MSTP receives STP BPDUs, the port sends only STP BPDU packets. If a port receives RSTP BPDUs, the port sends MSTP BPDUs because RSTP can process MSTP BPDUs.

A port connected to a bridge running STP or RSTP is considered a boundary port of the MSTP region and the bridge as belonging to a different region.

An MSTP region can be considered as a virtual bridge. The implication is that other MSTP regions and STP and RSTP single-instance spanning trees cannot discern the topology or constitution of a MSTP region. The only bridge they are aware of is the regional root of the CIST instance.

Summary of Guidelines

Careful planning is essential for the successful implementation of MSTP. This section reviews all the rules and guidelines mentioned in earlier sections, plus a few new ones:

- An AT-8500 Series switch can support up to 16 spanning tree instances, including the CIST, at a time.
- A MSTI can contain any number of VLANs.
- A VLAN can belong to only one MSTI at a time.
- An MSTI ID can be from 1 to 15.
- The CIST ID is 0. You cannot change this value.
- A switch port can belong to more than one spanning tree instance at a time. This allows you to assign a port as a untagged and tagged member of VLANs that belong to different MSTIs. What makes this possible is a port's ability to be in different MSTP states for different MSTIs simultaneously. For example, a port can be in the MSTP blocking state for one MSTI and the forwarding state for another spanning tree instance.
- A router or Layer 3 network device is required to forward traffic between VLANs.
- A network can contain any number of regions and a region can contain any number of AT-8500 Series switches.
- An AT-8500 Series switch can belong to only one region at a time.
- A region can contain any number of VLANs.
- All of the bridges in a region must have the same configuration name, revision level, VLANs, and VLAN to MSTI associations.
- An MSTI cannot span multiple regions.
- Each MSTI must have a regional root for locating loops in the instance. MSTIs can share the same regional root or have different roots. A regional root is determined by the MSTI priority value and a bridge's MAC address.
- The regional root of a MSTI must be in the same region as the MSTI.

- The CIST must have a regional root for communicating with other regions and single-instance spanning trees.
- MSTP is compatible with STP and RSTP.
- A port transmits CIST information even when it's associated with another MSTI ID. However, in determining network loops, MSTI takes precedence over CIST. (This is explained more in "Associating VLANs to MSTIs" on page 518.)

Note
The AT-S62 implementation of MSTP complies with the IEEE 802.1s standard and is compatible with versions from other vendors that conform to the standard.

Associating VLANs to MSTIs

Allied Telesyn recommends that you assign all VLANs on a switch to an MSTI. You should not leave a VLAN assigned to just the CIST, including the Default_VLAN. This is to prevent the blocking of a port that should be in the forwarding state. The reason for this guideline is explained here.

An MSTP BPDU contains the instance to which the port transmitting the packet belongs. By default, all ports belong to the CIST instance. So CIST would be included in the BPDU. If the port is a member of a VLAN that has been assigned to another MSTI, that information is also included in the BPDU.

This is illustrated in Figure 175. Port 8 in Switch A is a member of a VLAN assigned to MSTI ID 7 while Port 1 is a member of a VLAN assigned to MSTI ID 10. The BPDUs transmitted by port 8 to Switch B would indicate that the port is a member of both CIST and MSTI 7, while the BPDUs from Port 1 would indicate the port is a member of the CIST and MSTI 10.

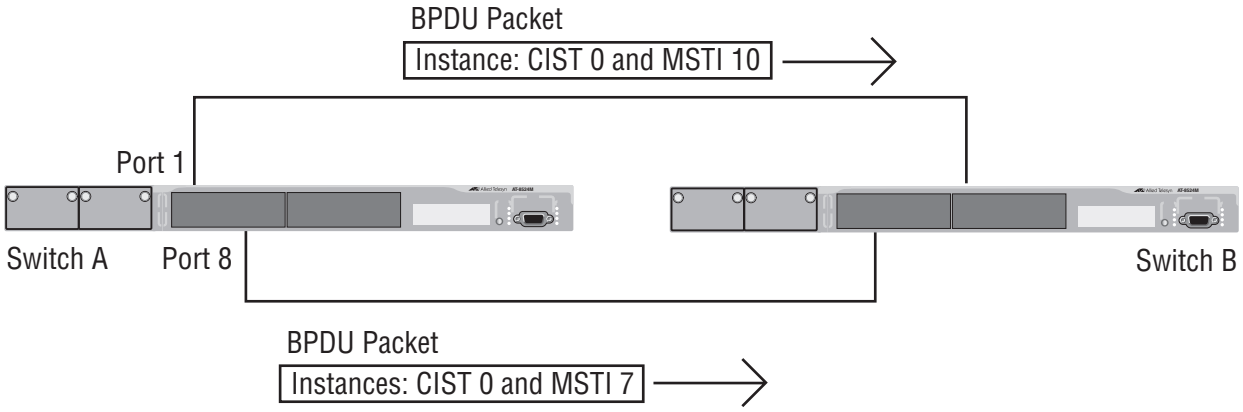

Figure 175. CIST and VLAN Guideline - Example 1

At first glance, it might appear that since both ports belong to CIST, a loop would exist between the switches and that MSTP would block a port to stop the loop. However, within a region, MSTI takes precedence over CIST. When Switch B receives a packet from Switch A, it uses MSTI, not CIST, to determine whether a loop exists. And since both ports on Switch A belong to different MSTIs, Switch B determines that no loop exists.

A problem can arise if you assign some VLANs to MSTIs while leaving others just to CIST. The problem is illustrated in Figure 176. The network is the same as the previous example. The only difference is that the VLAN containing Port 8 on Switch A has not been assigned to an MSTI, and belongs only to CIST with its MSTI ID 0.

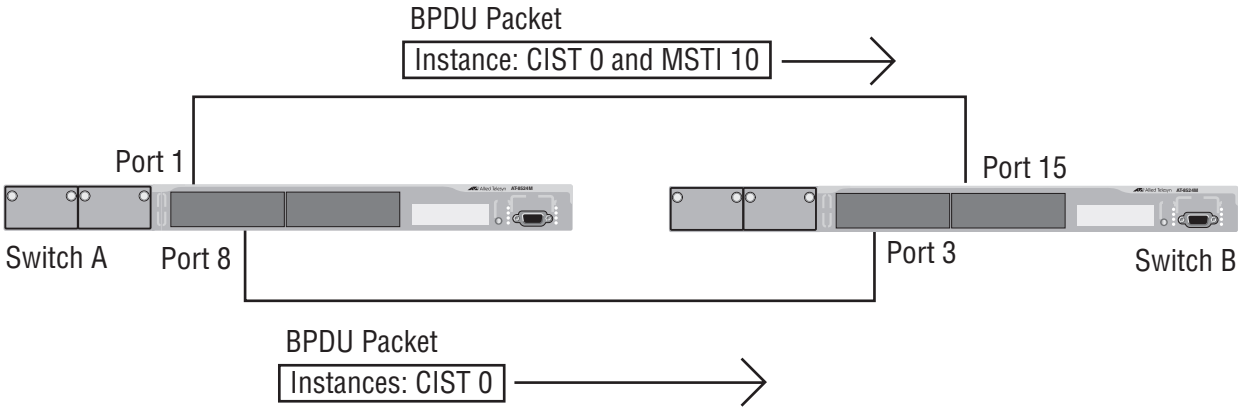

Figure 176. CIST and VLAN Guideline - Example 2

When port 3 on Switch B receives a BPDU, the switch notes the port sending the packet belongs only to CIST. Consequently, Switch B uses CIST in determining whether a loop exists. The result would be that the switch would determine that a loop exists because the other port is also receiving BPDU packets from CIST 0. Switch B would block a port to cancel the loop.

To avoid this issue, always assign all VLANs on a switch, including the Default_VLAN, to an MSTI. This guarantees that all ports on the switch have an MSTI ID and that helps to ensure that loop detection is based on MSTI, not CIST.

Connecting VLANs Across Different Regions

Special consideration needs to be taken into account when connecting different MSTP regions or an MSTP region and a single-instance STP or RSTP region. Unless planned properly, VLAN fragmentation can occur between the VLANS of your network.

As mentioned previously, only the CIST can span regions. A MSTI cannot. Consequently, you may run into a problem if you use more than one physical data link to connect together various parts of VLANs that reside in

bridges in different regions. The result can be a physical loop, which spanning tree disables by blocking ports.

This is illustrated in Figure 177. The example show two switches, each residing in a different region. Port 5 in Switch A is a boundary port. It is an untagged member of the Accounting VLAN, which has been associated with MSTI 4. Port 15 is a tagged and untagged member of three different VLANs, all associated to MSTI 12.

If both switches were a part of the same region, there would be no problem since the ports reside in different spanning tree instances. However, the switches are part of different regions and MSTIs do not cross regions. Consequently, the result would be that spanning tree would determine that a loop exists between the regions, and Switch B would block a port.

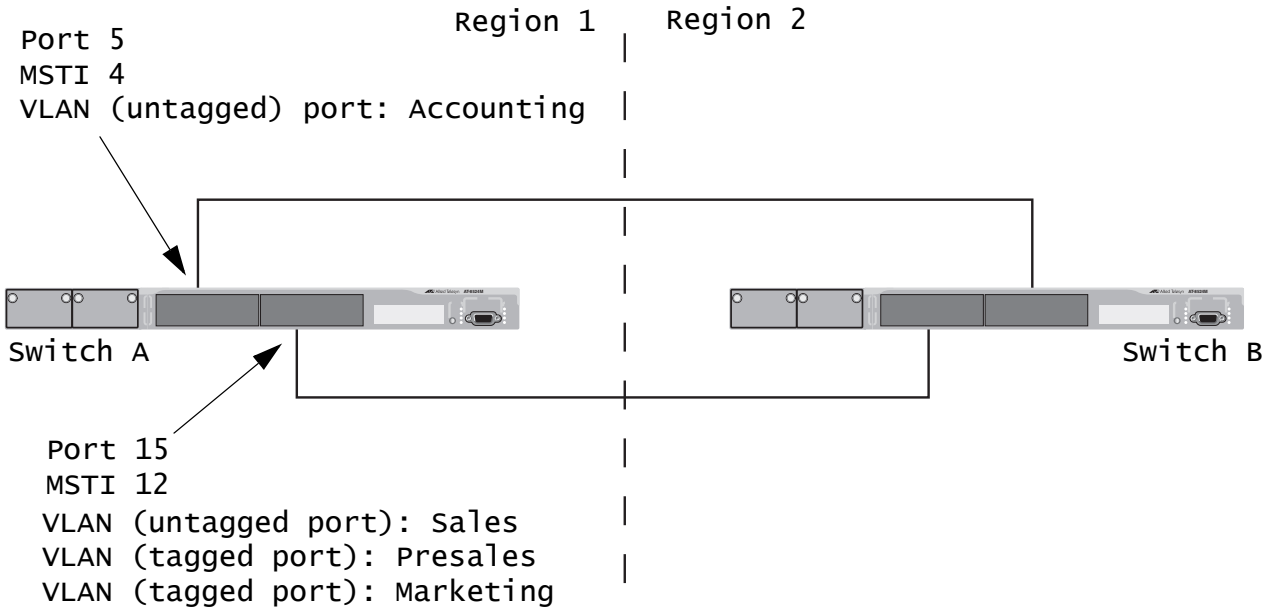

Figure 177. Spanning Regions - Example 1

There are several ways to address this issue. One is to have only one MSTP region for each subnet in your network.

Another approach is to group those VLANs that need to span regions into the same MSTI. Those VLANs that do not span regions can be assigned to other MSTIs.

Here is an example. Let's assume that you have two regions that contain the following VLANS:

Region 1 VLANs	Region 2 VLANs
Sales	Hardware Engineering
Presales	Software Engineering
Marketing	Technical Support
Advertising	Product Management
Technical Support	CAD Development
Product Management	Accounting
Project Management	
Accounting	

The two regions share three VLANs: Technical Support, Product Management, and Accounting. You could group those VLANs into the same MSTI in each region. For instance, for Region 1 you might group the three VLANs in MSTI 11 and in Region 2 you could group them into MSTI 6. Once grouped, you can connect the VLANs across the regions using a link of tagged ports.

Selecting MSTP as the Active Spanning Tree Protocol

To select and activate MSTP as the active spanning tree protocol on the switch, or to disable spanning tree, perform the following procedure:

1. From the Main Menu, type **3** to select Spanning Tree Configuration.

   The Spanning Tree Menu is shown in Figure 163 on page 493.

2. To change the active version of spanning tree on the switch, type **2** to select Active Protocol Version.

   The following prompt is displayed:

   ```
   Enter new value (S-STP, R-RSTP, M-MSTP):
   ```

3. Type **M** to select MSTP.

   Note
   A change to the active spanning tree is automatically saved on the switch.

   Now that you have selected MSTP as the active spanning tree on the switch, you can configure its parameter settings.

4. To enable or disable spanning tree, type **1** to select Spanning Tree Status.

   The following prompt is displayed:

   ```
   Enter new value (E-Enable, D-Disable):
   ```

5. Type **E** to enable spanning tree or **D** to disable it. The default is disabled.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring MSTP Bridge Settings

This section contains the procedure for configuring a bridge's MSTP settings.

Note
You cannot configure the MSTP parameters until you have selected the protocol as the active spanning tree protocol on the switch. For instructions, refer to "Selecting MSTP as the Active Spanning Tree Protocol" on page 522.

1. From the Main Menu, type **3** to select Spanning Tree Menu.

   The Spanning Tree Menu is shown in Figure 163 on page 493.

2. From the Spanning Tree Menu, type **3** to select Configure Active Protocol.

   The MSTP Menu is shown in Figure 178.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                    11:20:02 02-Jan-2006
                                 MSTP Menu
1 - Force Version .......... MSTP
2 - Hello Time ............. 2/2 (Configured/Actual)
3 - Forwarding Delay ....... 15/15 (Configured/Actual)
4 - Max Age ................ 20/20 (Configured/Actual)
5 - Max Hops ............... 20
6 - Configuration Name .....
7 - Revision Level ......... 0
8 - Bridge Identifier ...... 32768/00:30:84:00:00:02
9 - Root Identifier ........ 32768/00:30:84:00:00:02
A - Root Path Cost ......... 0

C - CIST Menu
M - MSTI Menu
V - VLAN-MSTI Association Menu
P - MSTP Port Parameters

R - Return to Previous Menu

Enter your selection?
```

Figure 178. MSTP Menu

Menu selections 1 to 9 are described below. Selections C, M, V, and P are described in later sections in this chapter.

The hello time, forwarding delay, and max age parameters will have two values if MSTP is enabled on the switch (for example, Forwarding Delay .. 15/15). The first number is the configured value on the switch for the parameter and the second is the value the switch obtained from the root bridge and is actually using for the parameter. The switch displays only the configured values for these parameters if multiple spanning tree is not enabled on the switch.

3. Adjust the MSTP settings as needed. Changes are immediately activated on the switch. The selections are described below.

1 - Force Version
This selection determines whether the bridge operates with MSTP or in an STP-compatible mode. If you select MSTP, the bridge operates all ports in MSTP, except for those ports that receive STP or RSTP BPDU packets. If you select Force STP Compatible, the bridge uses its MSTP parameter settings, but sends only STP BPDU packets from the ports.

2 - Hello Time
The time interval between generating and sending configuration messages by the bridge. The range of this parameter is 1 to 10 seconds. The default is 2 seconds. This value is active only if the bridge is selected as the root bridge of the network.

3 - Forwarding Delay
The waiting period before a bridge changes to a new state, for example, becomes the new root bridge after the topology changes. If the bridge transitions too soon, not all links may have yet adapted to the change, possibly resulting in a network loop. The range is 4 to 30 seconds. The default is 15 seconds. This setting applies only to ports running in the STP-compatible mode.

4 - Max Age
The length of time after which stored bridge protocol data units (BPDUs) are deleted by the bridge. This parameter applies only if the bridged network contains an STP or RSTP single-instance spanning tree. Otherwise, the bridges use the Max Hop counter to delete BPDUs.

All bridges in a single-instance bridged LAN use this aging time to test the age of stored configuration messages called bridge protocol data units (BPDUs). For example, if you use the default of 20, all bridges delete current configuration messages after 20 seconds. The range of this parameter is 6 to 40 seconds. The default is 20 seconds.

In selecting a value for maximum age, the following must be observed:

MaxAge must be greater than (2 x (HelloTime + 1))

MaxAge must be less than (2 x (ForwardingDelay - 1))

5 - Max Hops
MSTP regions use this parameter to discard BPDUs. The Max Hop counter in a BPDU is decremented every time the BPDU crosses an MSTP region boundary. Once the counter reaches zero, the BPDU is deleted. The range is 1 to 40 hops. The default is 20.

6 - Configuration Name
The name of the MSTP region. The range is 0 (zero) to 32 alphanumeric characters in length. The name, which is case-sensitive, must be the same on all bridges in a region. Examples include Sales Region and Production Region.

7 - Revision Level
The revision level of an MSTP region. The range is 0 (zero) to 255. This is an arbitrary number that you assign to a region. The revision level must be the same on all bridges in a region. Different regions can have the same revision level without conflict.

8 - Bridge Identifier
The bridge identifier of the switch. The identifier consists of the switch's CIST priority value and MAC address, separated by a slash (/). To change the switch's priority value, refer to "Configuring the CIST Priority" on page 526. The MAC address of the switch cannot be changed.

9 - Root Identifier
The bridge identifier of the root bridge of the CIST spanning tree domain. The identifier consists of the root switch's bridge or CIST priority value and MAC address, separated by a slash (/). If this MAC address is the same as the current bridge's MAC address, then the switch is functioning as a root bridge. If the two MAC addresses are different, then a different switch is functioning as the root bridge. This parameter is only displayed with MSTP is enabled.

A - Root Path Cost
The cost of the path from the current switch to the root switch of the spanning tree domain. If the current switch is the root switch, root path cost will be "0". This value cannot be changed and is only displayed when MSTP is activated on the switch.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring the CIST Priority

This procedure explains how to adjust the bridge's CIST priority.

> **Note**
> You cannot configure MSTP parameters until you have selected the protocol as the active spanning tree protocol on the switch. For instructions, refer to "Selecting MSTP as the Active Spanning Tree Protocol" on page 522.

This procedure starts from the MSTP Menu. If you do not know how to access the menu, perform steps 1 and 2 in "Configuring MSTP Bridge Settings" on page 523.

To change the CIST priority, do the following:

1. From the MSTP Menu, type to select **C** to select CIST Menu.

   The CIST Menu is shown in Figure 179.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                         Production Switch
User: Manager                                    11:20:02 02-Jan-2006
                        CIST Configuration
CIST Priority ..............  32768
Associated VLANs ...........  1,2,4,11

1 - Modify CIST Priority

R - Return to Previous Menu
Enter your selection?
```

Figure 179. CIST Configuration Menu

   The CIST Priority field in the menu displays the current value for this MSTP parameter. This number is used in determining the root bridge of the network spanning tree. This number is analogous to the RSTP bridge priority value. The bridge in the network with the lowest priority number is selected as the root bridge. If two or more bridges have the same bridge or CIST priority values, the bridge with the numerically lowest MAC address becomes the root bridge.

   The Associated VLANs field displays the VIDs of the VLANs that are currently associated with CIST and have not been assigned to a MSTI.

2. To change the CIST priority, type **1**.

The following prompt is displayed:

```
Enter new priority [the value will be multiplied by
4096]: [0 to 15] ->
```

3. Enter the increment that represents the new CIST priority value. The range is 0 (zero) to 61,440 in increments of 4,096, with 0 being the highest priority. For a list of the increments, refer to Table 11, "Bridge Priority Value Increments" on page 485.

   The change is immediately implemented on the switch.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Creating, Deleting, and Modifying MSTIs

The following procedures explain how to create, delete, and modify spanning tree instances.

> **Note**
> You cannot configure MSTP parameters until you have selected the protocol as the active spanning tree protocol on the switch. For instructions, refer to “Selecting MSTP as the Active Spanning Tree Protocol” on page 522.

This procedure starts from the MSTP Menu. If you do not know how to access the menu, perform steps 1 and 2 in “Configuring MSTP Bridge Settings” on page 523.

1. From the MSTP Menu, type **M** to select MSTI Menu.

   The MSTI Configuration menu is shown in Figure 180.

```
           Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                    11:20:02 02-Jan-2006
                            MSTI Configuration
MSTI  |  Priority |  Regional Root ID |   Path Cost |  Associated VLANs
-------------------------------------------------------------------
1       32768      00A0D2 1454B3       0            1,2
2       32768      00A0D2 1454B3       0            4,11
1 - Create MSTI
2 - Delete MSTI
3 - Modify MSTI

U - Update Display
R - Return to Previous Menu
Enter your selection?
```

Figure 180. MSTI Configuration Menu

The fields in the table are defined below:

MSTI
Lists the MSTI IDs existing on the switch.

Priority
Specifies the MSTI priority value for the MSTI. The steps in this procedure explain how you can assign this value when you create an MSTI ID and how to modify the value for an existing MSTI ID.

Regional Root ID
Identifies the regional root for the MSTI by its MAC address.

Path Cost
Specifies the path cost from the bridge to the regional root. If the bridge is the regional root, the value is 0.

Associated VLANs
Specifies the VIDs of the VLANs that have been associated with the MSTI ID.

The table does not include the CIST. The table is empty if no MSTI IDs have been created.

Creating an MSTI

To create an MSTI, do the following:

1. From the MSTI Menu, type **1** to select Create MSTI.

   The following prompt is displayed:

   ```
   Enter the MSTI ID to be created: [1 to 15] ->
   ```

2. Enter an ID number for the new MSTI. The range is 1 to 15. You can create only one MSTI at a time.

   The following prompt is displayed:

   ```
   Enter new priority [the value will be multiplied by 4096]
   [0 to 15] -> 8
   ```

3. Enter a MSTI priority number for the new MSTI. This parameter is used in selecting a regional root. The range is 0 (zero) to 61,440 in increments of 4,096, with 0 being the highest priority. For a list of the increments, refer to Table 11, "Bridge Priority Value Increments" on page 485.

   The following prompt is displayed:

   ```
   Enter the list of VLANs to associate with this MSTI:
   ```

4. Enter the VIDs of the VLANs you want to associate with this MSTI. You can specify more than one VLAN at a time (for example, 4,6,11) To view VIDs, refer to "Displaying VLANs" on page 569. If you do not want to associate any VLANs with the MSTI at this time, just press Return.

   The MSTI is created by the switch and is activated immediately.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an MSTI

To delete an MSTI, do the following:

1. From the MSTI Menu, type **2** to select Delete MSTI.

   The following prompt is displayed:

   ```
   Enter the MSTI ID to be deleted: [1 to 15] ->
   ```

2. Enter the ID number of the MSTI you want to delete. The range is 1 to 15. (You cannot delete CIST, which has a value of 0.) You can delete only one MSTI at a time.

   The selected MSTI is deleted from the switch. All associated VLANs are returned to CIST.

3. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an MSTI

To change the priority value of an MSTI or the associated VLANs, do the following:

1. From the MSTI Menu, type **3** to select MSTI Configuration Menu.

   The following prompt is displayed:

   ```
   Enter the MSTI ID to be modified: [1 to 15] ->
   ```

2. Enter the ID number of the MSTI you want to modify. The range is 1 to 15. You can modify only one MSTI at a time.

   The following prompt is displayed:

   ```
   Enter new priority [the value will be multiplied by 4096]
   [0 to 15] -> 8
   ```

3. Enter a new MSTI priority number. This parameter is used in selecting a regional root. The range is 0 (zero) to 61,440 in increments of 4,096, with 0 being the highest priority. The default is increment 8 for a value of 32768. For a list of the increments, refer to Table 11, “Bridge Priority Value Increments” on page 485.

   The following prompt is displayed:

   ```
   Enter the list of VLANs to associate with this MSTI:
   ```

4. Enter the VIDs of the VLANs you want to associate with this MSTI. You can specify more than one VLAN at a time (for example, 4,6,11). The new VLAN associates overwrite the current VLAN associations. To add new VLANs while retaining the existing VLAN associations, you must enter the VIDs of the new and existing VLANs. To remove VLAN associations, reenter the VLAN ID list, omitting the VIDs of

those VLANs you no longer want associated with the MSTI. If you do not want to change the current associates, just press Return. To view the VIDs of the VLANs on the switch, refer to "Displaying VLANs" on page 569.

The MSTI modifications are immediately activated on the switch.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Associating VLANs to MSTI IDs

When you create a new MSTI, you are given the opportunity to associate VLANs to it. But once a MSTI is created, there might come a time when you want to add more VLANs, or perhaps remove VLANs from it. This procedure explains how to associate VLANs on the switch to an existing MSTI and also how to remove VLANs. Before performing this procedure, note the following:

- You must create an MSTI before you can assign VLANs to it. To create a MSTI ID, refer to “Creating, Deleting, and Modifying MSTIs” on page 528.
- You can assign a VLAN to only one MSTI. By default, a VLAN, when created, is associated with the CIST instance, which has a MSTI ID of 0.
- An MSTI can contain any number of VLANs.
- You can also associate VLANs to an MSTI by performing the procedure “Modifying an MSTI” on page 530.

Note
You cannot configure MSTP parameters until you have selected the protocol as the active spanning tree protocol on the switch. For instructions, refer to “Selecting MSTP as the Active Spanning Tree Protocol” on page 522.

This procedure starts from the MSTP Menu. If you do not know how to access the menu, perform steps 1 and 2 in “Configuring MSTP Bridge Settings” on page 523.

To add or remove a VLAN from an MSTI, do the following:

1. From the MSTP Menu, type **V** to select VLAN-MSTI Association Menu.

The VLAN-MSTI Association Menu is shown in Figure 181.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                      11:20:02 02-Jan-2006
                        VLAN-MSTI Association
MSTI/CIST      Associated VLANs
------------------------------------------------------------

0
4              1,2
5              6
7              7,22


1 - Add VLANs to MSTI
2 - Delete VLANs from MSTI
3 - Set VLAN to MSTI association
4 - Clear VLAN to MSTI association

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 181. VLAN-MSTI Association Menu

The fields in the table are defined below:

MSTI / CIST
Lists the CIST and existing MSTI IDs on the switch.

Associated VLANs
Specifies the VIDs of the VLANs associated with the CIST and MSTIs. For instance, referring to the figure above, the VLANs with the VIDs 7 and 22 are assigned to MSTI 7.

Adding VLAN Associations to an MSTI

This procedure adds new VLANs associations to an MSTI while retaining the existing associations. If you want to add VLAN associations but not retain the existing ones, perform the procedure "Replacing VLAN Associations to an MSTI" on page 534.

To associate a VLAN to an MSTI, do the following:

1. From the VLAN-MSTI Association Menu, type **1** to select Add VLANs to MSTI Association.

   The following prompt is displayed:

   ```
   Enter the MSTI ID <enter 0 for CIST> [0 to 15] ->
   ```

2. Enter the MSTI ID to which you want to associate a VLAN.

A prompt similar to the following is displayed:

```
Enter the list of VLANs:
```

3. Enter the VLAN ID of the virtual LAN you want to associate with the MSTI. You can enter more than one VLAN at a time (for example, 2,4,7). The new VLAN associations are added to the existing associations in the MSTI. To view VIDs, refer to “Displaying VLANs” on page 569.

   New VLAN associations are immediately implemented on the switch.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Removing VLAN Associations from an MSTI

To remove a VLAN from an MSTI, do the following:

1. From the VLAN-MSTI Association Menu, type **2** to select Delete VLANs from MSTI.

   The following prompt is displayed:

   ```
   Enter the MSTI ID <enter 0 for CIST> [0 to 15] ->
   ```

2. Enter the ID number of the MSTI where you want to remove a VLAN associate. (You cannot remove VLANs from CIST using this procedure. To remove a VLAN from CIST, you must assign it to an MSTI.)

   A prompt similar to the following is displayed:

   ```
   Enter the list of VLANs:
   ```

3. Enter the VID of the virtual LAN you want removed from the MSTI. You can specify more than one VLAN at a time (for example, 2,4,7) To view the VIDs of the VLANs on the switch, refer to “Displaying VLANs” on page 569.

   A VLAN removed from an MSTI is automatically returned to CIST.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Replacing VLAN Associations to an MSTI

To associate VLANs to an MSTI while removing the existing VLAN associates, do the following:

1. From the VLAN-MSTI Association Menu, type **3** to select Set VLANs to MSTI Association.

   The following prompt is displayed:

   ```
   Enter the MSTI ID <enter 0 for CIST> [0 to 15] ->
   ```

2. Enter the ID number of the MSTI you want to associate a VLAN.
3. A prompt similar to the following is displayed:

   ```
   Enter the list of VLANs:
   ```

4. Enter the VLAN ID of the virtual LAN that you want to associate with the MSTI. You can enter more than one VLAN at a time (for example, 2,4,7) (To view VIDs, refer to "Displaying VLANs" on page 569.)

   The existing VLANs associations are removed from the MSTI when the new VLANs are added. The removed VLANs are returned to CIST.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Removing All VLAN Associations from an MSTI

To remove all VLAN associations from an MSTI, do the following:

1. From the VLAN-MSTI Association Menu, type **4** to select Clean VLAN to MSTI Association.

   The following prompt is displayed:

   ```
   Enter the MSTI ID [1 to 15] ->
   ```

2. Enter the ID number of the MSTI whose VLAN associations you want to remove. (You cannot remove VLANs from CIST using this procedure. To remove a VLAN from CIST, you must assign the VLAN to an MSTI.)

   All VLAN associations are immediately removed from the MSTI and are returned to CIST.

3. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring MSTP Port Settings

As explained in “Ports in Multiple MSTIs” on page 512, MSTP port settings are divided into two groups. The parameters in the first group are set just once on a port. The setting for a generic port parameter applies to all MSTIs in which the port is a member. These settings are:

- External path cost
- Point-to-point designation
- Edge port designation

The procedure for setting these parameters is in “Configuring Generic MSTP Port Settings” on page 536.

The second group of port parameters can be set independently for each MSTI in which the port is a member. This means that you can assign a port a different value to an MSTI-specific parameter for each spanning tree instance where the port is a member. These parameters are:

- Internal path cost
- Priority

To set these parameters, refer to “Configuring MSTI-specific Port Parameters” on page 538.

Configuring Generic MSTP Port Settings

To configure the external path cost of a port or to designate whether the port is an edge or point-to-point port, perform the following procedure:

1. From the MSTP Menu, type **P** to select MSTP Port Parameters.

   The MSTP Port Parameters menu is shown in Figure 182.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                      11:20:02 02-Jan-2006
                         MSTP Port Parameters
1 - Configure Generic Port Settings
2 - Configure Per Spanning Tree Port Settings
3 - Display MSTP Port Configuration
4 - Display MSTP Port State

R - Return to Previous Menu

Enter your selection?
```

Figure 182. MSTP Port Parameters Menu

2. Type **1** to select Configure Generic Port Settings.

   The following prompt is displayed:

   ```
   Start port to configure:  [1 to 26] ->
   ```

3. Enter the number of the port you want to configure. To configure a range of ports, enter the first port of the range.

   The following prompt is displayed:

   ```
   End port to configure:  [1 to 26] -> 4
   ```

4. Enter the last port of the range. To configure just one port, enter the same port here as in Step 3.

   The Configure Generic Port Settings menu is shown in Figure 183.

```
        Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                  11:20:02 02-Jan-2006
                  Configure Generic Port Settings
1 - Port External Path Cost ..... Auto
2 - Point-to-Point ............... Auto Detect
3 - Edge Port .................... Yes

R - Return to Previous Menu

Enter your selection?
```

Figure 183. Configure MSTP Port Settings Menu

5. Adjust the port settings as needed. The parameters are described below:

   1- Port External Path Cost
   The port cost of the port if the port is connected to a bridge which is a member of another MSTP region or is running STP or RSTP. The range is 0 to 200,000,000. The default setting is Auto, which sets port cost depending on the speed of the port. Table 17 lists the MSTP port costs with the Auto setting when the port is not a member of a trunk.

Table 17. Auto External Path Costs

Port Speed	Port Cost
10 Mbps	2,000,000
100 Mbps	200,000
1000 Mbps	20,000

Table 18 lists the MSTP port costs with the Auto setting when the port is part of a port trunk.

Table 18. Auto External Path Trunk Costs

Port Speed	Port Cost
10 Mbps	20,000
100 Mbps	20,000
1000 Mbps	2,000

2 - Point-to-Point
This parameter defines whether the port is functioning as a point-to-point port. For an explanation of this parameter, refer to “Point-to-Point Ports and Edge Ports” on page 489.

3 - Edge Port
This parameter defines whether the port is functioning as an edge port. For an explanation of this parameter, refer to “Point-to-Point Ports and Edge Ports” on page 489.

Parameter changes are immediately activated on the port.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring MSTI-specific Port Parameters

This procedure explains how to set a port’s priority and internal path cost. These parameters can be set independently on a port for each MSTI where a port is a member. To configure the parameters, perform the following procedure:

1. From the MSTP Menu, type **P** to select MSTP Port Parameters.

   The MSTP Port Parameters menu is shown in Figure 182 on page 536.

2. Type **2** to select Configure Per Spanning Tree Port Settings.

   The following prompt is displayed:

   ```
   Enter Spanning Tree (CIST/MSTI) List :
   ```

3. Enter the ID number of the CIST or MSTI containing the VLAN with the port to be configured is assigned. You can specify more than one ID number.

   You can enter the ID number of an MSTI where the port is not a member. This allows you to pre-configure the parameter in the event you later add the port as a member of the MSTI through a VLAN assignment.

The following prompt is displayed:

```
Start port to configure:  [1 to 26] -> 1
```

4. Enter the number of the port to be configured. To configure a range of ports, enter the first port of the range.

   The following prompt is displayed:

```
End port to configure:  [1 to 26] -> 1
```

5. Enter the last port of the range. To configure just one port, enter the same port here as in Step 3.

   Configure Per Spanning Tree Port Settings Menu is shown in Figure 184.

```
        Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                  11:20:02 02-Jan-2006
            Configure Per Spanning Tree Port Settings
Spanning Tree List: 4
Configuring Ports: 7-7

1 - Port Priority ............... 128
2 - Port Internal Path Cost ..... Auto Update

R - Return to Previous Menu

Enter your selection?
```

Figure 184. Configure Per Spanning Tree Port Settings Menu

The Spanning Tree List displays the ID numbers of the MSTIs you specified.

6. Adjust the port settings as needed. The selections are described below:

   1 - Port Priority
   This parameter functions as a tie breaker when two or more ports are determined to have equal costs to the regional root bridge. The range is 0 to 240 in increments of 16. The default value is 8 (priority value 128). For a list of the increments, refer to Table 16, "Port Priority Value Increments" on page 488.

   2- Port Internal Path Cost
   The port cost of the port if the port is connected to a bridge which is part of the same MSTP region. The range is 0 to 200,000,000. The default setting is Auto-detect, which sets port cost depending on the speed of the port. Table 19 lists the MSTP port costs with Auto Update.

Table 19. MSTP Auto Update Port Costs

Port Speed	Port Cost
10 Mbps	2,000,000
100 Mbps	200,000
1000 Mbps	20,000

Table 20 lists the MSTP port costs with Auto Update when the port is part of a port trunk.

Table 20. MSTP Auto Update Port Trunk Costs

Port Speed	Port Cost
10 Mbps	20,000
100 Mbps	20,000
1000 Mbps	2,000

Parameter changes are immediately activated on the port.

7. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying MSTP Port Settings and Status

The MSTP Port Parameters menu, shown in Figure 182 on page 536, has two selections for displaying a variety of MSTP port information. The two menu selections are described below. (To display the menu, from the MSTP Menu, type **P** to select MSTP Port Parameters.)

2 - Display MSTP Port Configuration
This selection displays a menu that contains the current port settings for the following MSTP parameters:

- Edge-Port
- Point-to-Point Port
- External or Internal Port Cost
- Port Priority

3 - Display MSTP Port State
This selection displays a menu that contains the following MSTP operating status for a port:

- State - Identifies the MSTP state of the port. Possible states are: discarding, learning, and forwarding. A state of disabled means the port has not established a link with its end node.
- MSTI-ID - The MSTI ID of the VLAN containing the port. (The MSTI ID for a regional boundary port is always 0, even if the VLAN containing the port has been associated with a MSTI other than CIST.)
- Role - Indicates the MSTP role of the port. Possible roles are: root, alternate, backup, and designated.
- Internal Port Cost - The port cost when the port is connected to a bridge in the same region.
- Version - Indicates whether the port is operating in MSTP mode or STP-compatible mode.

Section V

Virtual LANs

The chapters in this section explain virtual LANs (VLANs). The chapters include:

Chapter 24

Port-based and Tagged Virtual LANs

This chapter contains background information on port-based and tagged virtual LANs (VLANs). It also contains the procedures for creating, modifying, and deleting VLANs from a local or Telnet management session.

This chapter contains the following sections:

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 management software and so be able to 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 could create separate VLANs for the different departments 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 will decrease.

  VLANs improve network performance because VLAN traffic stays within 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 not destined for them. It also frees up bandwidth within all the logical workgroups.

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

- Increased security

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

- Simplified network management

  VLANs can also 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 wiring at the switches.

  But with VLANS, you can change the LAN segment assignment of an end node connected to the switch through the switch's AT-S62

management software. VLAN memberships can be changed any time through the management software without moving the workstations physically, or having to change group memberships by moving cables from one switch port to another.

Additionally, 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-8500 Series switch supports the following types of VLANs you can create yourself:

- Port-based VLANs
- Tagged VLANs

These VLANs are described in the following sections.

Port-based VLAN Overview

As explained in the “VLAN Overview” on page 546, a VLAN consists of a group of ports on one or more Ethernet switches 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 a Fast Ethernet Switch that form a logical Ethernet segment. Each port of a port-based VLAN can belong to only one VLAN at a time.

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. A port-based VLAN can also span switches and consist of ports from multiple Ethernet switches.

Note
The AT-8500 Series switch is preconfigured with one port-based VLAN. All ports on the switch are members of this VLAN, called the Default_VLAN.

The parts that make up a port-based VLAN are:

- VLAN name
- VLAN Identifier
- Untagged ports
- Port VLAN Identifier

VLAN Name

To create a port-based VLAN, you must give it a name. The name should reflect the function of the network devices that are be members of the VLAN. Examples include Sales, Production, and Engineering.

VLAN Identifier

Each VLAN in a network must have a unique number assigned to it. This number is called the VLAN identifier (VID). This number uniquely identifies a VLAN in the switch and the network.

If a VLAN consists only of ports located on one physical switch in your network, you assign it a VID different from all other VLANs in your network.

If a VLAN spans multiple switches, then the VID for the VLAN on the different switches should be the same. The switches are then able to recognize and forward frames belonging to the same VLAN even though the VLAN spans multiple switches.

For example, if you had a port-based VLAN titled Marketing that spanned three AT-8500 Series switches, you would assign the Marketing VLAN on each switch the same VID.

You can assign this number manually or allow the management software to do it automatically. If you allow the management software to do it automatically, it will select the next available VID. This is acceptable when you are creating a new, unique VLAN.

If you are creating a VLAN on a switch that will be part of a larger VLAN that spans several switch, then you will need to assign the number yourself so that the VLAN has the same VID on all switches.

Untagged Ports

You need to specify which ports on the switch are to be members of a port-based VLAN. Ports in a port-based VLAN are referred to as *untagged ports* and the frames received on the ports as *untagged frames*. The names derive from the fact that the frames received on a port will not contain any information that indicates VLAN membership, and that VLAN membership will be determined solely by the port's PVID. (There is another type of VLAN where VLAN membership is determined by information within the frames themselves, rather than by a port's PVID. This type of VLAN is explained in "Tagged VLAN Overview" on page 555.)

A port on a switch can be an untagged member of only one port-based VLAN at a time. An untagged port cannot be assigned to two port-based VLANs simultaneously.

Port VLAN Identifier

Each port in a port-based VLAN must have a port VLAN identifier (PVID). The switch associates a frame to a port-based VLAN by the PVID assigned to the port on which the frame is received, and forwards the frame only to those ports with the same PVID. Consequently, all ports of a port-based VLAN must have the same PVID. Additionally, the PVID of the ports in a VLAN must match the VLAN's VID.

For example, if you were creating a port-based VLAN on a switch and you had assigned the VLAN the VID 5, the PVID for each port in the VLAN would need to be assigned the value 5.

Some switches and switch management programs require that you assign the PVID value for each port manually. However, the AT-S62 management software performs this task automatically. The software

automatically assigns a PVID to a port, making it identical to the VID of the VLAN to which the port is a member, when you assign the port as an untagged member to a VLAN.

General Rules for Creating a Port-based VLAN

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

- Each port-based VLAN must be assigned a unique VID. If a particular VLAN spans multiples switches, each part of the VLAN on the different switches should be assigned the same VID.
- A port can be an untagged member of only one port-based VLAN at a time.
- Each port must be assigned a PVID. This value must be the same for all ports in a port-based VLAN and it must match the VLAN's VID. This value is automatically assigned by the AT-S62 management software.
- A port-based VLAN that spans multiple switches requires a port on each switch where the VLAN is located to function as an interconnection between the switches where the various parts of the VLAN reside.
- If there are end nodes in different VLANs that need to communicate with each other, a router or Layer 3 switch is required to interconnect the VLANs.
- A port set to the 802.1x authenticator or supplicant role must be changed to the 802.1x none role before you can change its untagged VLAN assignment. After the VLAN assignment is made, the port's role can be changed back again to authenticator or supplicant, if desired.
- Deleting an untagged port from the Default VLAN without assigning it to another VLAN results in the port being an untagged member of no VLAN.

Drawbacks of Port-based VLANs

There are several drawbacks to port-based VLANs:

- It is not easy to share network resources, such as servers and printers, across multiple VLANs. A router or Layer 3 switch must be added to the network to provide a means for interconnecting the port-based VLANs. The introduction of a router into your network could create security issues from unauthorized access to your network.
- A VLAN that spans several switches requires a port on each switch for the interconnection of the various parts of the VLAN. For example, a VLAN that spans three switches would require one port on each switch to interconnect the various sections of the VLAN. In network configurations where there are many individual VLANs that span switches, many ports could end up being used ineffectively just to interconnect the various VLANs.

Port-based Example 1

Figure 185 illustrates an example of one AT-8524M Fast Ethernet Switch with three port-based VLANs. (For purposes of the following examples, the Default_VLAN is not shown.)

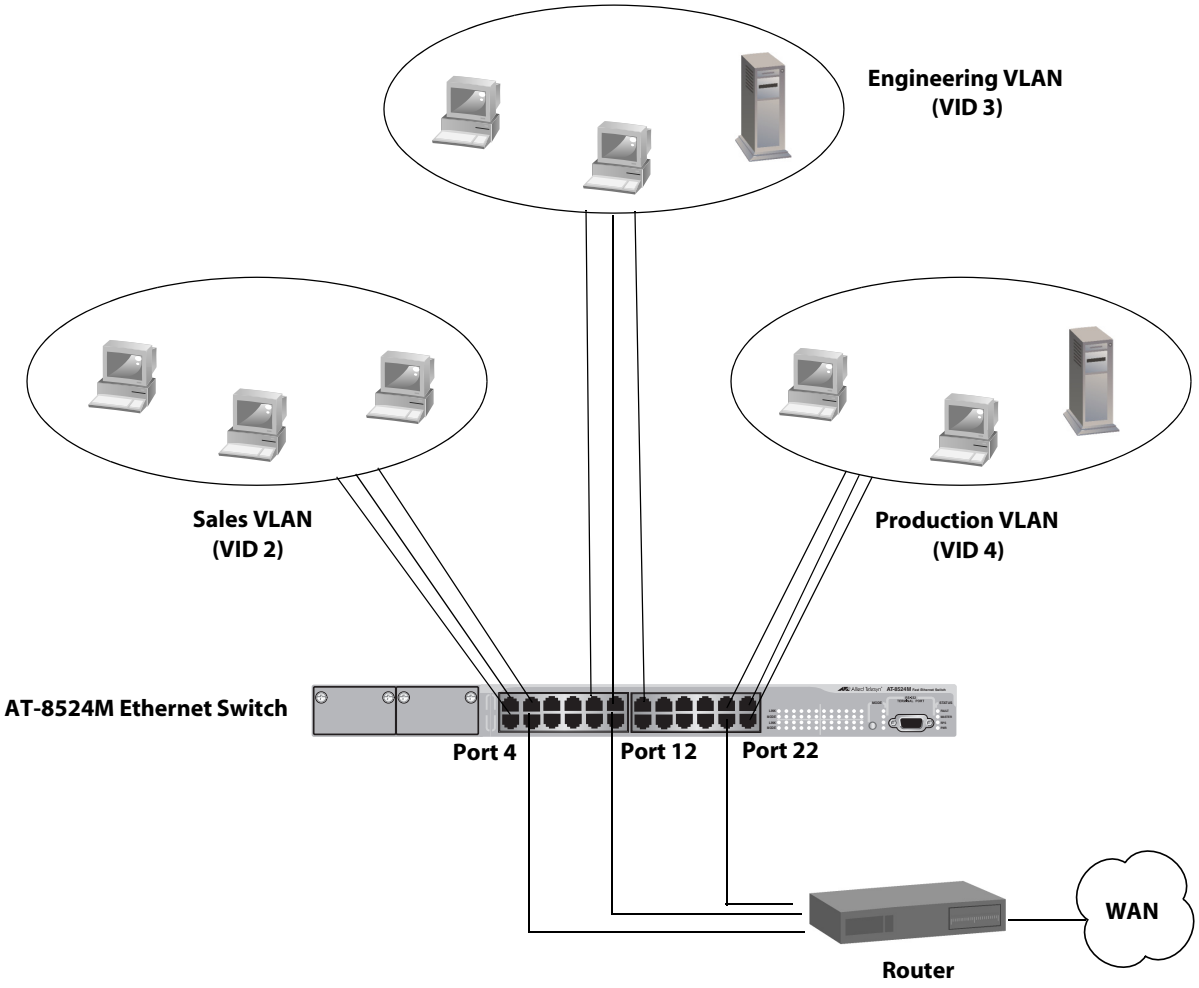

Figure 185. Port-based VLAN - Example 1

The table below lists the port assignments for the Sales, Engineering, and Production VLANs on the switch.

	Sales VLAN (VID 2)	Engineering VLAN (VID 3)	Production VLAN (VID 4)
AT-8524M Switch	Ports 1 - 4 (PVID 2)	Ports 9, 11 - 13 (PVID 3)	Ports 21 - 24 (PVID 4)

Each VLAN has been assigned a unique VID. You assign this number when you create a VLAN.

The ports have been assigned PVID values. The management software automatically assigns the PVIDs when you create the VLAN. The PVID of a port is the same as the VID to which the port is an untagged member.

In the example, each VLAN has one port connected to the router. The router interconnects the various VLANs and functions as a gateway to the WAN.

Port-based Example 2

Figure 186 illustrates more port-based VLANs. In this example, two VLANs, Sales and Engineering, span two Ethernet switches.

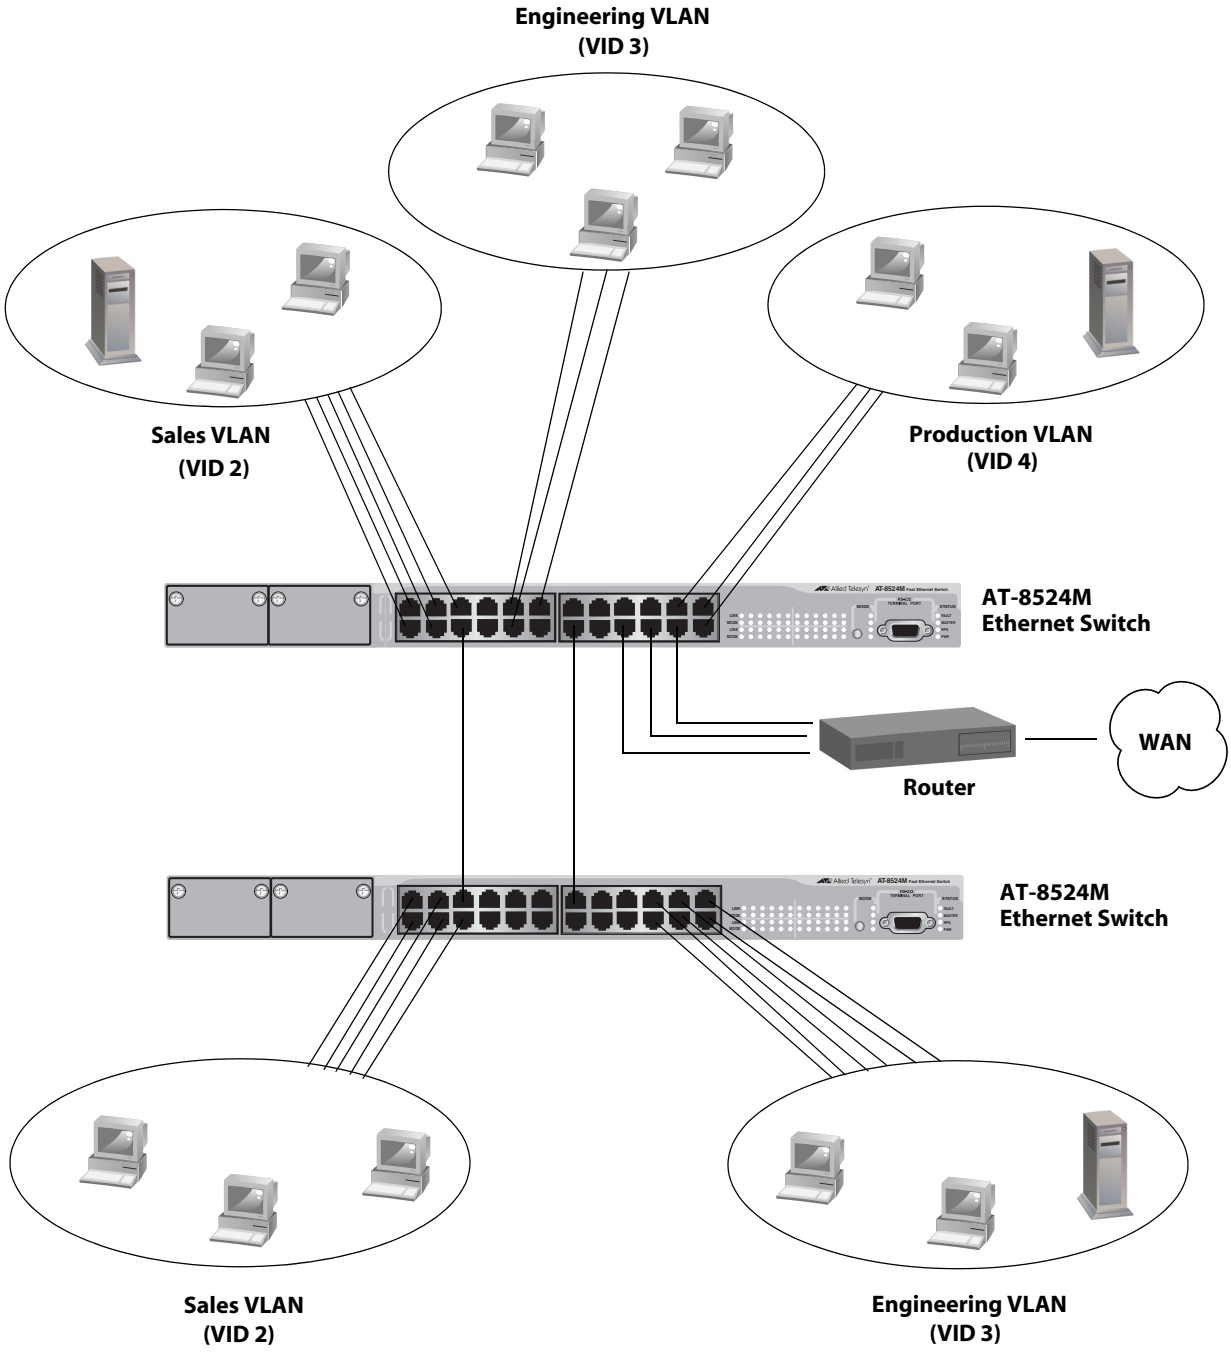

Figure 186. Port-based VLAN - Example 2

The table below lists the port assignments for the Sales, Engineering, and Production VLANs on the switches:

	Sales VLAN (VID 2)	**Engineering VLAN (VID 3)**	**Production VLAN (VID 4)**
AT-8524M Switch (top)	Ports 1 - 6, 18 (PVID 2)	Ports 9 - 11, 14, 20 (PVID 3)	Ports 21 - 24 (PVID 4)
AT-8524M Switch (bottom)	Ports 1 - 6 (PVID 2)	Ports 13, 19-24 (PVID 3)	none

- Sales VLAN - This VLAN spans both switches. It has a VID value of 2 and consists of seven untagged ports on the top switch and six untagged ports on the bottom switch.

  The two parts of the VLAN are connected by a direct link from port 6 on the top switch to port 5 on the bottom switch. This direct link allows the two parts of the Sales VLAN to function as one logical LAN segment.

  Port 18 on the top switch connects to the router. This port allows the Sales VLAN to exchanged Ethernet frames with the other VLANs and to access the WAN.

- Engineering VLAN - The workstations of this VLAN are connected to ports 9 to 11 on the top switch and ports 19 to 24 on the bottom switch.

  Since this VLAN spans multiple switches, it needs a direct connection between its various parts to provide a communications path. This is provided in the example with a direct connection from port 14 on the top switch to port 13 on the bottom switch.

  This VLAN uses port 20 on the top switch as a connection to the router and the WAN.

- Production VLAN - This is the final VLAN in the example. It has the VLAN of 4 and its ports have been assigned the PVID also of 4.

  The nodes of this VLAN are connected only to the top switch. So this VLAN does not require a direct connection to the bottom switch. However, it uses port 22 as a connection to the router.

Tagged VLAN Overview

The second type of user-configured VLAN is the *tagged VLAN*. VLAN membership in a tagged VLAN is determined by information within the frames that are received on a port. This differs from a port-based VLAN, where the PVIDs assigned to the ports determine VLAN membership.

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 VID of the VLAN to which the frame belongs (IEEE 802.3ac standard). As explained earlier in this chapter in "VLAN Identifier" on page 548, 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 that are members of the VLAN whose VID matches the tag in the frame.

A port receiving or transmitting 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.

The benefit of a tagged VLAN is that the tagged ports can belong to more than one VLAN at one time. This can greatly simplify the task of adding shared devices to the network. For example, a server can be configured to accept and return packets from many different VLANs simultaneously.

Tagged VLANs are also useful where multiple VLANs span across switches. You can use one port per switch to connect all VLANs on the switch to another switch.

The IEEE 802.1Q standard deals with how this tagging information is used to forward the traffic throughout the switch. The handling of frames tagged with VIDs coming into a port is straightforward. If the incoming frame's VID tag matches one of the VIDs of a VLAN of which the port is a tagged member, the frame is accepted and forwarded to the appropriate ports. If the frame's VID does not match any of the VLANs that the port is a member of, the frame is discarded.

The parts of a tagged VLAN are much the same as those for a port-based VLAN. They are:

- ❒ VLAN Name
- ❒ VLAN Identifier
- ❒ Tagged and Untagged Ports
- ❒ Port VLAN Identifier

Note
For an explanation of VLAN name and VLAN identifier, refer back to VLAN Name and “VLAN Identifier” on page 548.

Tagged and Untagged Ports

You need to specify which ports will be members of the VLAN. In the case of a tagged VLAN, this will usually be a combination of both untagged ports and tagged ports. You specify which ports are tagged and which untagged when you create the VLAN.

An untagged port, whether a member of a port-based VLAN or a tagged VLAN, can be in only one VLAN at a time. However, a tagged port can be a member of more than one VLAN. A port can also be an untagged member of one VLAN and a tagged member of different VLANs simultaneously.

Port VLAN Identifier

As explained earlier in the discussion on port-based VLANs, the management software automatically adjusts the PVID of a port when a port is made an untagged member of a VLAN. It is adjusted to match the VLAN’s VID.

Because a tagged port determines VLAN membership by examining the tagged header within the frames that it receives, you might conclude that there is no need for a PVID. However, the PVID is used if a tagged port receives an untagged frame—a frame without any tagged information. The port forwards the frame based on the port’s PVID. This is only in cases where an untagged frame arrives on a tagged port. Otherwise, the PVID of a port is ignored on a tagged port.

General Rules for Creating a Tagged VLAN

Below is a summary of the rules to observe when creating 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.
- An untagged port can be an untagged member of only one VLAN at a time.
- The AT-8500 Series switch can support up to 255 tagged VLANS.

Tagged VLAN Example

Figure 187 illustrates how tagged ports can be used to interconnect IEEE 802.1Q-based products.

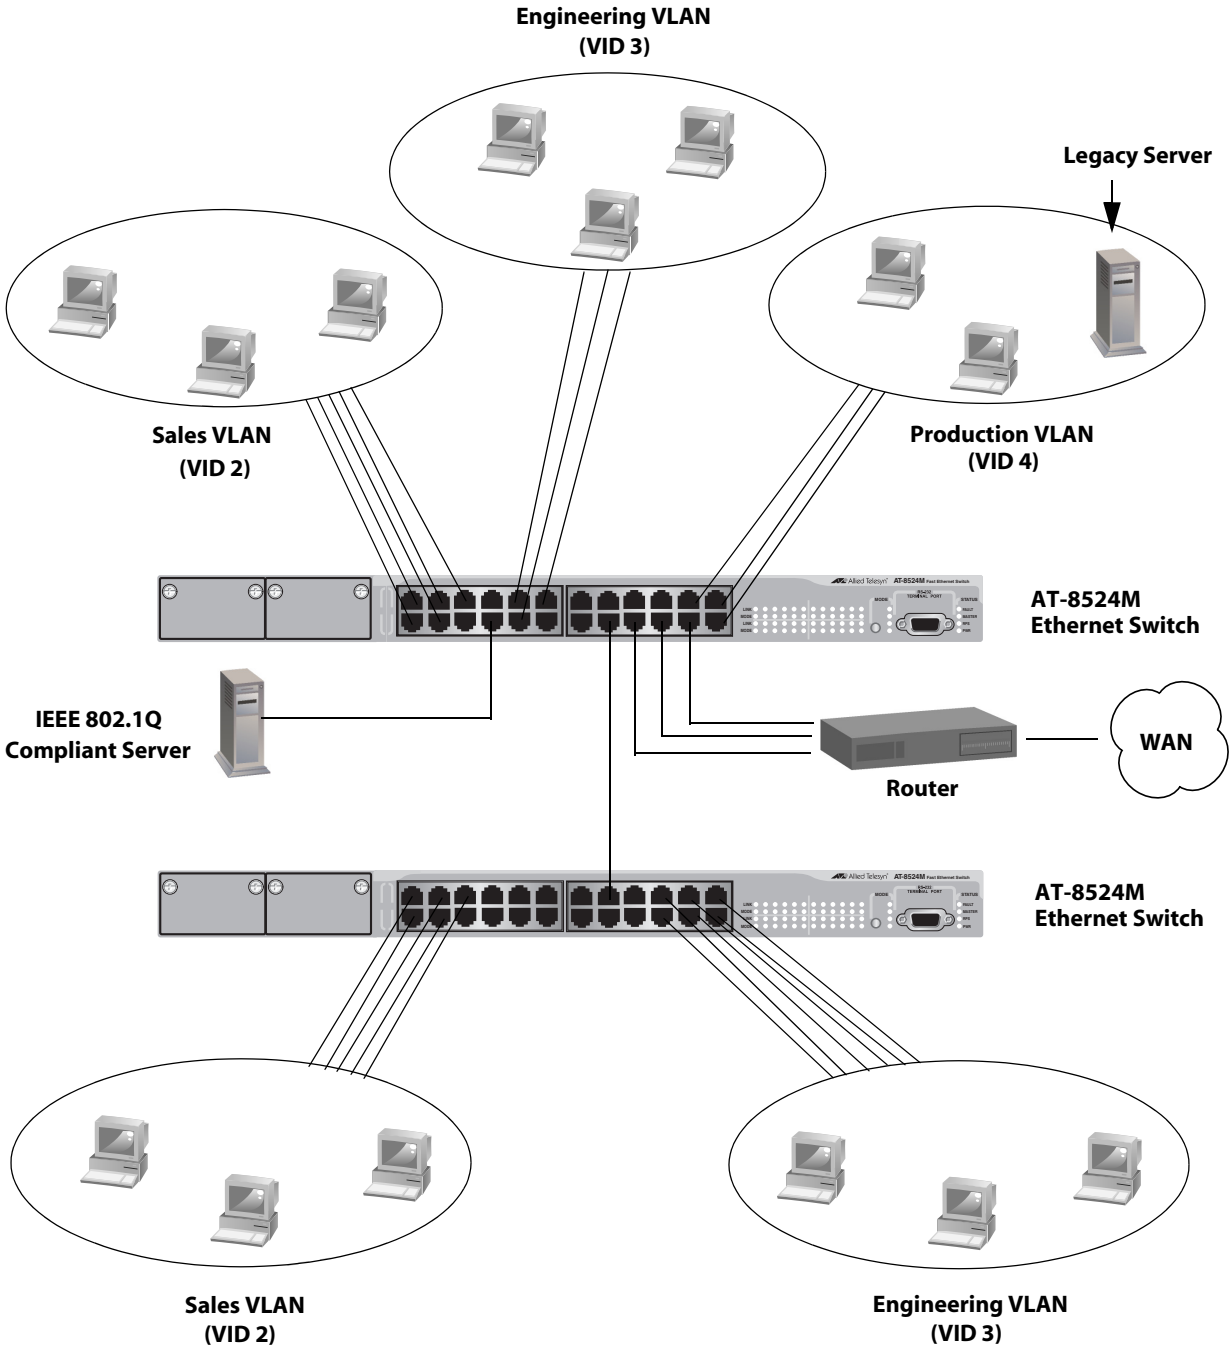

Figure 187. Example of a Tagged VLAN

The port assignments for the VLANs are as follows:

	Sales VLAN (VID 2)		**Engineering VLAN (VID 3)**		**Production VLAN (VID 4)**	
	Untagged Ports	Tagged Ports	Untagged Ports	Tagged Ports	Untagged Ports	Tagged Ports
AT-8524 M Switch (top)	1 to 5, 18 (PVID 2)	8, 16	9 to 11, 20 (PVID 3)	8, 16	21 to 24 (PVID 4)	8
AT-8524 M Switch (bottom)	1 to 5 (PVID 2)	15	19 to 24 (PVID 3)	15	none	none

This example is nearly identical to the “Port-based Example 2” on page 553. Tagged ports have been added to simplify network implementation and management.

One of the tagged ports is port 8 on the top switch. This port has been made a tagged member of the three VLANs. It is connected to an IEEE 802.1Q-compliant server, meaning the server can handle frames from multiple VLANs. Now all three VLANs can access the server without having to go through a router or other interconnection device.

It is important to note that even though the server is accepting frames from and transmitting frames to more than one VLAN, data separation and security remain.

Two other tagged ports are used to simplify network design in the example. They are port 16 on the upper switch and port 15 on the lower switch. These ports have been made tagged members of the Sales and Engineering VLANs so that they can carry traffic from both VLANs, simultaneously. These ports provide a common connection that enables different parts of the same VLAN to communicate with each other while maintaining data separation between VLANs.

In comparison, the Sales and Engineering VLANs in the “Port-based Example 2” on page 553 each had to have its own individual network link between the switches to connect the different parts of the VLANs. But with tagged ports, you can use one data link to carry data traffic from several VLANs, while still maintaining data separation and security. The tagged frames, when received by the switch, are delivered only to those ports that belong to the VLAN from which the tagged frame originated.

Creating a Port-based or Tagged VLAN

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

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188.

```
        Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch

User: Manager                                11:20:02 02-Jan-2006

                          VLAN Configuration


1 - Ingress Filtering Status ........ Enabled
2 - VLANs Mode ...................... User Configured VLANs
3 - Management VLAN ................. 1 (Default_VLAN)
4 - Configure VLANs
5 - Show VLANs
6 - Show PVIDs
7 - Configure GARP-GVRP

R - Return to Previous Menu


Enter your selection?
```

Figure 188. VLAN Configuration Menu

2. From the VLAN Configuration menu, type **4** to select Configure VLANs.

Note
If option "4 - Configure VLANs" is not displayed in the menu, the switch is running in a multiple VLAN mode. To change a switch's VLAN mode, refer to "Selecting a VLAN Mode" on page 612.

The Configure VLANs menu is shown in Figure 189.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch

User: Manager                          11:20:02 02-Jan-2006

                      Configure VLANs


1 - Create VLAN
2 - Modify VLAN
3 - Delete VLAN
4 - Reset to Default VLAN

R - Return to Previous Menu


Enter your selection?
```

Figure 189. Configure VLANs Menu

3. From the Configure VLANs menu, type **1** to select Create VLAN.

   The Create VLAN menu is shown in Figure 190.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch

User: Manager                          11:20:02 02-Jan-2006

                         Create VLAN

1 - VLAN Name .............
2 - VLAN ID (VID) ........ 2
3 - Tagged Ports .........
4 - Untagged Ports .......
5 - Protected Ports ...... No

C - Create VLAN
R - Return to Previous Menu

Enter your selection?
```

Figure 190. Create VLAN Menu

4. Type **1** to select VLAN Name and enter a name for the new VLAN.

   The name can be from one to twenty alphanumeric characters in length. The name should reflect the function of the nodes that will be a part of the VLAN (for example, Sales or Accounting). The name cannot contain spaces or special characters, such as asterisks (*) or exclamation points (!).

   If the VLAN will be unique in your network, then the name should be unique as well. If the VLAN will be part of a larger VLAN that spans

multiple switches, then the name for the VLAN should be the same on each switch where nodes of the VLAN are connected.

Note
A VLAN must be assigned a name.

5. Type **2** to select VLAN ID (VID) and enter a VID value for the new VLAN. The permitted range of the VID value is 1 to 4094.

Note
A VLAN must have a VID.

The management software will use the next available VID number on the switch as the default value. If this VLAN will be unique in your network, then its VID should also be unique. If this VLAN will be part of a larger VLAN that spans multiple switches, than the VID value for the VLAN should be the same on each switch. For example, if you are creating a VLAN called Sales that will span three switches, the Sales VLAN on each switch should be assigned the same VID value.

The switch is only aware of the VIDs of the VLANs that exist on the device, and not those that might already be in use in the network. For example, if you add a new AT-8500 Series switch to a network that already has VLANs using VIDs 2 through 24, the AT-S62 software will still use VID 2 as the default value for the first VLAN you create on the new switch, even though that VID number is already being used by another VLAN on the network. To prevent inadvertently using the same VID for two different VLANs, you should keep a list of all your network VLANs and their VID values.

6. If the VLAN will contain tagged ports, type **3** to select Tagged Ports and specify the ports. If this VLAN will not contain any tagged ports, leave this field empty.

   You can specify the ports individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9).

7. Type **4** to select Untagged Ports and specify the ports on the switch to function as untagged ports in the VLAN. If this VLAN will not contain any untagged ports, leave this field empty.

   You can specify the ports individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9).

   When assigning ports to a VLAN, note the following:

   - An untagged port is automatically removed from its current untagged VLAN assignment when assigned to a new VLAN.

- A port set to the 802.1x authenticator or supplicant role must be changed to the 802.1x none role before you can change its untagged VLAN assignment. After the VLAN assignment is made, you can return the port's role to authenticator or supplicant, if desired.

Note
Option 5, Protected Ports, in the Create VLAN menu is not used to create port-based and tagged VLANs. It must be left in the "No" default setting. This option is used to create protected ports VLANs, as explained in Chapter 27, "Protected Ports VLANs" on page 615.

8. Type **C** to select Create VLAN.

   The following message is displayed:

   ```
   SUCCESS - Press any key to continue.
   ```

   The AT-S62 software creates the new VLAN. The new VLAN is now ready for network use.

9. Press any key.

   The VLAN Configuration menu in Figure 188 on page 559 is redisplayed.

10. To verify that the VLAN was created correctly, type **5** to select Show VLANs.

11. Check to see that the VLAN contains the appropriate ports.

12. Repeat this procedure to create additional VLANs.

13. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Note
When you create a new VLAN, ports designated as untagged ports of the new VLAN are automatically removed from their current untagged VLAN assignment. For example, if you are creating a new VLAN on a switch that contains only the Default_VLAN, the ports that you specify as untagged ports of the new VLAN are automatically removed from the Default_VLAN.

Tagged ports are not removed from any current VLAN assignments because tagged ports can belong to more than one VLAN at a time.

Example of Creating a Port-based VLAN

The following procedure creates the Sales VLAN illustrated in "Port-based Example 1" on page 551. This VLAN will be assigned a VID of 2 and will consist of four untagged ports, Ports 1 to 4. The VLAN will not contain any tagged ports.

To create the Sales VLAN, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **4** to select Configure VLANs.

   The Configure VLANs menu is shown in Figure 189 on page 560.

3. From the Configure VLANs menu, type **1** to select Create VLAN.

   The Create VLAN menu is shown in Figure 190 on page 560.

4. Type **1** to select VLAN Name and enter "Sales".

5. Type **2** to select VLAN ID (VID) and enter "2". This is the VID value for the new VLAN.

6. Type **4** to select Untagged Ports and enter "1-4". These are the untagged ports of the VLAN. Press Return.

   Note
   Option 5, Protected Ports, in the Create VLAN menu is not used to create port-based or tagged VLANs. It must be left in the "No" default setting. This option is explained in Chapter 27, "Protected Ports VLANs" on page 615.

7. Type **C** to select Create VLAN.

8. After the switch displays the prompt notifying you that it created the VLAN, press any key.

   The new Sales VLAN has now been created.

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Example of Creating a Tagged VLAN

The following procedure creates the Engineering VLAN in the top switch illustrated in “Tagged VLAN Example” on page 557. This VLAN will be assigned a VID of 3. It will consist of four untagged ports, Ports 9 to 11 and 20, and two tagged ports, Ports 8 and 16.

To create the example Engineering VLAN, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **4** to select Configure VLANs.

   The Configure VLANs menu is shown in Figure 189 on page 560.

3. From the Configure VLANs menu, type **1** to select Create VLAN.

   The Create VLAN menu is shown in Figure 190 on page 560.

4. Type **1** to select VLAN Name and enter “Engineering”.

5. Type **2** to select VLAN ID (VID) and enter “3”. This is the VID value for the new VLAN.

6. Type **3** to select Tagged Ports and enter “8,16”. These are the tagged ports of the VLAN on the switch.

7. Type **4** to select Untagged Ports and enter “9-11, 20”. These are the untagged ports of the VLAN.

   Note
   Option 5, Protected Ports, in the Create VLAN menu is not used to create port-based or tagged VLANs. It must be left in the “No” default setting. This option is explained in Chapter 27, “Protected Ports VLANs” on page 615.

8. Type **C** to select Create VLAN.

9. After the switch displays the prompt notifying you that it created the VLAN, press any key.

   The new Engineering VLAN has now been created.

10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a VLAN

You can use this procedure to add or remove ports from a port-based or tagged VLAN. You can also use this procedure to change a VLAN's name.

Note
To modify a VLAN, you need to know its VID. To view VLAN VIDs, refer to "Displaying VLANs" on page 569.

To modify a VLAN, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **4** to select Configure VLANs.

   The Configure VLANs menu is shown in Figure 189 on page 560.

   Note
   If option "4 - Configure VLANs" is not displayed in the menu, the switch is running a multiple VLAN mode. To change a switch's VLAN mode, refer to "Selecting a VLAN Mode" on page 612.

3. From the Configure VLANs menu, type **2** to select Modify VLAN.

   The Modify VLAN menu is shown in Figure 191.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch

User: Manager                          11:20:02 02-Jan-2006

                        Modify VLAN

1 - VLAN ID (VID) ........
2 - Change GARP VLAN

R - Return to Previous Menu

Enter your selection?
```

Figure 191. Modify VLAN Menu

   Option 2 - Change GARP VLAN is described in "Converting a Dynamic GVRP VLAN" on page 596.

4. Type **1** to select VLAN ID (VID).

The following prompt is displayed:

```
Enter new value -> [1 to 4096] ->
```

5. Enter the VID of the VLAN you want to modify.

   The Modify VLAN menu expands to contain all relevant information about the VLAN, as shown in Figure 192.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                          11:20:02 02-Jan-2006

                        Modify VLAN

1 - VLAN Name ............... Sales
2 - VLAN ID (VID) ........... 3
3 - Tagged Ports ............ 7,9
4 - Untagged Ports .......... 20-24
5 - Protected Ports ......... No

M - Modify VLAN
R - Return to Previous Menu

Enter your selection?
```

Figure 192. Expanded Modify VLAN Menu

6. Change the VLAN's information as desired.

   The selections in the menu are described below:

   1 - VLAN Name
   Use this selection to change the name of a VLAN. The name can be from one to twenty characters in length. The name should reflect the function of the nodes that will be a part of the VLAN (for example, Sales or Accounting). The name cannot contain spaces or special characters, such as asterisks (*) or exclamation points (!).

   When changing a VLAN's name, observe the following guidelines:

   - A VLAN's new name cannot be the same as the name of another VLAN on the same switch. For example, if the switch already contains a VLAN called Sales, you cannot change an existing VLAN's name to Sales.
   - You cannot change the name of the Default_VLAN.

   Note
   A VLAN must have a name.

   2 - VLAN ID (VID)
   This is the VLAN's VID value. You cannot change this value.

3 - Tagged Ports
Use this selection to add or remove tagged ports from the VLAN. You can specify the ports individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9).

When adding or removing tagged ports, observe the following guidelines:

- The new list of tagged ports will replace the existing tagged ports.
- If the VLAN contains tagged ports and you want to remove them all, enter 0 (zero) for this value.

4 - Untagged Ports
Use this selection to add or remove untagged ports from the VLAN. You can specify the ports individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9).

When adding or removing untagged ports, observe the following guidelines:

- The new list of untagged ports will replace the existing list.
- If you want to remove all untagged ports from the VLAN, enter 0 (zero) for this value.
- You cannot change the name of the Default_VLAN, nor can you directly remove untagged ports from the Default_VLAN. Instead, you must assign the port as an untagged port to another VLAN.
- A port set to the 802.1x authenticator or supplicant role must be changed to the 802.1x none role before you can change its untagged VLAN assignment. After the VLAN assignment is made, you can return the port's role to authenticator or supplicant, if desired.
- An untagged port removed from a VLAN is automatically returned to the Default_VLAN as an untagged port.

Note
Option 5, Protected Ports, in the Modify VLAN menu is not used to modify port-based and tagged VLANs. It should be left in the "No" default setting. This option is used to modify protected ports VLANs, as explained in Chapter 27, "Protected Ports VLANs" on page 615.

7. After making the desired changes, type **M** to select Modify VLAN.

   The following message is displayed:

```
SUCCESS
Please make sure to manually update any static
multicast MAC address(es) entries for this VLAN.
Press any key to continue...
```

   The VLAN has been modified and is now ready for network operations.

Any untagged ports removed from a VLAN are automatically returned to the Default_VLAN as untagged ports.

If you added or removed from the VLAN a port with one or more static MAC addresses assigned to it, you must update the static addresses by deleting their entries from the MAC address table and reentering them again using the VID of the VLAN to which the port has been moved to. For information on how to add static MAC addresses, refer to “Adding Static Unicast and Multicast MAC Addresses” on page 130. For instructions on how to delete addresses, refer to “Deleting Unicast and Multicast MAC Addresses” on page 132.

8. Press any key.

   The Modify VLAN menu in Figure 191 on page 565 is displayed again.

9. Repeat this procedure starting with Step 4 to modify other VLANs, or return to the Main Menu and type **S** to select Save Configuration Changes.

Displaying VLANs

To view the name, VID number, and member ports of all the VLANs on a switch, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **5** to select Show VLANs.

   An example of the Show VLANs menu is shown in Figure 193.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch

User: Manager                                        11:20:02 02-Jan-2006

                                Show VLANs

VID    VLAN Name      VLAN Type    Protocol   Member Ports
-------------------------------------------------------------------

1      Default_VLAN   Port Based              Untagged
                                                Configured: 5-6,11-17,25-26
                                                Actual:     5-6,11-17,25-26
                                              Tagged:
4      Sales          Port Based              Untagged
                                                Configured: 1-4,7-10
                                                Actual:     1-4,7-10
                                              Tagged: 17
7      Production     Port Based              Untagged
                                                Configured: 18-24
                                                Actual:     18-24
                                              Tagged:17

U - Update Display
D - Detail Information Display
R - Return to Previous Menu

Enter your selection?
```

Figure 193. Show VLANs Menu

Note
Selection D, Detail Information Display, only applies to MAC address-based VLANs.

The menu contains the following columns of information:

VID - The VLAN ID.

VLAN Name - The name of the VLAN.

VLAN Type
The VLAN type. The possible settings are:

Port Based - The VLAN is a port-based or tagged VLAN.

GARP - The VLAN was automatically created by GARP.

Protected - The VLAN is a protected ports VLAN.

Protocol - If this column is blank, the VLAN is a port-based, tagged, or protected ports VLAN. If it contains “GARP,” the VLAN or the port is a dynamic GVRP VLAN or a dynamic GVRP port of a static VLAN.

Member Port(s)
The untagged and tagged ports of a VLAN. (These fields will be blank for a MAC address-based VLAN.) The untagged ports of a VLAN are listed as follows.

- Configured: The untagged ports assigned to the VLAN when the VLAN was created or modified.
- Actual: The current untagged ports of the VLAN. If you are not using 802.1x Port-based Network Access Control, both the Configured and Actual untagged ports of a VLAN will always be the same.

  If you are using 802.1x and you assigned a Guest VLAN to an authenticator port or you associated an 802.1x supplicant to a VLAN on the authentication server, it is possible for a port to be in different VLAN than the virtual LAN where it was originally assigned as an untagged port. In these situations, the Configured and Actual port lists can differ, with the Actual list detailing the ports that are currently functioning as untagged ports of the VLAN.

  For example, if a particular port is listed as a Configured member of a VLAN, but not as an Actual member, that would mean either the port is currently a part of a Guest VLAN or the supplicant who logged on the port was associated with a VLAN assignment on the authentication server.

Deleting a VLAN

This procedure deletes port-based and tagged VLANs from the switch. All untagged ports in a deleted VLAN are returned to the Default_VLAN.

> **Note**
> To delete a VLAN, you need to know its VID. To view VLAN VIDs, refer to "Displaying VLANs" on page 569.

To delete a VLAN, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **4** to select Configure VLANs.

   The Configure VLANs menu is shown in Figure 189 on page 560.

   > **Note**
   > If option "4 - Configure VLANs" is not displayed in the menu, the switch is running a multiple VLAN mode. To change a switch's VLAN mode, refer to "Selecting a VLAN Mode" on page 612.

3. From the Configure VLANs menu, type **3** to select Delete VLAN.

   The Delete VLAN menu is shown in Figure 194.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                        11:20:02 02-Jan-2006

                      Delete VLAN

1 - VLAN ID (VID) ........

R - Return to Previous Menu

Enter your selection?
```

Figure 194. Delete VLAN Menu

4. Type **1** to select VLAN ID (VID).

   The following prompt is displayed:

   ```
   Enter new value -> [2 to 4094] ->
   ```

5. Enter the VID of the VLAN you want to delete. You can specify only one VID at a time.

   Note
   You cannot delete the Default_VLAN, which has a VID of 1.

   The Delete VLAN menu expands to contain all relevant information about the VLAN, as shown in Figure 195.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch

User: Manager                         11:20:02 02-Jan-2006

                         Delete VLAN

1 - VLAN Name ............... Sales
2 - VLAN ID (VID) ........... 3
3 - VLAN Type ............... Port Based
4 - Tagged Ports ............ 7,9
5 - Untagged Ports .......... 20-24
6 - Protected Ports ......... No

D - Delete VLAN
R - Return to Previous Menu

Enter your selection?
```

Figure 195. Expanded Delete VLAN Menu

6. Type **D** to delete the VLAN or **R** to cancel the procedure.

   If you select to delete the VLAN, the following confirmation prompt is displayed:

```
Are you sure you want to delete this VLAN [Yes/No] ->
```

7. Type **Y** to delete the VLAN or **N** to cancel the procedure. Press Return.

   If you select Yes, the VLAN is deleted and the following message is displayed:

```
SUCCESS
Please make sure to manually delete any static
multicast MAC address(es) entries for this VLAN
Press any key to continue ...
```

   All untagged ports in the deleted VLAN are returned to the Default_VLAN as untagged ports.

   Any static addresses assigned to the ports of the VLAN are now obsolete, since the VLAN has been deleted. Those addresses should be deleted from the MAC address table. For instructions on how to delete addresses, refer to “Deleting Unicast and Multicast MAC Addresses” on page 132.

8. Press any key.
9. Repeat this procedure starting with Step 4 to delete other VLANs.
10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting All VLANs

This section contains the procedure for deleting all port-based and tagged VLANs, except the Default_VLAN, on a switch. To delete selected VLANs, perform the procedure “Deleting a VLAN” on page 571.

To delete all VLANs on a switch, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **4** to select Configure VLANs.

   The Configure VLANs menu is shown in Figure 189 on page 560.

   Note
   If option “4 - Configure VLANs” is not displayed in the menu, the switch is running in a multiple VLAN mode. To change a switch’s VLAN mode, refer to “Selecting a VLAN Mode” on page 612.

3. From the Configure VLANs menu, type **4** to select Reset to Default VLAN.

   The following prompt is displayed:

   ```
   This operation deletes ALL user created VLANs!
   Do you want to continue [Yes/No] ->
   ```

4. Type **Y** to delete all VLANs or **N** to cancel the procedure. Press Return.

   If you select Yes, all port-based and tagged VLANs are deleted and the following message is displayed:

   ```
   SUCCESS
   Please make sure to manually update any static
   multicast MAC address(es) entries.
   Press any key to continue...
   ```

   All tagged and untagged ports are returned to the Default_VLAN as untagged ports.

   Any static addresses assigned to the ports of the VLANs are now obsolete, except for the Default_VLAN, since the VLANs have been deleted. Those addresses should be deleted from the MAC address table. For instructions on how to delete addresses, refer to “Deleting Unicast and Multicast MAC Addresses” on page 132.

5. Press any key.

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying PVIDs

The following procedure displays a menu that lists the PVIDs for all the ports on the switch. To display the PVID settings on the switch, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **6** to select Show PVIDs.

   The Show PVIDs menu is shown in Figure 196.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                      11:20:02 02-Jan-2006

                        Show PVIDs

Port   PVID
---------------


01      1
02      1
03      1
04      1
05      1
06      1
07      1


N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 196. Show PVIDs & Priorities Menu

The PVID column displays the current PVID value for each switch port.

Enabling or Disabling Ingress Filtering

There are rules a switch follows when it receives and forwards an Ethernet frame. There are rules for frames as they enter a port (called *ingress rules*) and rules for when a frame is transmitted out a port (called *egress rules*). A switch does not accept and forward a frame unless the frame passes the ingress and egress rules.

There are quite a few ingress and egress rules for Fast Ethernet switches. Fortunately, this discussion need only review the rules as they apply to tagged frames, because ingress filtering does not apply to untagged frames.

First, as a reminder, a tagged frame is an Ethernet frame that contains a tagged header. The header contains the VID of the VLAN to which the frame originated. For further information, refer to "Tagged VLAN Overview" on page 555.

The ingress rules are applied to tagged frames when ingress filtering is activated. The switch examines the tagged header of each tagged frame that enters a port and determines whether the tagged frame and the port that received the frame are members of the same VLAN. If they belong to the same VLAN, the port accepts the frame. If they belong to different VLANs, the port discards the frame.

As an example, assume that a tagged frame with a VID of 4 is received on a port that is a member of a VLAN also with a VID of 4. In this case, the port accepts the frame, because both the frame and the port belong to the same VLAN. If the frame and port belong to different VLANs, the frame is discarded.

How do the egress rules apply when ingress filtering is disabled? First, any tagged frame is accepted on any port on the switch. It does not matter whether the frame and the port belong to the same or different VLANs.

After the tagged frame is received, the switch examines the tagged header and determines if the VID in the header corresponds to any VLANs on the switch. If there is no corresponding VLAN, the switch discards the frame. If there is, the switch transmits the frame out the port to the destination node, assuming that the destination node's MAC address is in the MAC address table, or floods the port to all ports on the VLAN if the MAC address is not in the table.

In addition, each tagged frame contains a priority tag that informs the switch about the importance of the frame. Frames with a high priority are handled ahead of frames with a low priority.

Activating or deactivating ingress filtering has no effect on the switch's handling of priority tags. A switch will always examines a priority tag in a

tagged frame, without regard to the status of ingress filtering.

You can enable or disable ingress filtering on a per switch basis. You cannot set this per port. The default setting for ingress filtering is disabled.

To enable or disable ingress filtering, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **1** to select Ingress Filtering Status.

   The following prompt is displayed:

   ```
   Enter Ingress Filtering Status (E-Enable, D-
   Disable) ->
   ```

3. Type **E** to activate ingress filtering or **D** to disable the feature on the switch.

   A change to the status of ingress filtering is immediately activated on the switch.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Specifying a Management VLAN

The management VLAN is the VLAN on which an AT-8500 Series switch expects to receive management packets. This VLAN is important if you will be managing a switch remotely, using the enhanced stacking feature of the switch, or activating the BOOTP or DHCP client.

Management packets are packets generated by a management workstation when you manage a switch using the Telnet application protocol, SSH, or a web browser. The switch will act upon the management packets only if they are received by a port that is a member of the management VLAN.

The default management VLAN on an AT-8500 Series switch is the Default_VLAN. If you do not create any additional VLANs and link the switches together using untagged ports, then you do not need to specify a new management VLAN in order to remotely manage the devices.

However, if you create additional VLANs on your switches, it may be necessary for you to create a management communications path and then specify that path as the new management VLAN.

Below are several rules to observe when using this feature:

- The management VLAN must exist on each AT-8500 Series switch that you want to manage.
- Using the following procedure, you must specify the management VLAN in the AT-S62 software on each slave and master switch of an enhanced stack.
- The uplink and downlink ports on each switch that are functioning as the tagged or untagged data links between the switches must be either tagged or untagged members of the management VLAN.
- When managing a switch or enhanced stack remotely using Telnet, SSH or a web browser, the remote management workstation must be communicating with the switch through the management VLAN. (This rule does not apply when managing the switch locally through the RJ-45 terminal port.)
- If you activate the BOOTP or DHCP client on a switch, the BOOTP or DHCP server on the network must be communicating with the switch through a port that is a member of the management VLAN.

As an example, assume that you have an enhanced stack of seven AT-8500 Series switches with one master switch. If the uplink and downlink ports between the various switches are members of the Default_VLAN and if the management station is connected to a port of the Default_VLAN, you can manage all the switches because the Default_VLAN is the default management VLAN.

Now assume that you decide to create a VLAN called NMS with a VID of 24 for the sole purpose of remote network management. For this, you need to create the NMS VLAN on each AT-8500 Series switch that you want to manage remotely, being sure to assign each NMS VLAN the VID of 24. Then you need to be sure that the uplink and downlink ports connecting the switches together are either tagged or untagged members of the NMS VLAN. You also need to specify the NMS VLAN as the management VLAN on each switch using the management software. Finally, you must be sure to connect your management station to a port on a switch that is a tagged or untagged member of the management VLAN.

The best approach to changing the management VLAN on the switches of an enhanced stack is to establish a local management session with each switch and adjust it through the local session, rather than through enhanced stacking. Changing a switch's management VLAN through enhanced stacking ends your management session. You will not be able to reestablish the session until you change the management VLAN on the master switch.

Note
You cannot specify a management VLAN when the switch is operating in a multiple VLAN mode.

To specify a management VLAN, do the following:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **3** to select Management VLAN.

   The following prompt is displayed:

   ```
   Enter Management VLAN ID [1 to 4094] ->
   ```

3. Specify the VID of the VLAN that is to function as the management VLAN. This VLAN must already exist on the switch.

   The following prompt is displayed:

   ```
   SUCCESS
   Press any key to continue ...
   ```

4. Press any key.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Chapter 25

GARP VLAN Registration Protocol

This chapter describes the GARP VLAN Registration Protocol (GVRP). It contains the following sections:

Basic Overview of GARP VLAN Registration Protocol (GVRP)

The GARP VLAN Registration Protocol (GVRP) allows network devices to share VLAN information. The main purpose of GVRP is to allow switches to automatically discover some of the VLAN information that would otherwise have to be manually configured in each switch.

This can be helpful in networks where VLANs span more than one switch. Without GVRP, you must manually configure your switches to ensure that the various parts of a VLAN can communicate across the different switches. GVRP, an application of the Generic Attribute Registration Protocol (GARP), can perform this for you automatically.

The AT-S62 management software uses GVRP protocol data units (PDUs) to share VLAN information among GVRP-active devices. The PDUs contain the VID numbers of the VLANs on the switch. A PDU contains the VIDs of all the VLANs on the switch, not just the VID to which the transmitting port is a member.

When a switch receives a GVRP PDU on a port, it examines the PDU to determine the VIDs of the VLANs on the device that sent it. It then does the following:

- If a VLAN does not exist on the switch, it creates the VLAN and adds the port as a tagged member to the VLAN. A VLAN created by GVRP is called a *dynamic GVRP VLAN*.
- If the VLAN already exists on the switch but the port that received the PDU is not a member, the switch adds the port as a tagged member of the VLAN. A port that has been added by GVRP to a static VLAN (that is a user-created VLAN) is called a *dynamic GVRP port*.

You cannot modify a dynamic GVRP VLAN. Once created, only GVRP can modify or delete it. A dynamic GVRP VLAN exists only so long as there are active nodes in the network that belong to the VLAN. If all nodes of a dynamic GVRP VLAN are shutdown and there are no active links, the VLAN is deleted from the switch.

A dynamic GVRP port in a static VLAN remains a member of the VLAN as long as there are active VLAN members. If all members of the VLAN become inactive or there are no active links, GVRP removes the dynamic port from the VLAN, but does not delete the VLAN if the VLAN is a static VLAN (i.e., user created).

Figure 197 provides an example of how GVRP works.

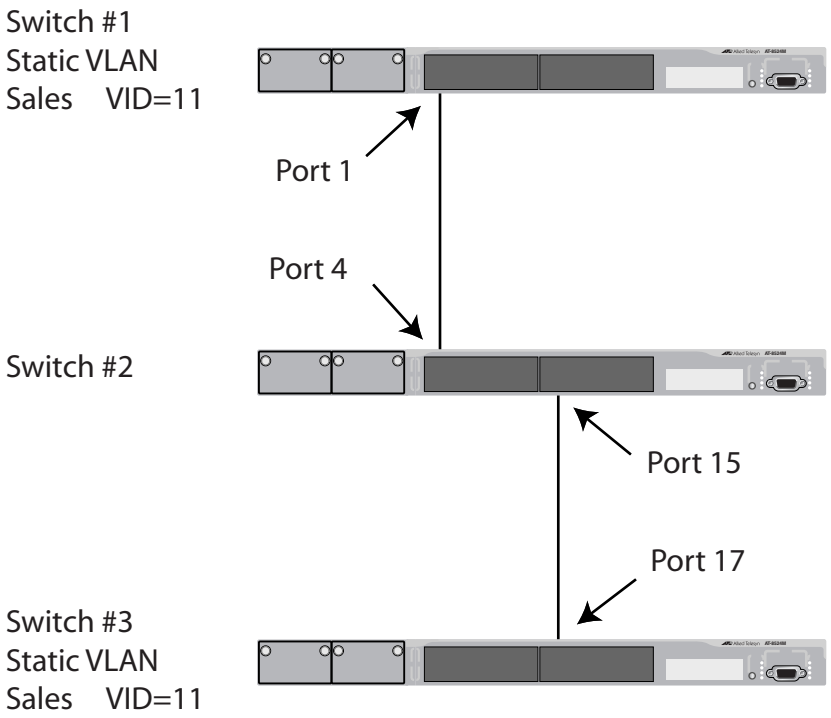

Figure 197. GVRP Example

Switches #1 and #3 contain the Sales VLAN, but Switch #2 does not. Consequently, the end nodes of the two parts of the Sales VLANs are unable to communicate with each other.

Without GVRP, you would need to configure Switch #2 by creating the Sales VLAN on the switch and adding ports 4 and 15 on the switch as members of the VLAN. If you happen to have a large network with a large number of VLANs, manually configuring the devices can be cumbersome and time consuming.

GVRP can make the configurations for you. Here is how GVRP would resolve the problem in the example.

1. Port 1 on Switch #1 sends a PDU to Port 4 on Switch #2, containing the VIDs of all the VLANs on the switch. One of the VIDs in the PDU would be that of the Sales VLAN, VID 11.

2. Switch #2 examines the PDU it receives on Port 4 and notes that it does not have a VLAN with a VID 11. So it creates the VLAN as a dynamic GVRP VLAN and assigns it a VID 11 and the name GVRP_VLAN_11. (The name of a dynamic GVRP VLAN has the prefix "GVRP_VLAN_", followed by the VID number.) The switch then adds Port 4, the port that received the PDU, as a tagged member of the VLAN.

3. Switch #2 sends a PDU out port 15 containing all of the VIDs of the VLANs on the switch, including the new GVRP_VLAN_11 VLAN with its VID of 11. (It should be noted that port 15 is not yet a member of the VLAN. Ports are added to VLANs when they receive, not send a PDU.)

4. Switch #3 receives the PDU on port 17 and, after examining it, notes that one of the VLANs on Switch #2 has the VID 11, which matches the VID of an already existing VLAN on the switch. So it does not create the VLAN since it already exists. It then determines whether the port that received the PDU, in this case port 17, is a member of the VLAN. If it is not a member, it automatically adds the port to the VLAN as an tagged dynamic GVRP port. If the port is already a member of the VLAN, then no change is made.

5. Switch #3 sends a PDU out port 17 to Switch #2.

6. Switch #2 receives the PDU on port 15 and then adds the port as a tagged dynamic GVRP port to the dynamic GVRP_VLAN_11 VLAN.

There is now a communications path for the end nodes of the Sales VLAN on Switches #1 and #3. GVRP created a new dynamic GVRP VLAN, GVRP_VLAN_11, with a VID of 11 on Switch #2 and added ports 4 and 15 to the VLAN as tagged dynamic GVRP ports.

Guidelines

Here are guidelines to observe when using this feature:

- ❒ GVRP is supported with STP and RSTP, or without spanning tree. GVRP is not supported with MSTP.
- ❒ GVRP is supported when the switch is operating in the user-configure VLAN mode, which is the VLAN mode for creating your own tagged and port-based VLANs. GVRP is not supported in either of the Multiple VLAN modes.
- ❒ Both ports that constitute a network link between the switch and the other device must be running GVRP.
- ❒ You cannot modify or delete a dynamic GVRP VLAN.
- ❒ You cannot remove a dynamic GVRP port from a static or dynamic VLAN.
- ❒ GVRP is only aware of those VLANs that have active nodes, or where at least one end node of a VLAN has established a valid link with a switch. GVRP is not aware of a VLAN if there are no active end nodes or if no end nodes have established a link with the switch.
- ❒ Resetting a switch erases all dynamic GVRP VLANs and dynamic GVRP port assignments. The switch relearns the dynamic assignments as it receives PDUs from the other switches.
- ❒ GVRP has three timers that you can set: join timer, leave timer, and leave all timer. The values for these timers must be set the same on all switches running GVRP. Timers with different values on different

switches can result in GVRP incompatibility problems.

- You can convert dynamic GVRP VLANs and dynamic GVRP port assignments to static VLANs and static port assignments. The procedure for this is found in "Modifying a VLAN" on page 565.
- The default port settings on the switch for GVRP is active, meaning that the ports participate in GVRP. Allied Telesyn recommends disabling GVRP on those ports that are connected to GVRP-inactive devices, which are nodes that do not feature GVRP.
- PDUs are transmitted only from those switch ports where GVRP is enabled.

GVRP and Network Security

GVRP should be used with caution because it can expose your network to unauthorized access. A network intruder could access restricted parts of the network by connecting to a switch port running GVRP and transmitting a bogus GVRP PDU containing VIDs of restricted VLANs. GVRP would make the switch port a member of the VLANs and that could give the intruder access to restricted areas of your network.

To protect against this type of network intrusion, you should consider the following:

- Activating GVRP only on those switch ports that are connected to other devices that support GVRP. Do not activate GVRP on ports connected to GVRP-inactive devices, or on ports that are not being used.
- Converting all dynamic GVRP VLANs and dynamic GVRP ports to static assignments, and then turning off GVRP on all switches. This preserves the new VLAN assignments while protecting against network intrusion. The procedure for converting dynamic VLANs to static VLANs is found in "Converting a Dynamic GVRP VLAN" on page 596.

GVRP-inactive Intermediate Switches

The presence of a GVRP-inactive switch between GVRP-active devices may impact the ability of GVRP to automatically configure the VLANs in your switches. You might need to take this into account when implementing GVRP in your network.

One of the problems posed by the introduction of a GVRP-inactive device is that a GVRP-inactive device will probably not forward PDUs, thus preventing the GVRP-active switches from sharing VLAN information. PDUs are management packets, intended for a switch's CPU. In all likelihood, a GVRP-inactive switch will discard the PDUs it receives on its ports because the CPU will not recognize their function.

Another issue is that even if the GVRP-inactive switch does forward GVRP PDUs, it will not automatically create the VLANs. Consequently, even if GVRP-active switches on either side of a GVRP-inactive switch receive the PDUs and create the necessary VLANs, the intermediate switch may block the VLAN traffic, unless you manually modify its VLANs and port assignments.

Technical Overview of Generic Attribute Registration Protocol (GARP)

The purpose of the Generic Attribute Registration Protocol (GARP) is to provide a generic framework whereby devices in a bridged LAN, for example, end stations and switches, can register and de-register *attribute* values, such as VLAN Identifiers, with each other. In doing so, the attributes are propagated to devices in the bridged LAN, and these devices form a "reachability" tree that is a subset of an active topology. For a bridged LAN, the active topology is normally that created and maintained by the Spanning Tree Protocol (STP).

To use GARP, a GARP application must be defined. The AT-S62 management software has one GARP application presently implemented, GVRP.

The GARP application specifies what the attribute represents.

GARP defines the architecture, rules of operation, state machines and variables for the registration and de-registration of attribute values. By itself, GARP is not directly used by devices in a bridged LAN. It is the applications of GARP that perform meaningful actions. The use of GVRP allows dynamic filter entries for VLAN membership to be distributed among the forwarding databases of VLAN-active switches.

A GARP Participant in a switch or an end station consists of a GARP Application component, and a *GARP Information Declaration* (GID) component associated with each port of the switch. One such GARP Participant exists per port, per GARP Application. The propagation of information between GARP Participants for the same Application in a switch is carried out by the *GARP Information Propagation* (GIP) component. Protocol exchanges take place between GARP Participants by means of LLC Type 1 services, using the group MAC address and PDU format defined for the GARP Application concerned.

Every instance of a GARP application includes a database to store the values of the attributes. Within GARP, attributes are mapped to GID indexes.

The architecture of GARP is shown in Figure 198.

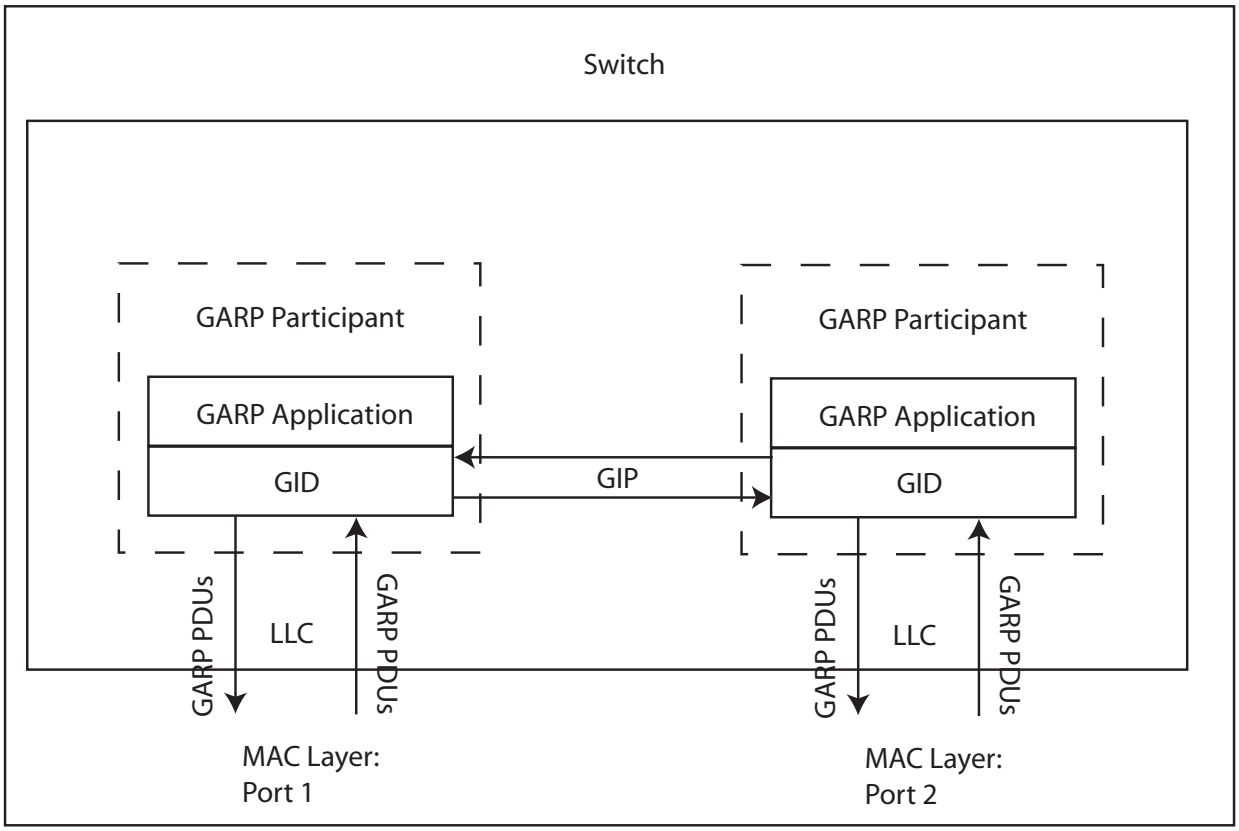

Figure 198. GARP Architecture

The GARP Application component of the GARP Participant is responsible for defining the semantics associated with the parameter values and operators received in GARP PDUs, and for generating GARP PDUs for transmission. The Application makes use of the GID component, and the state machines associated with the operation of GID, in order to control its protocol interactions.

An instance of GID consists of the set of state machines that define the current registration and declaration state of all *attribute* values associated with the GARP Participant. Separate state machines exist for the

Applicant and Registrar. This is shown in Figure 199.

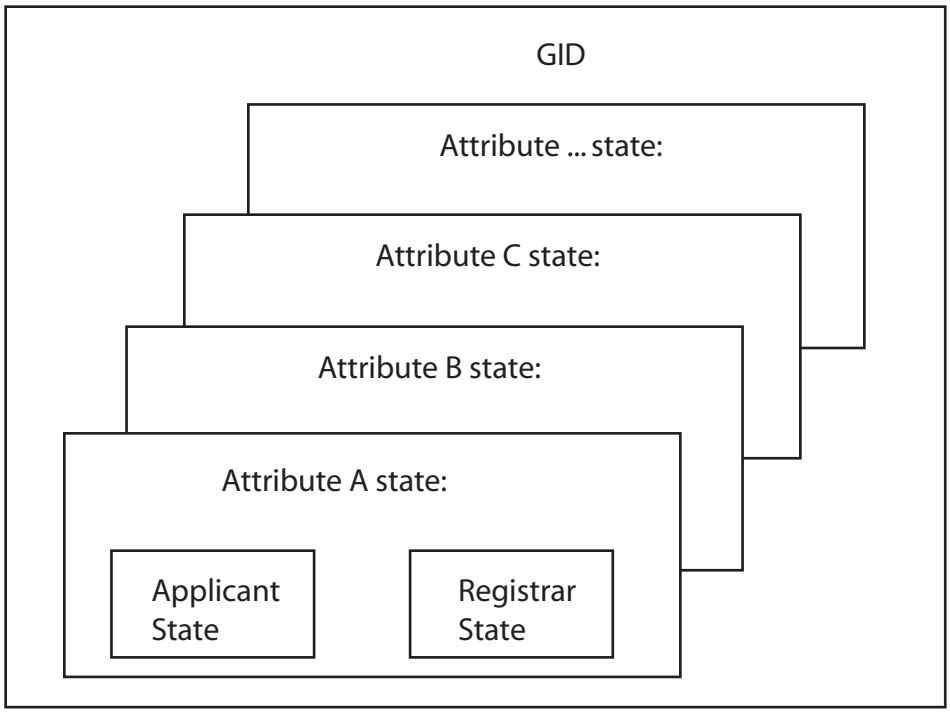

Figure 199. GID Architecture

GARP registers and de-registers *attribute* values through GARP messages sent at the GID level. A GARP Participant that wishes to make a declaration (an Applicant registering an *attribute* value) sends a JoinIn or JoinEmpty message. An Applicant that wishes to withdraw a declaration (de-registering an *attribute* value) sends a LeaveEmpty or LeaveIn message. Following the de-registration of an *attribute* value, the Applicant sends a number of Empty messages. The purpose of the Empty message is to prompt other Applicants to send JoinIn/JoinEmpty messages. For the GARP protocol to be resilient against multiple lost messages, a LeaveAll message is available. Timers are used in the state machines to generate events and control state transitions.

The job of the Applicant is twofold:

- To ensure that this Participant's declarations are registered by other Participants' Registrars
- To ensure that other Participants have a chance to re-declare (rejoin) after anyone withdraws a declaration (leaves).

The Applicant is therefore looking after the interests of all would-be Participants. This allows the Registrar to be very simple.

The job of the Registrar is to record whether an attribute is registered, in the process of being de-registered, or is not registered for an instance of GID.

To control the Applicant state machine, an Applicant Administrative Control parameter is provided. This parameter determines whether or not the Applicant state machine participates in GARP protocol exchanges. The default value has the Applicant participating in the exchanges.

To control the Registrar state machine, a Registrar Administrative Control parameter is provided. Basically, this parameter determines whether or not the Registrar state machine listens to incoming GARP messages. The default value has the Registrar listening to incoming GARP messages.

The propagation of information between GARP Participants for the same Application in a switch is carried out by the GIP component. The operation of GIP is dependent upon STP being enabled on a port, as only ports in the STP Forwarding state are eligible for membership to the GIP connected ring. Ports in the GIP connected ring propagate GID Join and Leave requests to notify each other of attribute registrations and de-registrations. The operation of GIP allows ports in the switch to share information between themselves and the LANs/end stations to which the ports are connected.

If a port enters the STP Forwarding state and the GARP application that the port belongs to is enabled, then the port is added to the GIP connected ring for the GARP application. All attributes registered by other ports in the GIP connected ring is propagated to the recently connected port. All attributes registered by the recently connected port is propagated to all other ports in the GIP connected ring.

Similarly, if a port leaves the STP Forwarding state and the GARP application that the port belongs to is enabled, then the port is removed from the GIP connected ring for the GARP application. Prior to removal, GID leave requests are propagated to all other ports in the GIP connected ring if the port to be removed has previously registered an attribute and no other port in the GIP connected ring has registered that attribute. The operations of GIP can be enabled or disabled by user command.

Configuring GVRP

This section contains the procedure for configuring GVRP. The timers in the following menus are in increments of centi seconds, which are hundredths of a second.

To configure GVRP, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **8** to select Configure GARP-GVRP.

   The GARP-GVRP Menu is shown in Figure 200.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch

User: Manager                                  11:20:02 02-Jan-2006

                            GARP-GVRP Menu

1 - GVRP Status ........... Disabled
2 - GVRP GIP Status ....... Enabled
3 - GVRP Join Timer ....... 20
4 - GVRP Leave Timer ...... 60
5 - GVRP Leave All Timer .. 1000

P - GVRP Port Parameters
O - Other GVRP Parameters Menu
D - Reset GVRP to Defaults

R - Return to Previous Menu

Enter your selection?
```

Figure 200. GARP-GVRP Menu

3. Type **1** - GVRP Status to enable or disable GVRP.

   The following prompt is displayed:

   ```
   Enter your new value (E-Enabled, D-Disabled):
   ```

4. Choose one of the following:

   E to enable GVRP.

   D to disable GVRP. This is the default setting.

5. Type **2** - GVRP GIP Status to enable or disable GIP.

The following prompt is displayed:

```
Enter your new value (E-Enabled, D-Disabled):
```

6. Choose one of the following:

   E to enable GIP.

   D to disable GIP.

> **Note**
> Do not disable GIP if you intend to use GVRP. GIP is required to propagate VLAN information among the ports of the switch.

> **Caution**
> The following steps change the three GVRP timers. The settings for these timers must be the same on all GVRP-active devices in your network.

7. Type **3** - GVRP Join Timer to change the value of the Join Timer.

   The following prompt is displayed:

   ```
   Enter new value (in centi seconds): [10 to 60] -> 20
   ```

8. Enter a new value for the Join Timer field in centi seconds which are one hundredths of a second. The default is 20 centiseconds.

   If you change this field, it must in relation to the GVRP Leave Timer according to the following equation:

   Join Timer <= (2 x (GVRP Leave Timer))

9. Type **4** - GVRP Leave Timer to enter a new value for this field.

   The following prompt is displayed:

   ```
   Enter new value (in centi seconds): [30 to 180] -> 60
   ```

10. Type **5** - GVRP Leave All Timer to enter a new value for this field. The default is 60 centiseconds.

    The following prompt is displayed:

    ```
    Enter new value (in centi seconds): [500 to 3000] -> 1000
    ```

11. Enter a value in centiseconds. The default is 1000 centiseconds.

Enabling or Disabling GVRP on a Port

This procedure enables and disables GVRP on a switch port. The default setting for GVRP on a port is enabled. Only those ports where GVRP is enabled transmit PDUs.

> **Note**
> To protect against unauthorized access to restricted areas of your network, Allied Telesyn recommends disabling GVRP on unused ports and those ports that are connected to GVRP-inactive devices. For further information, refer to "GVRP and Network Security" on page 585.

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **8** to select Configure GARP-GVRP.

   The GARP-GVRP menu is shown in Figure 200 on page 591.

3. Type **P** - GVRP Port Parameters to configure the switch ports.

   The GVRP Port Parameters Menu is shown in Figure 201.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                         Production Switch

User: Manager                                  11:20:02 02-Jan-2006

                        GVRP Port Parameters

1 - Configure GVRP Port Settings
2 - Display GVRP Port Configuration

R - Return to Previous Menu

Enter your selection?
```

Figure 201. GVRP Port Parameters Menu

4. Type **1** to configure GVRP Port Settings.

   The following prompt is displayed:

   ```
   Enter port-list:
   ```

5. Enter a port. You can configure more than one port at a time.

The Configure GVRP Port Settings Menu is shown in Figure 202.

```
        Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch

User: Manager                                11:20:02 02-Jan-2006

                    Configure GVRP Port Settings

Configuring Port 1-2

1 - Port Mode ............. Normal

R - Return to Previous Menu

Enter your selection?
```

Figure 202. Configure GVRP Port Settings Menu

6. Type **1** - Port Mode.

   The following prompt is displayed:

   ```
   Enter mode (0-Normal, 1-None): [0 to 1] -> 0
   ```

7. Type either **0** to select Normal or **1** to select None. A setting of Normal means the port processes and propagates GVRP information. This is the default setting. A setting of None prevents the port from processing GVRP information and from transmitting PDUs.

   A change to GVRP port mode is immediately activated on a port.

8. If you want to view the current port settings, from the GVRP Port Parameters menu, type **2** to display the GVRP port configuration.

   The Display GVRP Port Configuration Menu is shown in Figure 203.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch

User: Manager                                  11:20:02 02-Jan-2006

                    Display GVRP Port Configuration

GARP Port Parameters

Mode Normal ............. 1-2
Mode None ............... 3-26

U - Update
R - Return to Previous Menu

Enter your selection?
```

Figure 203. Display GVRP Port Configuration Menu

9. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

   Your changes are saved.

Converting a Dynamic GVRP VLAN

This procedure converts a dynamic GVRP VLAN into a static VLAN. You can perform this procedure to permanently retain the VLANs the switch learned through GVRP.

Note
This procedure cannot convert a dynamic GVRP port in a static VLAN into a static port. For that you must manually modify the static VLAN, specifying the dynamic port as either a tagged or untagged member of the VLAN.

To convert a dynamic GVRP VLAN to a static VLAN, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **4** to select Configure VLANs.

   The Configure VLANs menu is shown in Figure 189 on page 560.

   Note
   If option “4 - Configure VLANs” is not displayed in the menu, the switch is running a multiple VLAN mode. To change a switch’s VLAN mode, refer to “Selecting a VLAN Mode” on page 612.

3. From the Configure VLANs menu, type **2** to select Modify VLAN.

   The Modify VLAN menu is shown in Figure 191 on page 565.

4. Type **2** to select Change GARP VLAN.

   The following prompt is displayed:

   ```
   Enter VLAN ID: [1 to 4096] ->
   ```

5. Enter the VID of the dynamic GVRP VLAN you want to convert into a static VLAN. You can specify only one VLAN at a time.

   The dynamic GVRP VLAN is changed to a static VLAN. To confirm this, refer to “Displaying VLANs” on page 569.

6. Return to the Main Menu and type **S** to select Save Configuration Changes.

Displaying GVRP Parameters and Statistics

To display GVRP counters, database, state machine, and GIP connected ports ring, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **8** to select Configure GARP-GVRP.

   The GARP-GVRP Menu is shown in Figure 200 on page 591.

3. From the GARP-GVRP Menu, select **O** - Other GVRP Parameters Menu.

   The Other GARP Port Parameters Menu is shown in Figure 204.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch

User: Manager                                    11:20:02 02-Jan-2006

                      Other GARP Port Parameters

1 - Display GVRP Counters
2 - Display GVRP Database
3 - Display GIP Connected Ports Ring
4 - Display GVRP State Machine

R - Return to Previous Menu

Enter your selection?
```

Figure 204. Other GARP Port Parameters Menu

Each option is reviewed in a separate subsection below.

GVRP Counters

Option 1 - Display GVRP Counters in the Other GARP Port Parameters displays the GVRP Counters Menu (page 1) as shown in Figure 205.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch

User: Manager                                    11:20:02 02-Jan-2006

                              GVRP Counters

Receive:                                  Transmit:
--------                                  ---------
Total GARP Packets          41            Total GARP Packets    166
Invalid GARP Packets        0


Discarded:
-----------
GARP Disabled               0             GARP Disabled         0
Port Not Listening          0             Port Not Sending      3117
Invalid Port                0
Invalid Protocol            0
Invalid Format              0
Database Full               0

N - Next Page
U - Updated Display
R - Return to Previous Menu

Enter your selection?
```

Figure 205. GVRP Counters Menu (page 1)

The statistics span two menus. To display the second menu, type **N** to select Next Page. The second menu is shown in Figure 206. The information in both menus is for display purposes only.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch

User: Manager                                      11:20:02 02-Jan-2006

                              GVRP Counters

Receive:                                Transmit:
--------                                ---------
GARP Messages:
---------------
LeaveAll                      7         LeaveAll                    77
JoinEmpty                     0         JoinEmpty                   58
JoinIn                        68        JoinIn                      285
LeaveEmpty                    0         LeaveEmpty                  1
LeaveIn                       0         LeaveIn                     0
Empty                         5         Empty                       21
Bad Message                   0
Bad Attribute                 0


P - Previous Page
U - Updated Display
R - Return to Previous Menu

Enter your selection?
```

Figure 206. GVRP Counters Menu (page 2)

The GVRP counters in the menus are defined in Table 21.

Table 21. GVRP Counters

Parameter	Meaning
Receive: Total GARP Packets	Total number of GARP PDUs received by this GARP application.
Transmit: Total GARP Packets	Total number of GARP PDUs transmitted by this GARP application.
Receive: Invalid GARP Packets	Number of invalid GARP PDUs received by this GARP application.
Receive Discarded: GARP Disabled	Number of received GARP PDUs discarded because the GARP application was disabled.
Transmit Discarded: GARP Disabled	Number of GARP PDUs discarded because the GARP application was disabled. This counter is incremented when ports are added to or deleted from the GARP application arising from port movements in the underlying VLAN or STP.

Table 21. GVRP Counters

Parameter	Meaning
Receive Discarded: Port Not Listening	Number of GARP PDUs discarded because the port that received the PDUs was not listening, that is, MODE=NONE was set on the port.
Transmit Discarded: Port Not Sending	Number of GARP PDUs discarded because the port that the PDUs were to be transmitted on was not sending, that is, MODE=NONE was set on the port.
Receive Discarded: Invalid Port	Number of GARP PDUs discarded because the port that received the PDU does not belong to the GARP application.
Receive Discarded: Invalid Protocol	Number of GARP PDUs discarded because the GARP PDU contained an invalid protocol.
Receive Discarded: Invalid Format	Number of GARP PDUs discarded because the format of the GARP PDU was not recognized.
Receive Discarded: Database Full	Number of GARP PDUs discarded because the database for the GARP application was full, that is, the maximum number of attributes for the GARP application is in use.
Receive GARP Messages: LeaveAll	Number of GARP LeaveAll messages received by the GARP application.
Transmit: GARP Messages: LeaveAll	Number of GARP LeaveAll messages transmitted by the GARP application.
Receive GARP Messages: JoinEmpty	Total number of GARP JoinEmpty messages received for all attributes in the GARP application.
Transmit GARP Messages: JoinEmpty	Total number of GARP JoinEmpty messages transmitted for all attributes in the GARP application.
Receive GARP Messages: JoinIn	Total number of GARP JoinIn messages received for all attributes in the GARP application.
Transmit GARP Messages: JoinIn	Total number of GARP JoinIn messages transmitted for all attributes in the GARP application.

Table 21. GVRP Counters

Parameter	Meaning
Receive GARP Messages: LeaveEmpty	Total number of GARP LeaveEmpty messages received for all attributes in the GARP application.
Transmit GARP Messages: LeaveEmpty	Total number of GARP LeaveEmpty messages transmitted for all attributes in the GARP application.
Receive GARP Messages: LeaveIn	Total number of GARP LeaveIn messages received for all attributes in the GARP application.
Transmit GARP Messages: LeaveIn	Total number of GARP LeaveIn messages transmitted for all attributes in the GARP application.
Receive GARP Messages: Empty	Total number of GARP Empty messages received for all attributes in the GARP application.
Transmit GARP Messages: Empty	Total number of GARP Empty messages transmitted for all attributes in the GARP application.
Receive GARP Messages: Bad Message	Number of GARP messages that had an invalid Attribute Type value, an invalid Attribute Length value or an invalid Attribute Event value.
Receive GARP Messages: Bad Attribute	Number of GARP messages that had an invalid Attribute Value value.

GVRP Database

Option 2 - Display GVRP Database in the Other GARP Port Parameters displays the GVRP Database Menu as shown in Figure 207.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch

User: Manager                                    11:20:02 02-Jan-2006

                             GVRP Database

GARP Application: GVRP
GID index    VLAN ID           Used        GID index    VLAN ID   Used
-------------------------------------------------------------------------

  0          1                 Yes          1            3         Yes
  2          2                 Yes

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 207. GVRP Database Menu

The columns in the menu are defined in Table 22. The information is for viewing purposes only.

Table 22. GARP Database Parameters

Parameter	Meaning
GARP Application	Identifies the GARP application, that is, “GVRP”.
GID index	Value of the GID index corresponding to the attribute. GID indexes begin at 0. If the GARP application has no attributes presently registered, “No attributes have been registered” is displayed.
VLAN ID	Value of the attribute.
Used	Indicates whether the GID index is currently being used by any port in the GARP application. The definition of “used” is whether the Applicant and Registrar state machine for the GID index are in a non-initialized state, that is, not in {Vo, Mt} state. The value of this parameter is either “Yes” or “No”.

GIP Connected Ports Ring

Option 3 - Display GIP Connected Ports Ring in the Other GARP Port Parameters displays the GIP Connected Ports Ring Menu as shown in Figure 208.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch

User: Manager                                  11:20:02 02-Jan-2006

                      GIP Connected Ports Ring

GARP Application: GVRP
GIP Context ID: 0, STP ID: 0
------------------------------------------------------------

2 -> 8 -> 4

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 208. GIP Connected Ports Ring Menu

The information in the menu is defined in Table 23. This information is for viewing purposes only.

Table 23. GIP Connected Ports Ring Parameters

Parameter	Meaning
GARP Application	Identifies the GARP application, that is, "GVRP."
GIP Context ID	A number assigned to the instance for the GIP context.
STP ID	Present if the GARP application is GVRP; identifies the spanning tree instance associated with the GIP context.
Connected Ring	Ring of connected ports. Only ports presently in the STP Forwarding state are eligible for membership to the GIP connected ring. If no ports exist in the GIP connected ring, "No ports are connected" is displayed. If the GARP application has no ports, "No ports have been assigned" is displayed.

GVRP State Machine

Option 4 - Display GVRP State Machine in the Other GARP Port Parameters displays the GVRP State Machine Menu (page 1) as shown in Figure 209.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                            Production Switch

 User: Manager                                          11:20:02 02-Jan-2006

                               GVRP State Machine

 Enter a VLAN ID for displaying the state machine: [1 to 4094] -> 1
```

Figure 209. GVRP State Machine Menu (page 1)

Entering a VLAN ID displays the GVRP State Machine Menu (page 2) as shown in Figure 210.

```
                  Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                                  Production Switch

User: Manager                                                 11:20:02 02-Jan-2006

                                     GVRP State Machine

State Machine for VLAN: 1
 Port   App    Reg   | Port   App   Reg   | Port   App   Reg |  Port   App   Reg |
-------------------------------------------------------------------------------
 2.1    Qa     Fix   | 2.2    Qa    Fix   | 2.3    Qa    Fix |  2.4    Qa    Fix |
 2.5    Qa     Fix   | 2.6    Qa    Fix   | 2.7    Qa    Fix |  2.8    Qa    Fix |
 3.1    Qa     Fix   | 3.2    Qa    Fix   | 3.3    Qa    Fix |  3.4    Qa    Fix |
 8.1    Qa     Fix   | 8.2    Qa    Fix   | 8.3    Qa    Fix |  8.4    Qa    Fix |
 8.5    Qa     Fix   | 8.6    Qa    Fix   | 8.7    Qa    Fix |  8.8    Qa    Fix |

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 210. Display GVRP State Machine Menu (page 2)

The information in the menu is defined in Table 24. This information is for viewing purposes only.

Table 24. GVRP State Machine Parameters

Parameter	Meaning
Port	Port number on the switch; this port belongs to the GARP application. If the GARP application has no ports, “No ports have been assigned” is displayed.

Table 24. GVRP State Machine Parameters

Parameter	Meaning
App	Applicant state machine for the GID index on that particular port. One of:
	Normal Participant Management state:
	"Vo" Very Anxious Observer
	"Ao" Anxious Observer
	"Qo" Quiet Observer
	"Lo" Leaving Observer
	"Vp" Very Anxious Passive Member
	"Ap" Anxious Passive Member
	"Qp" Quiet Passive Member
	"Va" Very Anxious Active Member
	"Aa" Anxious Active Member
	"Qa" Quiet Active Member
	"La" Leaving Active Member
App (*Continued*)	*Non-Participant Management state:*
	"Von" Very Anxious Observer
	"Aon" Anxious Observer
	"Qon" Quiet Observer
	"Lon" Leaving Observer
	"Vpn" Very Anxious Passive Member
	"Apn" Anxious Passive Member
	"Qpn" Quiet Passive Member
	"Van" Very Anxious Active Member
	"Aan" Anxious Active Member
	"Qan" Quiet Active Member
	"Lan" Leaving Active Member
	The initialized state for the Applicant is Vo.

Table 24. GVRP State Machine Parameters

Parameter	Meaning	
Reg	Registrar state machine for the GID index on that particular port. One of:	
	“Mt”	Empty
	“Lv3”	Leaving substate 3 (final Leaving substate)
	“Lv2”	Leaving substate 2
	“Lv1”	Leaving substate 1
	“Lv”	Leaving substate (initial Leaving substate)
	“In”	In
	“Fix”	Registration Fixed
	“For”	Registration Forbidden
	The initialized state for the Registrar is Mt.	

Chapter 26

Multiple VLAN Modes

This chapter describes the multiple VLAN modes and how to select a mode.

This chapter contains the following sections:

Multiple VLAN Mode Overview

Multiple VLAN modes simplify the task of configuring the switch in network environments that require a high degree of network segmentation. In a multiple VLAN mode, the ports on a switch are prohibited from forwarding traffic to each other and can only forward traffic to a user designated uplink port. These configurations isolate the traffic on each port from all other ports, while providing the ports with access to an uplink port.

The AT-S62 software supports two types of multiple VLAN modes:

- 802.1Q-compliant Multiple VLAN mode
- Multiple VLAN mode (also referred to as non-802.1Q compliant Multiple VLAN mode)

Each mode uses a different technique to isolate the ports and their traffic. The first method uses VLANs while the second uses port mapping. The uplink port is also different in each mode. In one the port is a tagged port and in the other untagged. This is explained in the following subsections.

Note
The multiple VLAN mode feature is supported only in single switch (i.e. edge switch) environments. This means that cascading of switches while in a Multiple VLAN mode is not allowed.

Configuring multiple VLANs on a cascaded switch can possibly result in disconnection of network paths between switches unless the port used to link the switch (being configured for Multiple VLANs mode) is configured as uplink VLAN port.

Configuring multiple VLANs on cascaded switches can also affect Enhanced Stacking as the Master switch may not be able to detect member switches beyond the first cascaded switch.

802.1Q-Compliant Multiple VLAN mode

In this mode, each port is placed into a separate VLAN as an untagged port. The VLAN names and VID numbers are based on the port numbers. For example, the VLAN for Port 4 is named Client_VLAN_4 and is given the VID of 4, the VLAN for Port 5 is named Client_VLAN_5 and has a VID of 5, and so on.

The VLAN configuration is accomplished automatically by the switch. Once you have selected the mode and an uplink port, the switch forms the VLANs. It also assigns the PVID values as well. For example, the PVID for Port 4 is assigned as 4, to match the VID of 4.

A user designated port on the switch functions as an uplink port, which can be connected to a shared device, such as a router for access to a

WAN. This port is placed as a tagged port in each VLAN. Thus, while the switch ports are separated from each other in their individual VLANs, they all have access to the uplink port.

The uplink port also has its own VLAN, where it is an untagged member. This VLAN is called Uplink_VLAN.

Note
In 802.1Q Multiple VLAN mode, the device connected to the uplink port must be IEEE 802.1Q-compliant.

An example of the 802.1Q-compliant VLAN mode is shown in Table 25. The table shows the VLANs on an AT-8524M switch where Port 25, a port on an expansion module, has been selected as the uplink port.

Table 25. 802.1Q-Compliant Multiple VLAN Example

VLAN Name	VID	Untagged Port	Tagged Port
Client_VLAN_1	1	1	25
Client_VLAN_2	2	2	25
Client_VLAN_3	3	3	25
Client_VLAN_4	4	4	25
Client_VLAN_5	5	5	25
Client_VLAN_6	6	6	25
Client_VLAN_7	7	7	25
Client_VLAN_8	8	8	25
Client_VLAN_9	9	9	25
Client_VLAN_10	10	10	25
Client_VLAN_11	11	11	25
Client_VLAN_12	12	12	25
Client_VLAN_13	13	13	25
Client_VLAN_14	14	14	25
Client_VLAN_15	15	15	25
Client_VLAN_16	16	16	25
Client_VLAN_17	17	17	25
Client_VLAN_18	18	18	25

VLAN Name	VID	Untagged Port	Tagged Port
Client_VLAN_19	19	19	25
Client_VLAN_20	20	20	25
Client_VLAN_21	21	21	25
Client_VLAN_22	22	22	25
Client_VLAN_23	23	23	25
Client_VLAN_24	24	24	25
Uplink_VLAN	25	25	
Client_VLAN_26	26	26	25

This highly segmented configuration is useful in situations where traffic generated by each end node or network segment connected to a port on the switch needs to be kept separate from all other network traffic, but still have access to an uplink port to a WAN. Unicast traffic received by the uplink port is effectively directed to the appropriate port and end node, and is not directed to any other port on the switch.

The 802.1Q Multiple VLAN configuration is appropriate when the device connected to the uplink port is IEEE 802.1Q compatible, meaning that it can handle tagged packets.

When you select the 802.1Q-compliant VLAN mode, you are asked to specify the Uplink VLAN port. You can specify only one uplink port. The switch automatically configures the ports into the separate VLANs.

Note
The uplink VLAN is the management VLAN. Any remote management of the switch must be made through the uplink VLAN.

Non-802.1Q Compliant Multiple VLAN Mode

Unlike the 802.1Q-compliant VLAN mode, which isolates port traffic by placing each port in a separate VLAN, this mode forms one VLAN with a VID of 1 that encompasses all ports. Traffic isolation is established through port mapping. The result, however, is the same. Ports are permitted to forward traffic only to the designated uplink port and to no other port, even when they receive a broadcast packet.

With this mode the uplink port is untagged. You would want to use this mode when the device connected to the uplink port is not IEEE 802.1Q compatible, meaning that the device cannot handle tagged packets.

Note

When the uplink port receives a packet with a destination MAC address that is not in the MAC address table, the port will broadcast the packet to all switch ports. This can result in ports receiving packets that are not intended for them.

It should also be noted that a switch operating in this mode can be remotely managed through any port on the switch, not just the uplink port.

Selecting a VLAN Mode

The following procedure explains how to select a VLAN mode. Available modes are:

- ❒ User configured VLAN mode (port-based and tagged VLANs)
- ❒ IEEE 802.1Q Compliant Multiple VLAN mode
- ❒ Non-IEEE 802.1Q Compliant Multiple VLAN mode

> **Note**
> Any port-based or tagged VLANs you created are not retained when you change the VLAN mode from the user configured mode to a multiple VLAN mode and, at some point, reset the switch. The user configured VLAN information is lost and will need to be recreated if you return the switch to the user configured VLAN mode.

To select a VLAN mode, perform the following steps:

1. From the Main Menu, type **2** to select VLAN Configuration.
2. From the VLAN Configuration menu, type **2** to select VLAN Mode.

   The following prompt is displayed:

   ```
   Enter VLAN Mode (U-UserConfig, M-Multiple, Q-802.1Q
   Multiple VLANs) ->
   ```

3. Type **Q** to activate 802.1Q Multiple VLAN mode, **M** for Non-802.1Q compliant multiple VLAN mode, or **U** to create your own port-based and tagged VLANs. User configured is the default setting.

   If you enter **Q** or **M**, the following prompt is displayed:

   ```
   Enter Uplink VLAN Port number -> [1 to 26] ->
   ```

4. Enter the port number on the switch that will function as the uplink port for the other ports. You can specify only one port.

   The following prompt is displayed:

   ```
   SUCCESS
   Press any key to continue ...
   ```

   The new VLAN mode is now active on the switch.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying VLAN Information

To view the VLANs on the switch while the unit is operating in Multiple VLAN mode, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **6** to select Show VLANs.

   An example of the Show VLANs menu is shown in Figure 211.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch

                              Show VLANs

User: Manager                                        11:20:02 02-Jan-2006

VID    VLAN Name      Untagged (U) / Tagged (T)
-------------------------------------------------------------------


1    Client_VLAN_1             U: 1, 15
                               T:
2    Client_VLAN_2             U: 2, 15
                               T:
3    Client_VLAN_3             U: 3, 15
                               T:
4    Client_VLAN_4             U: 4, 15
                               T:
5    Client_VLAN_5             U: 5, 15
                               T:
6    Client_VLAN_6             U: 6, 15
                               T:
7    Client_VLAN_5             U: 7, 15
                               T:
8    Client_VLAN_6             U: 8, 15
                               T:

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 211. Show VLANs Menu, Multiple VLANS

The menu contains the following columns of information:

VID - The VLAN ID.

VLAN Name - The name of the VLAN.

Untagged (U) / Tagged (T) - The untagged and tagged ports that are part of the VLAN.

Chapter 27

Protected Ports VLANs

This chapter explains protected ports VLANs. It contains the following sections:

Protected Ports VLAN Overview

The purpose of a protected ports VLAN is to allow multiple ports on the switch to share the same uplink port but not share traffic with each other.

This feature has some of the same characteristics as the multiple VLAN modes described in the previous chapter, but it offers several advantages. One is that it provides more flexibility. With the multiple VLAN modes, you can select only one uplink port which is shared by all the other ports. Also, you are not allowed to modify the configuration.

With protected ports VLANs, you can create LAN segments that consist of more than one port and you can specify multiple uplink ports.

Another advantage is that the switch can support protected ports VLANs as well as port-based and tagged VLANs simultaneously, something that is not allowed with the multiple VLAN modes.

An important concept of this feature is *groups*. A group is a selection of one or more ports that function as a LAN segment within the VLAN. The ports in each group are independent of the ports in the other groups of the VLAN. The ports of a group can share traffic only amongst themselves and with the uplink port, but not with ports in other groups of the same VLAN.

A protected ports VLAN can consist of two or more groups and a group can consist of one or more ports. The ports of a group can be either tagged or untagged.

This type of VLAN also shares some common features with tagged VLANs, where one or more ports are shared by different LAN segments. But there are significant differences. First, all the ports in a tagged VLAN are considered a LAN segment, while the ports in a protected ports VLAN, though residing within a single VLAN, are subdivided into the smaller unit of groups, which represent the LAN segments.

Second, a tagged VLAN, by its nature, contains one or more tagged ports. These are the ports that are shared among one or more tagged VLANs. The device connected to a tagged port must be 802.1Q compliant and it must be able to handle tagged packets.

In contrast, the uplink port in a protected ports VLAN, which is shared by the ports in the different groups, can be either tagged or untagged. The device connected to it does not necessarily need to be 802.1Q compliant.

Note
For explanations of VIDs and tagged and untagged ports, refer to Chapter 24, ”Port-based and Tagged Virtual LANs” on page 545.

To create a protected ports VLAN, you perform many of the same steps that you do when you create a new port-based or tagged VLAN. You give it a name and a unique VID, and you indicate which of the ports will be tagged and untagged. What makes creating this type of VLAN different is that as part of the procedure you must create the individual groups within the VLAN by assigning the ports to their groups.

Here is an example of a protected ports VLAN. The first table lists the name of the VLAN, the VID, and the tagged and untagged ports. It also indicates which port will function as the uplink port, in this case port 22. The second table lists the different groups in the VLAN and the ports for each group.

Name	Internet_VLAN_1
VID	8
Untagged Ports in VLAN	1-10, 22, 25
Tagged Ports in VLAN	none
Uplink Port(s)	22

Group Number	Port(s)
1	1-2
2	3
3	4
4	5-7
5	8
6	9-10

Allied Telesyn recommends that you create tables similar to this before you create your own protected ports VLAN. You are prompted for this information when you create the VLAN, and having the tables handy will make the job easier.

Protected Ports VLAN Guidelines

Following are some guidelines for implementing protected ports VLANS:

- A switch can contain multiple protected ports VLANs.
- A protected ports VLAN should contain a minimum of two groups. A protected ports VLAN of only one group has little value. Create a port-based or tagged VLAN instead.

- A protected ports VLAN can contain any number of groups.
- A group can contain any number of ports.
- The ports of a group can be tagged or untagged.
- Each group must be assigned a unique group number on the switch. The number can be from 1 to 256.
- A protected ports VLAN can contain more than one uplink port.
- An uplink port can be either tagged or untagged.
- Uplink ports can be shared among more than one protected ports VLAN, but only if they are tagged.
- A switch can contain a combination of port-based and tagged VLANs and protected ports VLANs.
- A port that is a member of a group in a protected ports VLAN cannot be a member of a port-based or tagged VLAN.
- A group can be a member of more than one protected ports VLAN at a time. However, the port members of the group must be identical in both VLANs and the ports must be tagged.
- You cannot create protected ports VLANs when the switch is operating in a multiple VLAN mode.
- A port that is already an untagged member of a protected ports VLAN cannot be made an untagged member of another VLAN until it is first removed from its current VLAN assignment and returned to the Default_VLAN.

Creating a Protected Ports VLAN

To create a new protected ports VLAN, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.
2. From the VLAN Configuration Menu, type **4** to select Configure VLANs.

   Note
   If the menu does not include selection 4, Configure VLANs, the switch is running a multiple VLAN mode. To change the switch's VLAN mode, refer to "Selecting a VLAN Mode" on page 612.

3. From the Configure VLANs Menu, type **1** to select Create VLAN.

   The Create VLAN Menu is shown in Figure 212.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                   Production Switch

User: Manager                       11:20:02 02-Jan-2006

                        Create VLAN

1 - VLAN Name ............
2 - VLAN ID (VID) ........ 2
3 - Tagged Ports .........
4 - Untagged Ports .......
5 - Protected Ports ...... No

C - Create VLAN
R - Return to Previous Menu

Enter your selection?
```

Figure 212. Create VLAN Menu

4. Type **1** to select VLAN Name.

   The following prompt is displayed:

   ```
   Enter new value ->
   ```

5. Type a name for the new protected ports VLAN.

   The name can be from one to fifteen alphanumeric characters in length. The name should reflect the function of the nodes that will be a part of the protected ports VLAN (for example, InternetGroups). The name cannot contain spaces or special characters, such as asterisks (*) or exclamation points (!).

Note
A VLAN must be assigned a name.

6. Type **2** to select VLAN ID (VID.

   The following prompt is displayed:

   ```
   Enter new value -> [2 to 4094] ->
   ```

7. Type a VID value for the new VLAN. The range for the VID value is 2 to 4094.

   The AT-S62 management software uses the next available VID number on the switch as the default value. It is important to note that the switch is only aware of the VIDs of the VLANs that exist on the device, and not those that might already be in use in the network. For example, if you add a new AT-8500 Series switch to a network that already contains VLANs that use VIDs 2 through 24, the AT-S62 management software still uses VID 2 as the default value when you create the first VLAN on the new switch, even though that VID number is already being used by another VLAN on the network. To prevent inadvertently using the same VID for two different VLANs, you should keep a list of all your network VLANs and their VID values.

   Note
   A VLAN must have a VID.

8. If the VLAN will contain tagged ports, type **3** to select Tagged Ports and specify the ports. If this VLAN will not contain any tagged ports, leave this field empty.

   You can specify the ports individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9).

9. Type **4** to select Untagged Ports and specify the ports on the switch to function as untagged ports in the VLAN. If this VLAN will not contain any untagged ports, leave this field empty.

   You can specify the ports individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9).

10. Type **5** to select Protected Ports**.**

    The following prompt is displayed:

    ```
    Enter New Value [Yes/No] ->
    ```

11. To make this a protected ports VLAN, type **Y**. If you do not want this to be a protected ports VLAN and want it to be a port-based or tagged VLAN, type **N**.

12. Type **C** to select Create VLAN**.**

    The following prompt is displayed:

    ```
    Enter Uplink Ports (4 - 12) ->
    ```

    The prompt will show the ports that you specified as belonging to the VLAN.

13. Enter the port in the VLAN that will function as the uplink port for the different VLAN groups. You can select more than one uplink port.

    The following prompt is displayed:

    ```
    Enter Group Ports (4 - 11) ->
    ```

    The prompt lists the ports in the VLAN, minus the uplink port you specified in the previous step.

14. Specify the ports of one of the groups of the protected ports VLAN. This can be a few as one port or as many as all the remaining ports of the VLAN. You can specify the ports of the group individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9).

    The following prompt is displayed:

    ```
    Enter Group Number ->
    ```

15. Enter a group number for the port(s). Each group on the switch must be have a unique group number. The range is 0 to 256.

16. If there are ports within the VLAN that still need to be assigned to a group, the prompt in Step 13 is displayed again, showing the unassigned ports. You must repeat Steps 14 and 15, creating additional groups, until all of the ports in the VLAN have been assigned to a group.

    After you create all of the groups, the following prompt is displayed:

    ```
    SUCCESS - Press any key to continue.
    Press any key to continue.
    ```

    The new protected ports VLAN and its groups are now active on the switch.

17. Press any key to return to the Configure VLANs Menu.

18. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Protected Ports VLAN

Please note the following before you perform this procedure:

- To modify this type of VLAN, you must recreate it by reselecting the uplink port(s) and reassigning the ports to the groups. For this reason Allied Telesyn recommends that before you perform this procedure you first display the details of the protected ports VLAN you want to modify and write down on paper the VLAN's current configuration (i.e., uplink port and port to group assignments). This information will make it easier for you to recreate the current configuration, with whatever modifications you want to make, when you perform the procedure. To display a VLAN's configuration, refer to "Displaying a Protected Port VLAN" on page 626.
- If you are adding untagged ports, the ports must be untagged members of the Default_VLAN or a port-based or tagged VLAN. They cannot be members of another protected ports VLAN.
- An untagged port removed from a VLAN is automatically returned to the Default_VLAN.
- A port that is already an untagged member of a protected ports VLAN cannot be made an untagged member of another VLAN until it is first removed from its current VLAN assignment and returned to the Default_VLAN.

Note
To modify a VLAN, you need to know its VID. To view VLAN VIDs, refer to "Displaying a Protected Port VLAN" on page 626.

To modify a protected ports VLAN, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration Menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration Menu, type **4** to select Configure VLANs.

   The Configure VLANs Menu is shown in Figure 189 on page 560.

   Note
   If selection 4, Configure VLANs, is not displayed in the menu, the switch is running a multiple VLAN mode. To change a switch's VLAN mode, refer to "Selecting a VLAN Mode" on page 612.

3. From the Configure VLANs Menu, type **2** to select Modify VLAN.

The Modify VLAN Menu is shown in Figure 191 on page 565.

4. Type **1** to select VLAN ID (VID).

   The following prompt is displayed:

   ```
   Enter new value -> [1 to 4096] ->
   ```

5. Enter the VID of the VLAN you want to modify.

   The Modify VLAN Menu expands to contain all relevant information about the VLAN, as shown in Figure 213.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                      11:20:02 02-Jan-2006

                        Modify VLAN

1 - VLAN Name ............... Internet_1
2 - VLAN ID (VID) ........... 3
3 - Tagged Ports ............ 7,9
4 - Untagged Ports ......... 20-24
5 - Protected Ports ........ Yes

M - Modify VLAN
R - Return to Previous Menu

Enter your selection?
```

Figure 213. Expanded Modify VLAN Menu

6. Adjust the following parameters as necessary.

   1 - VLAN Name
   Use this selection to change the name of a VLAN. The name can be from one to fifteen alphanumeric characters in length. The name cannot contain spaces or special characters, such as asterisks (*) or exclamation points (!).

   When you change a VLAN's name, observe the following guidelines:

- ❒ A VLAN's new name cannot be the same as the name of another VLAN on the same switch.
- ❒ You cannot change the name of the Default_VLAN.

> **Note**
> A VLAN must have a name.

   2 - VLAN ID (VID)
   This is the VLAN's VID value. You cannot change this value.

3 - Tagged Ports
Use this selection to add or remove tagged ports from the VLAN. You can specify the ports individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9). The new list of tagged ports will replace the existing tagged ports.

4 - Untagged Ports
Use this selection to add or remove untagged ports from the VLAN. You can specify the ports individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9). The new list of untagged ports will replace the existing list of untagged ports.

5 - Protected Ports
This option cannot be changed. To convert a protected ports VLAN into a tagged or port-based VLAN, you must first delete it and then recreate it as a tagged or port-based VLAN.

7. After making the desired changes, type **M** to select Modify VLAN.

   The following prompt is displayed:

   ```
   Enter Uplink ports list(4 - 12) ->
   ```

   This prompt will differ depending on the ports you specified as part of the protected ports VLAN.

8. Enter the port in the VLAN that will function as the uplink port for the different VLAN groups. You can select more than one uplink port.

   The following prompt is displayed:

   ```
   Enter Group ports list (4 - 11) ->
   ```

   The prompt now lists the ports in the VLAN, minus the uplink port you specified in the previous step.

9. Specify the ports of one of the groups of the protected ports VLAN. This can be a small as one port or as many as all the remaining ports of the VLAN. You can specify the ports of the group individually (e.g., 2,3,5), as a range (e.g., 7-9), or both (e.g., 2,5,7-9).

   The following prompt is displayed:

   ```
   Enter Group Number ->
   ```

10. Enter a group number for the port(s). Each group on the switch must be given a unique group number.

11. If there are ports within the VLAN that still need to be assigned to a group, the prompt in Step 8 is displayed again, showing the unassigned ports. You must repeat Steps 9 and 10, creating additional groups, until all of the ports in the VLAN have been assigned to a group.

After you have created all of the groups, this prompt is displayed:

```
SUCCESS - Press any key to continue.
Press any key to continue.
```

The modified protected ports VLAN and its groups are now active on the switch.

12. Press any key to return to the Configure VLANs Menu.
13. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying a Protected Port VLAN

To view the name, VID number, and member ports of all the VLANs on a switch, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration menu, type **5** to select Show VLANs.

   The Show VLANs Menu is shown in Figure 214.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch

User: Manager                                        11:20:02 02-Jan-2006

                                  Show VLANs

VID    VLAN Name    VLAN Type    Protocol   Member Port(s)
-------------------------------------------------------------------

1   Default_VLAN    Port Based              U: 25
                                            T:
4   Sales           Port Based              U: 1-11
                                            T:
5   Internet_VLAN   Protected               U: 12-24
                                            T: 25

N - Next Page
U - Update Display
D - Detail Information Display
R - Return to Previous Menu

Enter your selection?
```

Figure 214. Show VLANs Menu

The menu contains the following columns of information:

VID - The VLAN ID.

VLAN Name - The name of the VLAN.

VLAN Type - A VLAN type of “Port Based” indicates a port-based or tagged VLAN. A VLAN type of “GARP” indicates a VLAN created automatically by GVRP. A VLAN type of “Protected” indicates a protected ports VLAN.

Protocol - If this column is blank, the VLAN is a port-based, tagged, or protected ports VLAN. If it contains “GARP,” the VLAN or the port is a dynamic GVRP VLAN or a dynamic GVRP port of a static VLAN.

Untagged (U) / Tagged (T) - The ports of the VLAN. Tagged ports are designated with a "T" and untagged ports with a "U."

3. To view additional information about a protected ports VLAN, type **D** to select Detail Information Display.

   The following prompt is displayed:

   ```
   Enter new value ->
   ```

4. Enter the VID of the protected ports VLAN whose information you want to view.

   An example of the Show VLANs window for a protect ports VLAN is shown in Figure 215.

```
          Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch

User: Manager                                    11:20:02 02-Jan-2006

                               Show VLANs

VID   VLAN Name     VLAN Type    Protocol   Untagged (U) / Tagged (T)
-------------------------------------------------------------------

5     Internet_VLAN  Protected              U: 12-24      ┐
                     Protected              T: 25         ┘ Section 1
                     Group                     Ports      ┐
                     ---------------------------------    │
                     Uplink                    25         │
                     1                         12-13      │
                     2                         14-15      │ Section 2
                     3                         16         │
                     4                         17         │
                     5                         18-20      ┘

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 215. Show VLANs Menu

Section 1 lists all the tagged and untagged ports in the protected ports VLAN.

Section 2 lists each group in the VLAN, starting with the uplink port(s). The groups are listed by group number followed by the port numbers. For example, in Figure 215 the uplink port for the VLAN is port 25 and Group 1 consists of ports 12 and 13.

Deleting a Protected Ports VLAN

All untagged ports in a deleted protected ports VLAN are automatically returned to the Default_VLAN.

To delete a protected ports VLAN, perform the following procedure:

1. From the Main Menu, type **2** to select VLAN Configuration.

   The VLAN Configuration Menu is shown in Figure 188 on page 559.

2. From the VLAN Configuration Menu, type **4** to select Configure VLANs.

   The Configure VLANs Menu is shown in Figure 189 on page 560.

   Note
   If option 4, Configure VLANs, is not displayed in the menu if the switch is running a multiple VLAN mode. To change a switch's VLAN mode, refer to "Selecting a VLAN Mode" on page 612.

3. From the Configure VLANs Menu, type **3** to select Delete VLAN.

   The Delete VLAN Menu is shown in Figure 216.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch

User: Manager                          11:20:02 02-Jan-2006

                         Delete VLAN

1 - VLAN ID (VID) ........

R - Return to Previous Menu

Enter your selection?
```

Figure 216. Delete VLAN Menu

4. Type **1** to select VLAN ID (VID).

   The following prompt is displayed:

   ```
   Enter new value -> [2 to 4094] ->
   ```

5. Enter the VID of the VLAN you want to delete. You can specify only one VID at a time.

Note
You cannot delete the Default_VLAN, which has a VID of 1.

The Delete VLAN Menu expands to contain all relevant information about the VLAN, as shown in Figure 217.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                      11:20:02 02-Jan-2006

                        Delete VLAN

1 - VLAN Name ............... Internet_VLAN
2 - VLAN ID (VID) ........... 3
3 - Tagged Ports ............ 25
4 - Untagged Ports .......... 12-24
5 - Protected Ports ......... Yes

D - Delete VLAN
R - Return to Previous Menu

Enter your selection?
```

Figure 217. Expanded Delete VLAN Menu

6. Type **D** to delete the VLAN or **R** to cancel the procedure.

   If you select to delete the VLAN, the following confirmation prompt is displayed:

   ```
   Are you sure you want to delete this VLAN [Yes/No] ->
   ```

7. Type **Y** to delete the VLAN or **N** to cancel the procedure. Press Return.

   If you select Yes, the VLAN is deleted and the following message is displayed:

   ```
   SUCCESS
   Please make sure to manually delete any static multicast
   MAC address(es) entries for this VLAN
   Press any key to continue ...
   ```

   All untagged ports in the deleted VLAN are returned to the Default_VLAN as untagged ports.

   Any static addresses assigned to the ports of the VLAN are now obsolete, because the VLAN has been deleted. Those addresses should be deleted from the MAC address table. For instructions on how to delete addresses, refer to "Deleting Unicast and Multicast MAC Addresses" on page 132.

8. Press any key.

9. Repeat this procedure starting with Step 4 to delete other VLANs.
10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Section VI

Port Security

The chapters in this section explain the port security features of the AT-8524M switch The chapters include:

Chapter 28

MAC Address-based Port Security

This chapter explains how you can use the dynamic and static MAC addresses learned and assigned to the ports of the switch to control which end nodes can forward packets through the device. The sections in this chapter include:

Note
This type of port security does not apply to ports located on optional GBIC and SFP modules.

MAC Address-based Port Security Overview

This feature can enhance the security of your network. You can use it to control which end nodes can forward frames through the switch, and so prevent unauthorized individuals from accessing your network or particular parts of the network.

This type of network security uses a frame's source MAC address to determine whether the switch should forward a frame or discard it. The source address is the MAC address of the end node that sent the frame.

There are four levels of port security:

- Automatic
- Limited
- Secured
- Locked

You set port security on a per port basis. Only one security level can be active on a port at a time.

Automatic

The Automatic security mode disables MAC address port security on a port. This is the default security level for a port.

Limited

With this security level you specify the maximum number of dynamic MAC addresses a port can learn. Packets are forwarded based only on the learned addresses. Packets with an unlearned address are discarded.

When the Limited security mode is initially activated on a port, all dynamic MAC addresses learned by the port are deleted from the MAC address table and the port begins to learn new addresses, up to the maximum allowed. After the port has learned its maximum number, it does not learn any new addresses and forwards a packet only if its source address had already been learned.

A dynamic MAC address learned on a port operating in the Limited security mode never times out from the MAC address table, even when the corresponding end node is inactive. However, the addresses are not retained when a switch is reset or power cycled. Instead, a port repeats the MAC address learning process, up to the specified maximum.

Static MAC addresses are retained by the port and are not included in the count of maximum dynamic addresses. You can continue to add static MAC addresses to a port operating with this security level, even after the port has already learned its maximum number of dynamic MAC addresses. A switch port can have up to 255 dynamic and static MAC addresses.

Secured

A port set to this security level forwards packets only using static MAC addresses. The port does not learn dynamic MAC addresses and deletes any it has already learned. The port discards an ingress packet if its source MAC address is not specified as a static address.

You must enter, either before or after activating this security level on a port, the static MAC addresses of the source end nodes to be allowed to forward frames through the port.

Locked

A port set to the Locked security level stops learning new source dynamic MAC addresses and forwards packets based on the source MAC addresses that it learned before being set to this security level. Packets with a source MAC address that it did not learn are discarded.

The dynamic addresses learned by the port prior to its being set to this security level are converted into static addresses and so are never timed out from the MAC address table, even when the corresponding end nodes are inactive. They are also retained when the switch is reset or power cycled.

You can add new static MAC addresses to a port operating under this security level.

Invalid Frames and Intrusion Actions

When a port receives an invalid frame, it performs an action referred to as the intrusion action, which defines the action taken by the port and switch.

Before defining the intrusion actions, it can help to understand first what constitutes an invalid frame for each security level:

- Limited Security Level - An invalid frame for this security level is an ingress frame with a source MAC address that the port did not learn before reaching its maximum number of allowed dynamic MAC addresses, or that was not assigned to the port as a static address.
- Secured Security Level - An invalid frame for this security level is an ingress frame whose source MAC address was not entered as a static address on the port.
- Locked - An invalid frame for this security level is an ingress frame whose source MAC address was not learned before the port was set to this security level, or was not assigned to it as a static address.

Intrusion action defines the action of a port when it receives an invalid frame. There is only one intrusion action for a port operating under the Secured or Locked security mode. The action is to discard the invalid frame.

The Limited security mode lets you specify one of the following intrusion actions:

- Discard the invalid frame.

- ❒ Discard the invalid frame and send an SNMP trap. (SNMP must be enabled on the switch for the trap to be sent.)
- ❒ Discard the invalid frame, send an SNMP trap, and disable the port.

Guidelines

Here are the guidelines to observe when using this type of port security:

- ❒ This security method only applies to ingress packets on a port and not to egress packets.
- ❒ MAC address security can be set from a local, Telnet, or SSH management session, but not from a web browser management session.
- ❒ You cannot use MAC address security and 802.1x port-based access control on a port at the same time. To configure a port as an Authenticator or Supplicant in 802.1x port-based access control, you must set its MAC address security level to Automatic.
- ❒ This type of port security is not supported on optional GBIC and SFP modules.

Configuring MAC Address-based Port Security

To set the port security level, perform the following procedure:

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **5** to select Port Security.

   The Port Security menu is shown in Figure 218.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                      11:20:02 02-Jan-2006

                      Port Security

1 - Configure Port Security
2 - Display Port Security

R - Return to Previous Menu

Enter your selection?
```

Figure 218. Port Security Menu

3. Type **1** to select Configure Port Security.

   The following prompt is displayed:

   ```
   Enter Port-List:
   ```

4. Enter the port where you want to set port security. You can specify one port or a range of ports (for example, 4-8).

   The Configure Port Security menu is shown in Figure 219.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                      11:20:02 02-Jan-2006

               Configure Port Security

 Configuring Port Security 4
1 - Security Mode ..................... Automatic

D - Set Default Port Security
R - Return to Previous Menu

Enter your selection?
```

Figure 219. Configure Port Security Menu #1

The menu displays the current security level on the selected port. If you are configuring a range of ports and the ports have different security levels, the menu displays the current security level of the lowest number port.

Note
The **D** - Select Default Port Security option in the menu sets the security mode for the port to the default value of Automatic.

5. Press **1** to change the port security on your specified port list.

   The following prompt appears:

   ```
   Enter new mode (A-Automatic, L-Limited, S-Secured,
   K-locKed):
   ```

6. Select the desired security level. For definitions of the security levels, refer to “MAC Address-based Port Security Overview” on page 634.

7. Do one of the following:

   - If you selected Automatic, which disables port security on the port, no further steps are required. Return to the Main Menu to save your change.
   - If you configured a port for Secure security level, remember to enter the static MAC addresses of the end nodes that can send packets through the port. For instructions on how to add static MAC addresses, refer to “Adding Static Unicast and Multicast MAC Addresses” on page 130.
   - If you selected Locked, no further steps are required. Return to the Main Menu to save your change. You can, if desired, add static addresses to a port operating in the Locked security mode. For instructions, refer to “Adding Static Unicast and Multicast MAC Addresses” on page 130.
   - If you selected Limited, several new menu options are added to the Configure Port Security menu, as shown in Figure 220. Continue with Step 8 for instructions on configuring a port operating under the Limited security level.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch

User: Manager                      11:20:02 02-Jan-2006

               Configure Port Security

Configuring Port Security 4

1 - Security Mode ..................... Limited
2 - Threshold ......................... 100
3 - Intruder Action ................... Discard
4 - Port Participating ................ No

D - Set Default Port Security
R - Return to Previous Menu

Enter your selection?
```

Figure 220. Configure Port Security Menu #2

8. If you selected the Limited security mode for the port, do the following to specify the maximum number of dynamic MAC addresses you want the port to be able to learn:

   a. Type **2** to select Threshold.

      The following prompt appears:

      ```
      Enter port security threshold: [1 to 255] -> 100
      ```

   b. Enter the maximum number of dynamic MAC addresses you want the port to be able to learn. The range is 1 to 255. The default is 100.

9. To set the intrusion action for a port in the limited security mode, do the following:

   a. Type **3** to select Intruder Action.

      The following prompt is displayed:

      ```
      Enter intruder action: (N-Discard, T-Trap, D-
      Disable):
      ```

   b. Select the desired action:

      N - Discard: The port discards invalid frames. This is the default.

      T - Trap: The port discards invalid frames and sends an SNMP trap.

      D - Disable: The port discards invalid frames, sends a SNMP trap, and disables the port.

10. If you selected the trap or disable intrusion action, type **4** to toggle the Port Participating option to Yes.

    This option applies only when the intrusion action is set to trap or disable. This option does not apply when intrusion action is set to discard. If this option is set to No when intrusion action is set to trap or disable, the port discards invalid packets, but it does not send the SNMP trap or disable the port. If you want the switch to send a trap and/or disable the port, you must sent this option to Yes.

11. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying Port Security Levels

To view the current security levels for the ports on the switch, do the following:

1. From the Main Menu, type **1** to select Port Configuration.
2. From the Port Configuration menu, type **5** to select Port Security.

   The Port Security menu is shown in Figure 218 on page 637.
3. From the Port Security menu, type **2** to select Display Port Security.

   The Display Port Security menu is shown in Figure 221.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                           Production Switch

User: Manager                                        11:20:02 02-Jan-2006

                         Display Port Security

Port   Security Mode    Threshold     Intruder Action    Participating
-------------------------------------------------------------------
1      Limited             6          Trap                   Yes
2      Limited             10         Trap                   Yes
3      Automatic           ---        ------                 ---
4      Locked              ---        ------                 ---
5      Automatic           ---        ------                 ---
6      Automatic           ---        ------                 ---
7      Automatic           ---        ------                 ---
8      Secured             ---        ------                 ---

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 221. Display Port Security Menu

This menu is for viewing purposes only. The columns in the menu are defined below:

Port
The number of the port.

Security Mode
The active security mode on the port.

Threshold
This column specifies the maximum number of dynamic MAC addresses the port will learn. This applies when a port is operating in the Limited security mode.

Intruder Action
The column specifies the action taken by a port if it receives an invalid frame while operating in the Limited security mode.

- Discard: The port discards invalid frames. This is the default.
- Send Trap: The port discards invalid frames and sends a trap.
- Disable Port: The port discards invalid frames, sends a trap, and disables the port.

Note
Though not reflected in the Display Port Security menu, ports operating in the Secure or Locked security mode discard all invalid frames. For further information, refer to “Invalid Frames and Intrusion Actions” on page 635.

Participating
This column applies only when the intrusion action for a port operating in the Limited security mode is set to trap or disable. This option does not apply when intrusion action is set to discard. If this option is set to No when intrusion action is set to trap or disable, the port discards invalid packets, but it does not send a trap or disable the port.

Chapter 29

802.1x Port-based Network Access Control

This chapter explains 802.1x Port-based Network Access Control and how you can use this feature to restrict access to the network ports on the switch. Sections are as follows:

IEEE 802.1x Port-based Network Access Control Overview

The AT-S62 management software offers you several different methods for protecting your network and its resources from unauthorized access. For instance, Chapter 28, “MAC Address-based Port Security” on page 633, explains how to restrict network access using the source MAC addresses of the end nodes in your network.

This chapter explains yet another way. This method, referred to as 802.1x port-based network access control, uses the RADIUS protocol to control who can send traffic through and receive traffic from a switch port. When implemented, this security method does not allow an end node to send or receive traffic through a port until the user of the node has been authenticated by a RADIUS server.

The benefit of this type of network security is obvious. You can use it to prevent unauthorized individuals from connecting a computer to a switch port or using an unattended workstation to access your network resources. Only those users whom you have designated as valid network users on the RADIUS server will be permitted to use the switch to access the network.

This port security method uses the RADIUS authentication protocol. The AT-S62 management software is shipped with RADIUS client software. If you have already read Chapter 34, “TACACS+ and RADIUS Authentication Protocols” on page 747, then you know that you can use the RADIUS client software on the switch, along with a RADIUS server on your network, to also create new manager accounts that control who can manage and change the AT-S62 parameter on the switch.

Note
RADIUS with Extensible Authentication Protocol (EAP) extensions is the only supported authentication protocol for 802.1x Port-based Network Access Control. This feature is not supported with the TACACS+ authentication protocol. The switch only supports one authentication protocol at a time. Consequently, if you want to implement 802.1 Port-based Network Access Control and also create new manager accounts as explained in Chapter 34, “TACACS+ and RADIUS Authentication Protocols” on page 747, you must use the RADIUS protocol.

Following are several terms to keep in mind when you use 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 supplicant has been validated by the RADIUS server.
- Authentication server - The authentication server is the network device that has the RADIUS server software. This is the device that does the actual authenticating of the supplicants.

The AT-8500 Series switch does not authenticate any of the supplicants connected to its ports. It's function is to act 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 initiates an authentication message 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/MD5 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 message, 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.

Port Roles

Part of the task of implementing this feature is specifying the roles of the ports on the switch. A port can have one of three roles:

- None
- Authenticator
- Supplicant

None Role

A switch port in the None role does not participate in port-based access control. Any device can connect to the port and send traffic through it and receive traffic from it without being validated. This port setting is appropriate if no validation is required for the network device connected to the port. This is the default setting for the switch ports.

> **Note**
> Because a RADIUS authentication server cannot authenticate itself, it must communicate with the switch through a port that is set to the None role.

Authenticator Role

Placing a switch port in the authenticator role activates port access control on the port. A port in the role of authenticator does not forward network traffic to or from the end node until the client has been authenticated by a RADIUS server.

Determining whether a switch port should be set to the authenticator role is straightforward. You should set a port on a switch to the authenticator role if you want the user of the end node connected to the port to be authenticated before being permitted to use the network.

Authentication Modes

The AT-8500 Series switch supports two authentication modes on an authenticator port.

- 802.1x username/password combination

  In this authentication mode, each supplicant connected to an authenticator port must be assigned a unique username and password combination on the RADIUS server. A supplicant must provide the information either manually or automatically when initially passing traffic through an authenticator port and during reauthentications. The 802.1x client software on the supplicant either prompts the user for the necessary information or provides the information automatically.

  Assigning unique username and password combinations to your network users and requiring the users to provide the information when they initially send traffic through the switch can enhance network security by limiting network access to only those supplicants who have been assigned valid combinations. Another advantage is that the

authentication is not tied to any specific computer or node. An end user can log on from any system and still be verified by the RADIUS server as a valid user of the switch and network.

This authentication method requires 802.1x client software on the supplicant nodes.

- MAC address-based authentication

  An alternative method is to use the MAC address of a node as the username and password combination for the device. The client is not prompted for this information. Rather, the switch extracts the source MAC address from the initial frames received from a supplicant and automatically sends the MAC address as both the username and password of the supplicant to the RADIUS server for authentication.

  The advantage to this approach is that the supplicant need not have 802.1x client software. The disadvantage is that because the client is not prompted for a username and password combination, it does not guard against an unauthorized individual from gaining access to the network through an unattended network node or by counterfeiting a valid network MAC address.

Operational Settings

A port in the authenticator role can have one of three possible operational settings:

- Auto - Activates port-based authentication. The port begins in the unauthorized state, forwarding only EAPOL frames and discarding all other traffic. 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 messages between the client and the RADIUS authentication server. After the supplicant is validated by the RADIUS server, the port begins forwarding all traffic to and from the supplicant. This is the default setting for an authenticator port.
- Force-authorized - Disables IEEE 802.1X port-based authentication and automatically places the port in the authorized state without any authentication exchange required. The port transmits and receives normal traffic without authenticating the client.

Note
A supplicant connected to an authenticator port set to force-authorized must have 802.1x client software if the port is configured for the 802.1x authentication mode. Though this setting precludes an authentication exchange, the supplicant must still have the client software. Supplicants without 802.1 client software cannot forward traffic through an authenticator port set to force-authorized.

- Force-unauthorized - Causes the port to remain in the unauthorized state, ignoring all attempts by the supplicant to authenticate. The port forwards EAPOL frames, but discards all other traffic. This setting is analogous to disabling a port.

As mentioned earlier, the switch itself does not authenticate the user names and passwords from the clients. That function is performed by the authentication server, which contains the RADIUS server software. The switch acts as an intermediary for the authentication server by denying access to the network by the client until the client has been validated by the authentication server.

Supplicant Role

A switch port in the supplicant role acts as a client. The port assumes it must log in by providing a valid user name and password to whatever device it is connected to, typically another switch port.

Figure 222 illustrates the port role. Port 11 on switch B has been set to the supplicant role. Now, whenever switch B is power cycled or reset and initiates a link with switch A, it must log on by providing a username and password. (You enter this information when you configure the port for the supplicant role.)

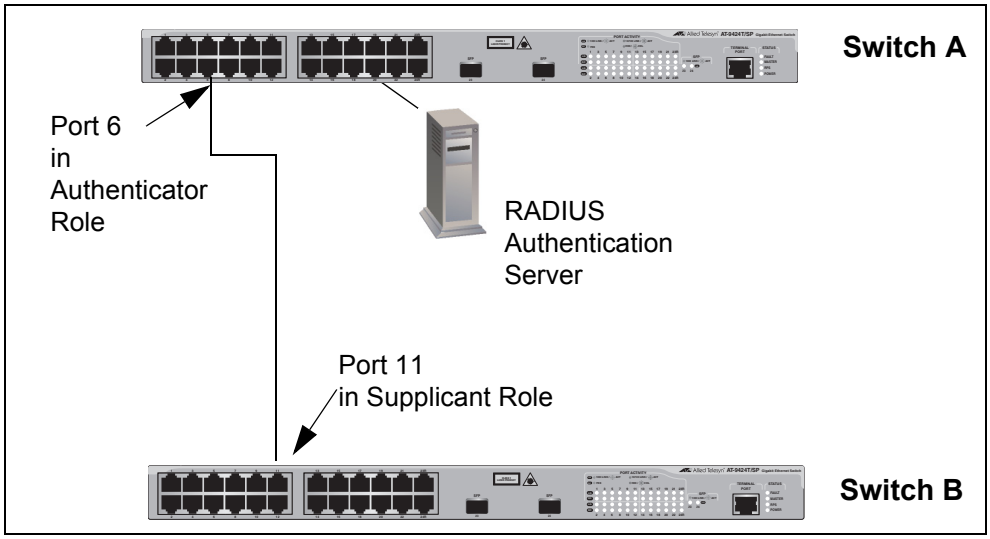

Figure 222. Example of the Supplicant Role

Authenticator Ports with Single and Multiple Supplicants

An authenticator port has two operating modes. The modes relate to the number of clients using the port and, in situations where an authenticator port is supporting more than one client, whether just one client or all the clients must log on to use the switch port.

The operating modes are:

- Single
- Multiple

Single Operating Mode

The Single operating mode is used in two situations. The first is when an authenticator port supports only one client. In this scenario, the switch allows only one client to log on and use the port.

You can also use the Single mode when an authenticator port supports more than one client, but where only one client needs to log on in order for all clients to use the port. This configuration can be useful in situations where you want to add 802.1x Port-based Network Access Control to a switch port that is supporting multiple clients, but want to avoid having to create individual accounts for all the clients on the RADIUS server.

This is referred to as "piggy-backing." After one client has successfully logged, the port permits the other clients to piggy-back onto the initial client's log on, allowing all clients to forward packets through the port.

To implement this configuration, you have to set the operating mode of an authenticator port to Single and also toggle the piggy-back mode feature. When piggy-back is disabled, only one client is allowed to log on and use the port. When this feature is enabled, an unlimited number of clients can use the port after one client has successfully logged on.

Note, however, that should the client who accomplished the initial log on fail to periodically reauthenticate or log out, the switch port reverts to the unauthenticated state. It bars all further traffic to and from all the clients on the port, until the initial client or another client logs on.

Here are several examples illustrating the Single operating mode and the piggy-back mode of an authenticator port. In Figure 223 on page 650, an authenticator port on a switch, in this case port 6, is connected to a single client. The authenticator port's operating mode is set to Single and the piggy-back feature is disabled so that only one client can use the port at any one time.

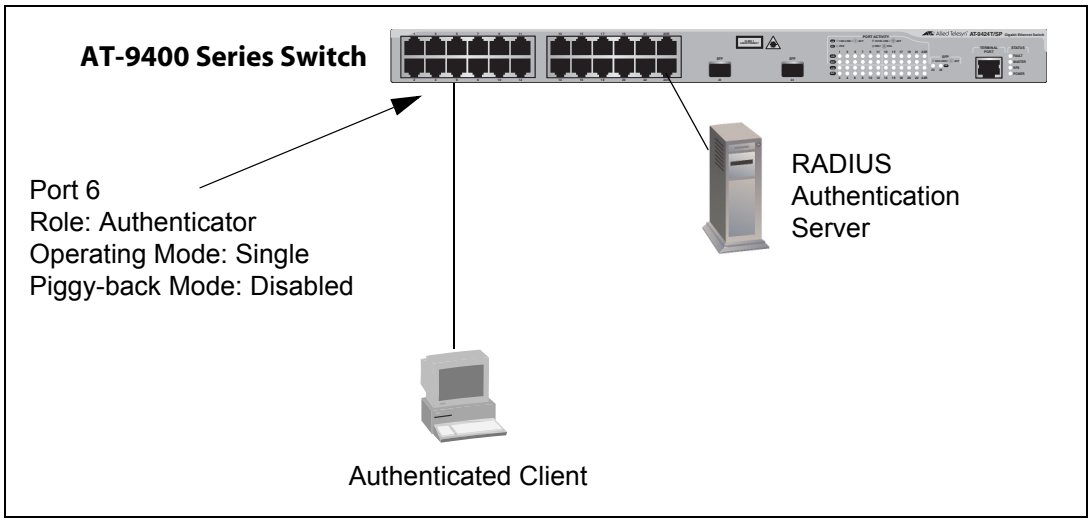

Figure 223. Authenticator Port in Single Operating Mode with a Single Client

The example in Figure 224 on page 651 illustrates a configuration where there are multiple clients connected to an authenticator port on the switch through an Ethernet hub or a non-802.1x-compliant Ethernet switch, such as an unmanaged switch. The operating mode of the authenticator port on the AT-8500 Series switch is set to Single and the piggy-back mode is enabled so that the port allows all clients to forward packets through it after one client logs on.

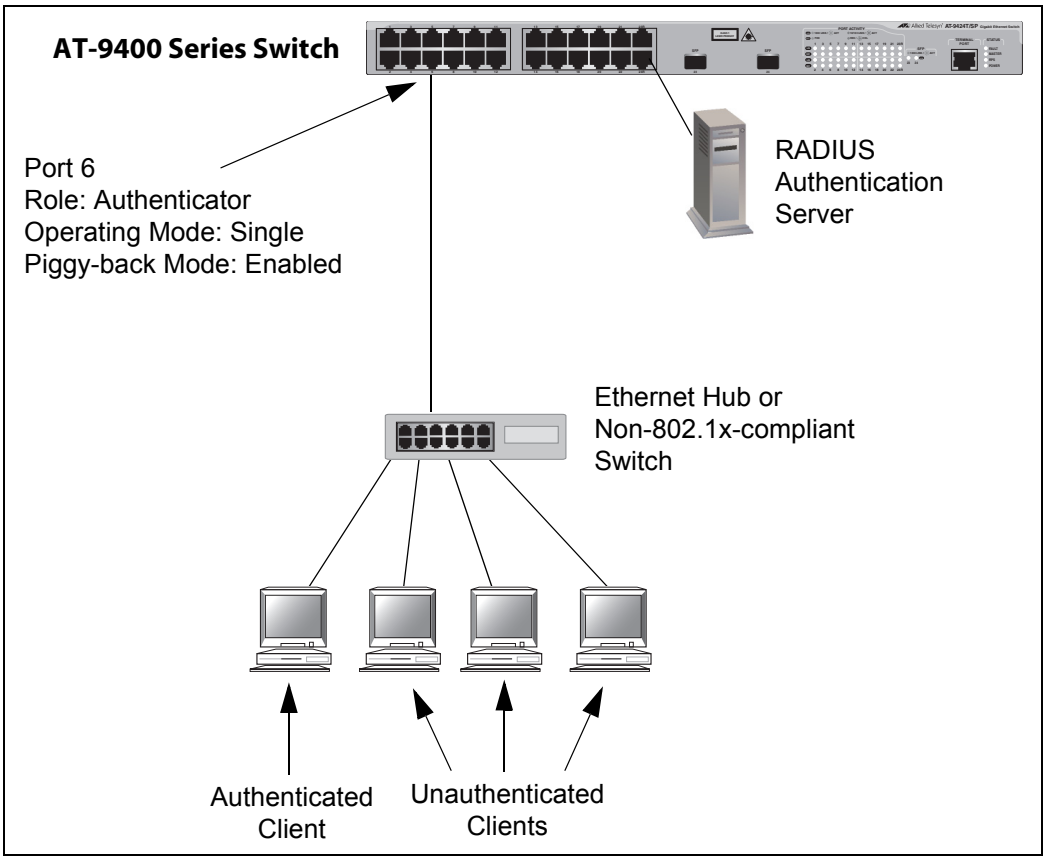

Figure 224. Single Operating Mode with Multiple Clients Using the Piggy-back Feature - Example 1

Because the piggy-back mode is activated on the authenticator port, only one client needs to be authenticated in order for all the clients to forward traffic through the port. If the port is using the 802.1x authentication method, then at least one client must have 802.1x client firmware and provide a username and password during authentication. (The other clients do not need 802.1x client firmware to forward traffic through the port after one client is authenticated.)

If the switch port is set to MAC address-based authentication, 802.1 client firmware is not required. The MAC address of the first client to forward traffic through the port is used for authentication. When that client is authenticated, all supplicants have access to the port.

As mentioned early, should the client who performed the initial log on fail to reauthenticate when necessary or log out, the port reverts to the unauthenticated state, blocking all traffic to and from all clients. Another client must be authenticated in order for all remaining clients to continue to forward traffic through the port.

If the clients are connected to an 802.1x-compliant device, such as another AT-8500 Series switch, you can automate the initial log on and reauthentications by configuring one of the switch ports as a supplicant. In this manner, the log on and reauthentications are performed automatically. eliminating the need for relying on an individual to perform the task. This scenario is illustrated in Figure 225.

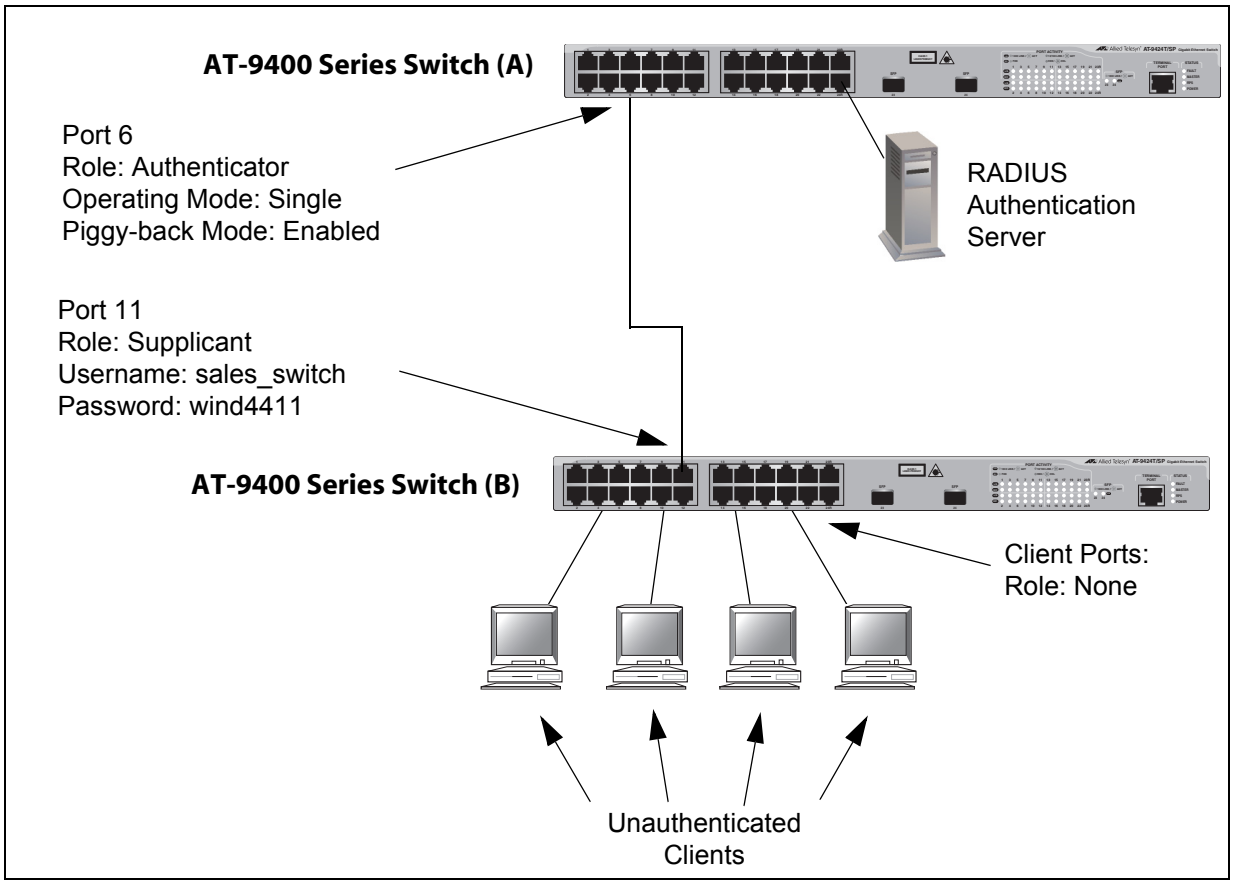

Figure 225. Single Operating Mode with Multiple Clients Using the Piggy-back Feature - Example 2

None of the workstations connected to switch B need to be authenticated or require 802.1x client software when accessing switch A because the log on to switch A and the subsequent reauthentications are performed automatically by the supplicant port on switch B, which is connected to an authenticator port on switch A with piggy-back mode enabled. It should be noted, however, that in this particular scenario the clients have full access to the resources of switch B even if the switch fails to log on or reauthenticate to switch A.

The example in the next figure again illustrates two 802.1x-compliant switches. The primary difference between this and the previous example is that the clients in the previous example did not have to log on to access switch B. In this example the clients have to log on to have any access at all to the network.

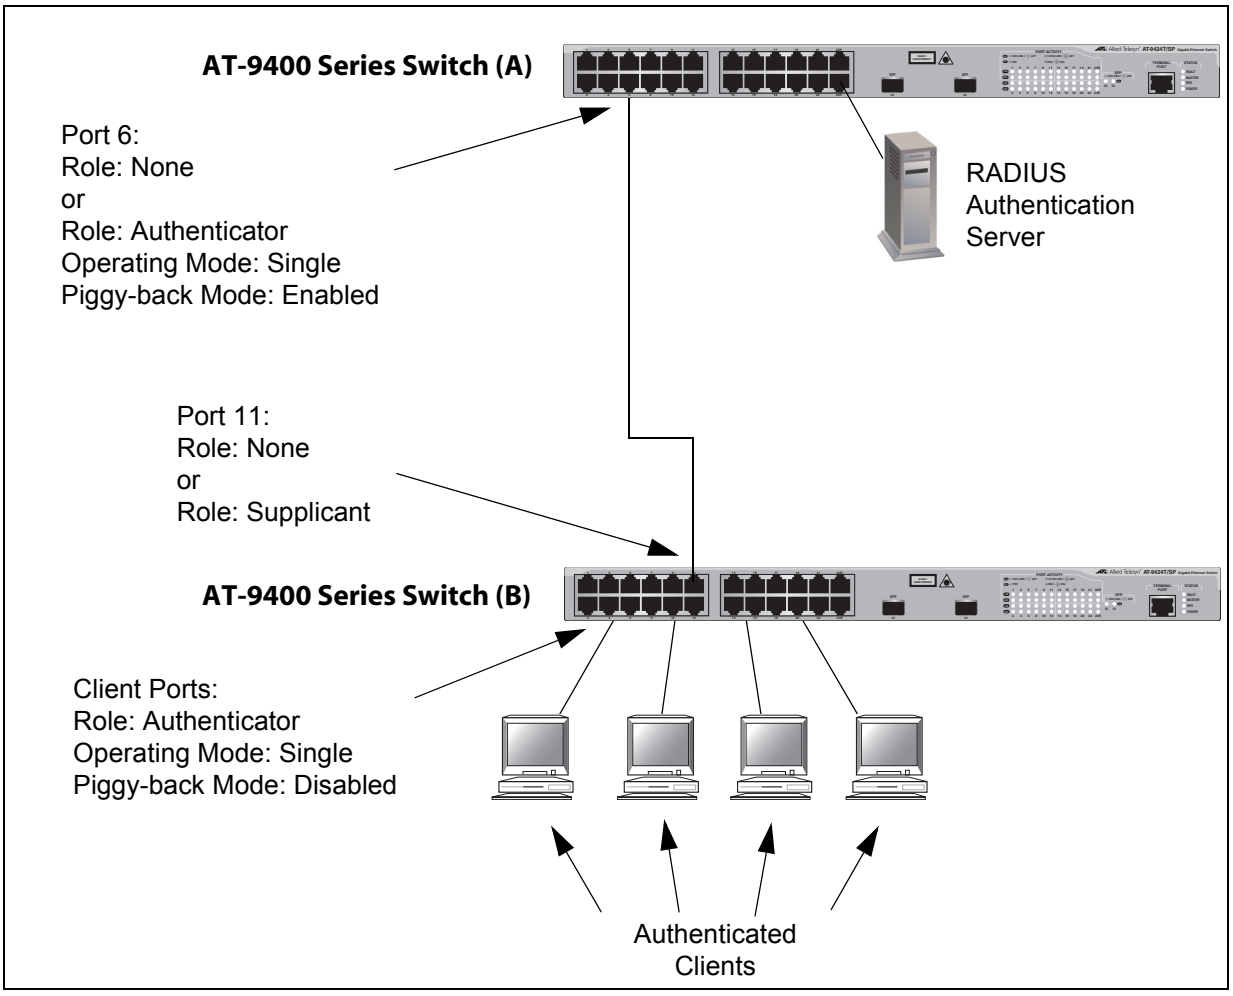

Figure 226. Single Operating Mode with Multiple Clients Using the Piggy-back Feature - Example 3

Multiple Operating Mode

The second type of operating mode for an authenticator port is the Multiple mode. You use this mode when a port is supporting more than one client and you want each client to log on individually before being permitted to use the port, perhaps as a means to increasing network security. An authenticator port in this mode can support up to a maximum of 20 clients, with a total maximum of 480 per switch. If you are using the 802.1x authentication method, you must provide each client with a separate username and password combination and the clients must provide their combinations to forward traffic through a switch port.

Selecting the Multiple mode for an authenticator port disables the piggy-back mode, because this operating mode does not permit piggy-backing.

An example of this authenticator operating mode is illustrated in Figure 227. The clients are connected to a hub or non-802.1x-compliant switch which is connected to an authenticator port on an AT-8500 Series switch. If the authenticator port is set to use the 802.1x authentication method, each client must be given a separate username and password combination to log on to and forward traffic through the AT-8500 Series switch. If the authentication method is MAC address-based, the authenticator port uses the MAC addresses of the clients as the username and password combinations. The port accepts and forwards traffic only from those clients whose MAC addresses have been entered on the RADIUS server and denies access to all other users.

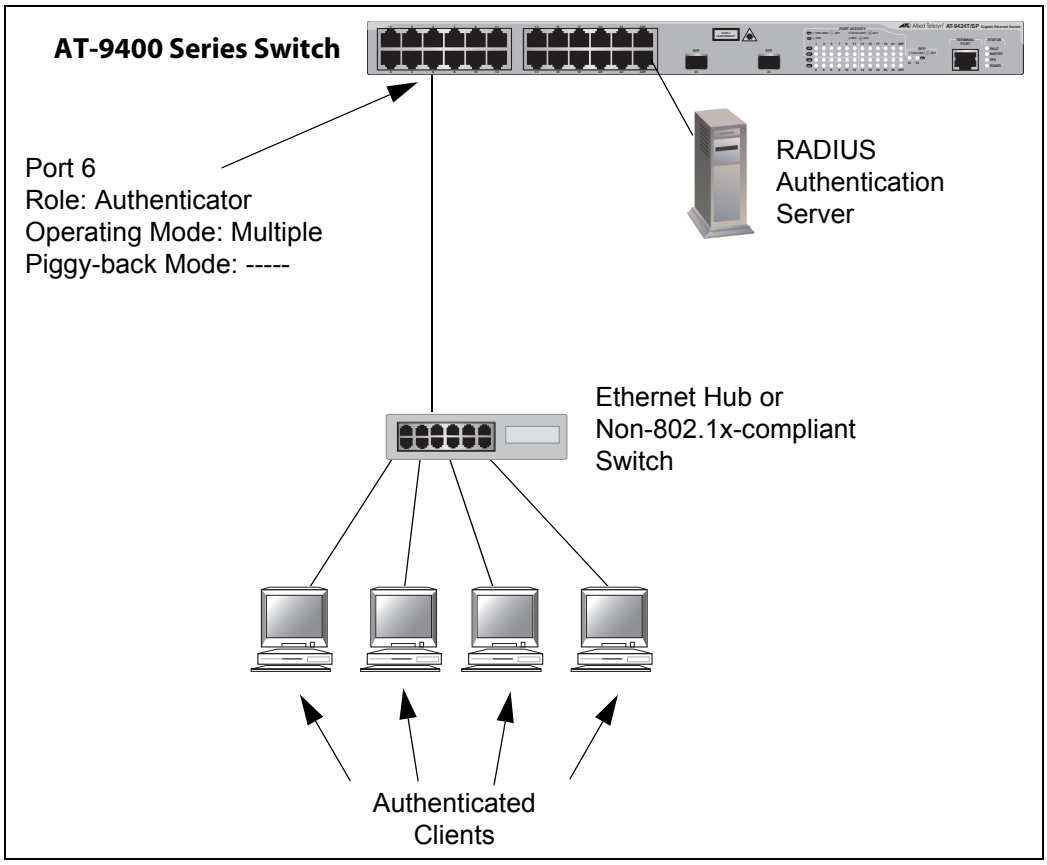

Figure 227. Authenticator Port in Multiple Operating Mode - Example 1

The next example of the multiple mode in Figure 228 shows two AT-8500 Series switches. The clients connected to switch B have to log on to port 6 on Switch A when they pass a packet to that switch for the first time.

There are several items to note when interconnecting two 802.1x-compliant devices using the Multiple operating mode of an authenticator port. In order for switch B in our example to pass the RADIUS messages to switch A, it must be able to log on to port 6 on switch A. That is why port 11 on the lower switch is configured as a supplicant. If its role is set to none, port 6 on switch A will discard the packets because switch B would

not be logged on to the port.

Also note that the ports where the clients are connected on switch B are set to the none role. This is because a client can log on only once. If, in this example, you were to make a client's port an authenticator, the client would have to log on twice when trying to access switch A, once on its port on switch B as well as the authenticator port on switch A. This is not permitted. Consequently, in our example the clients on switch B have full access to that switch, but are denied access to switch A until they log on to port 6 on switch A.

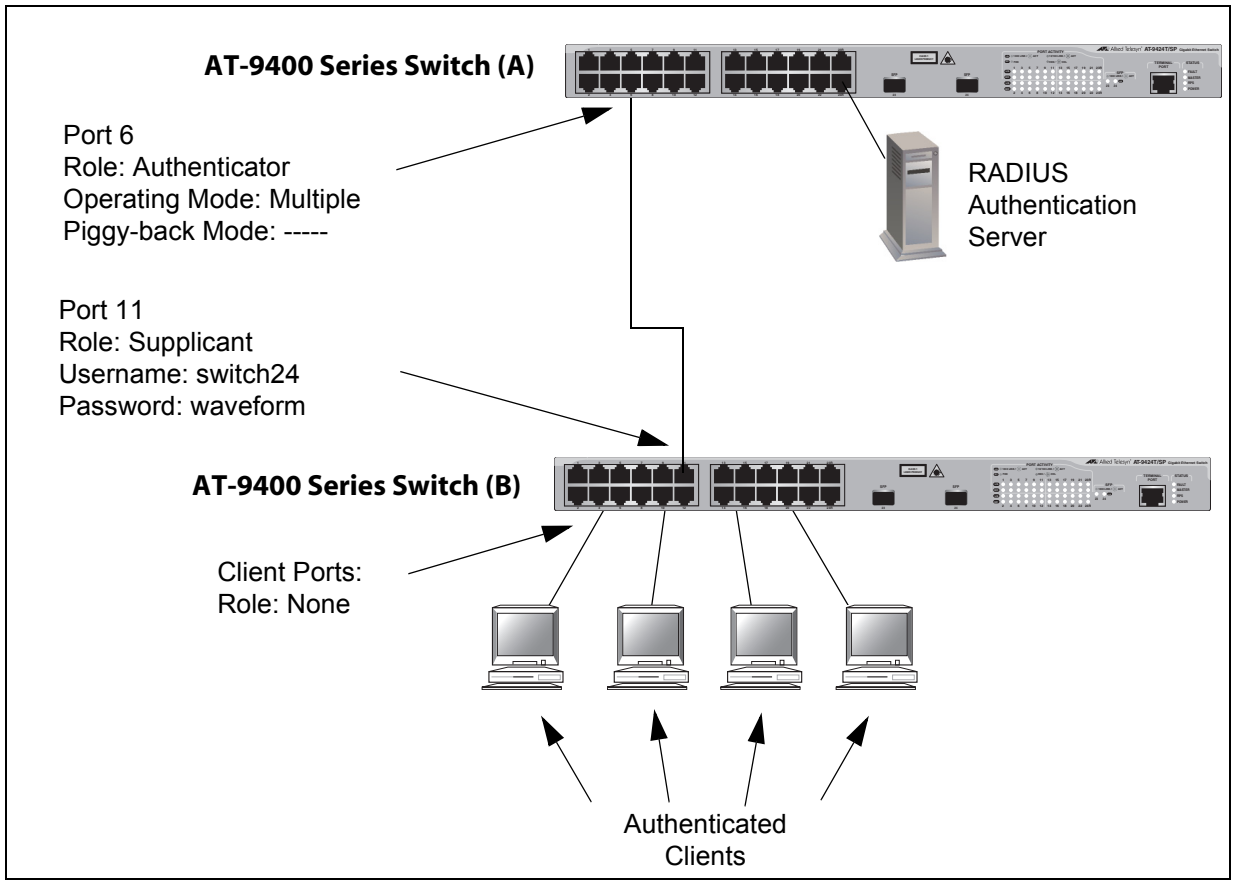

Figure 228. Authenticator Port in Multiple Operating Mode - Example 2

Supplicant and VLAN Associations

One of the challenges to managing a network is accommodating end users that roam. These are individuals whose work requires that they access the network resources from different points at different times. The difficulty arises in providing them with access to the same network resources and, conversely, restricting them from unauthorized areas, regardless of the workstation from where they access the network. A closely related issue is where a workstation is employed at various times by different individuals with unique requirements in terms of network resources and security levels.

Providing network users with access to their network resources while also maintaining network security is often achieved through the use of VLANs. As explained in “VLAN Overview” on page 546, a VLAN is an independent traffic domain where the traffic generated by the nodes within the VLAN is restricted to nodes of the same VLAN, unless there is a router or Layer 3 interconnection device. Different users are assigned to different VLANs depending on their resource requirements and security level.

The problem with a port-based VLAN is that VLAN membership is determined by the port on the switch to which the device is connected. If a different device that needs to belong to a different VLAN is connected to a port, it must be manually moved to the new VLAN using the management software.

With 802.1x port-based network access control, you can link a username and password combination or MAC address to a specific VLAN so that the switch automatically moves the port to the appropriate VLAN when a client logs on. This frees the network manager from having to reconfigure VLANs as end users access the network from different points or where the same workstation is used by different individuals at different times.

To use this feature, you have to enter a VLAN identifier, along with other information, when you create a user account on the RADIUS server. The server passes the identifier to the switch when a user logs on with a valid username and password combination or MAC address, depending on the authentication method. The information to provide on the RADIUS server is outlined in “Supplicant VLAN Attributes on the RADIUS Server” on page 657.

How the switch responses when it receives VLAN information during the authentication process can differ depending on the operating mode of the authenticator port.

Single Operating Mode

Here are the operating characteristics for the switch when an authenticator port is set to the Single operating mode:

- If the switch receives a valid VLAN ID or VLAN name from the RADIUS server, it moves the authenticator port to the designated VLAN and changes the port to the authorized state. If the piggy-back mode is disabled, only the authenticated supplicant is allowed to use the port. All other supplicants are denied entry. If the piggy-back mode is enabled, all clients are allowed access to the port and the same VLAN after the initial authentication.
- If the switch receives an invalid VLAN ID or VLAN name from the RADIUS server (e.g., the VID of a nonexistent VLAN), it leaves the port in the unauthorized state to deny access to the port.

Multiple Operating Mode

The initial authentication on an authenticator port running in the Multiple operating mode is handled in the same fashion as with the Single operating mode. If the switch receives a valid VLAN ID or name from the RADIUS server, it moves the authenticator port to the designated VLAN and changes the port to the authorized state.

How the switch handles subsequent authentications on the same port depends on how you set the Secure VLAN parameter. Your options are as follows:

- If you activate the Secure VLAN feature, only those supplicants with the same VLAN assignment as the initial supplicant are authenticated. Supplicants with a different VLAN assignment or with no VLAN assignment are denied access to the port.
- If you disable the Secure VLAN feature, all supplicants, regardless of their assigned VLANs, are authenticated. However, the port remains in the VLAN specified in the initial authentication.

Supplicant VLAN Attributes on the RADIUS Server

Here is the information that you need to configure on the RADIUS in order to associate a VLAN to a supplicant.

- Tunnel-Type
  The protocol to be used by the tunnel specified by Tunnel-Private-Group-Id. The only supported value is VLAN (13).
- Tunnel-Medium-Type
  The transport medium to be used for the tunnel specified by Tunnel-Private-Group-Id. The only supported value is 802 (6).
- Tunnel-Private-Group-ID
  The ID of the tunnel the authenticated user should use. This must be the name of VID of the VLAN of the switch.

Guest VLAN

An authenticator port in the unauthorized state typically accepts and transmits only 802.1x packets while waiting for an supplicant to be authenticated. 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 is automatically returned to the Guest VLAN.

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

RADIUS Accounting

The AT-S62 management software supports RADIUS accounting for switch ports set to the Authenticator role. This feature sends information to the RADIUS server about the status of its supplicants. You can view this information on the RADIUS server to monitor network activity and use.

The switch sends accounting information to the RADIUS server when one of the following events occur:

- Supplicant logs on
- Supplicant logs off
- A change in the status of an Authenticator port during an active Supplicant session (for example, the port is reset or is changed from the Authenticator role to None role while a Supplicant is logged on)

The information sent by the switch to the RADIUS server for an event includes:

- Port number where the event occurred
- The date and time when the event occurred
- The number of packets transmitted and received by the switch port during a supplicant's session. (This information is sent only when the client logs off.)

You can also configure the accounting feature to send interim updates so you can monitor which clients are still active.

Here are the guidelines to using the accounting feature:

- The AT-S62 management software supports the Network level of accounting, but not the System or Exec.
- This feature is only available for ports operating in the Authenticator role. No accounting is provided for ports operating in the Supplicant or None role.
- You must configure 802.1x Port-based Network Access Control as explained in this chapter and designate the Authenticator ports.
- You must also specify from one to three RADIUS servers. The instructions for this are in "Configuring RADIUS Authentication Protocol Settings" on page 755.

For instructions on configuring this feature, refer to "Configuring RADIUS Accounting" on page 676.

General Steps

Following are the general steps to implementing 802.1x Port-based Network Access Control and RADIUS accounting on the switch:

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 Telesyn. Funk Software Steel-Belted Radius and Free Radius have been verified as fully compatible with the AT-S62 management software.

   Note
   This feature is not supported with the TACACS+ authentication protocol.

2. Those clients connected to an authenticator port set to the 802.1x authentication method will need 802.1x client software. Microsoft WinXP client software and Meeting House Aegis client software have been verified as fully compatible with the AT-S62 management software. (802.1x client software is not required when an authenticator port is set to the MAC address-based authentication method.)

3. You must configure and activate the RADIUS client software in the AT-S62 management software. The default setting for the authentication protocol is disabled. You will need to provide the following information:

   - The IP addresses of up to three RADIUS servers.
   - The encryption key used by the authentication servers.

   The instructions for this step are in "Configuring RADIUS Authentication Protocol Settings" on page 755.

4. You must configure the port access control settings on the switch. This involves the following:

   - Specifying the port roles.
   - Configuring 802.1x port parameters.
   - Enabling 802.1x Port-based Network Access Control.

   The instructions for this step are found in this chapter.

5. If you want to use RADIUS accounting to monitor the clients connected to the switch ports, you must configure the service on the switch, as explained in "Configuring RADIUS Accounting" on page 676.

802.1x Port-based Network Access Control Guidelines

The following are general guidelines to using this feature:

- Ports operating under port-based access control do not support dynamic MAC address learning.
- The appropriate port role for a port on an AT-8500 Series switch connected to a RADIUS authentication server is None.
- The authentication server must be a member of the management VLAN. For information about the management VLAN, refer to “Specifying a Management VLAN” on page 579.
- The authentication method of an authenticator port can be either 802.1x username and password combination or MAC address-based, but not both.
- A supplicant must have 802.1x client software if the authentication method of a switch port is 802.1x username and password combination.
- A supplicant does not need 802.1x client software if the authentication method of an authenticator port is MAC address-based.
- An authenticator port set to the multiple operating mode can handle up to a maximum of 20 authenticated supplicants at one time.
- The switch can handle up to a maximum of 480 authenticated supplicants at one time. The switch stops accepting new authentications after the maximum is reached and starts accepting new authentications as supplicants log out or are timed out.
- An 802.1x 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 client 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 client logs off the network or fails to reauthenticate, at which point the address is removed. The address is not timed out, even when the node is inactive.

  Note
  End users of 802.1x Port-based Network Access Control should be instructed to always log off when they are finished with a work session. This can prevent an unauthorized individual from accessing the network through an unattended network workstation.

- Authenticator and supplicant ports must be untagged ports. They cannot be tagged ports of any VLAN.
- The MAC address-based port security setting for an authenticator port must be Automatic. This restriction does not apply to a supplicant port. For further information, refer to Chapter 28, “MAC Address-based Port Security” on page 633.

- An authenticator port cannot be part of a static port trunk, LACP port trunk, or port mirror.
- If a switch port set to the supplicant role is connected to a port on another switch that is not set to the authenticator role, the port, after a timeout period, assumes that it can send traffic without having to log on.
- GVRP must be disabled on an authenticator port.
- When 802.1x Port-based Network Access Control is activated on a switch, the feature polls all RADIUS servers specified in the RADIUS configuration. If three servers have been configured, the switch polls all three. If server 1 responds, all future requests go only to that server. If server 1 stops responding, the switch again polls all RADIUS servers. If server 2 responds, but not server 1, then all future requests go to servers 1 and 2. If only server 3 responds, then all future requests go to all three servers.
- In order to change the untagged VLAN assignment of an authenticator or supplicant port, you must first set its port role to none. After the VLAN assignment is made, you can change the port's role back again to authenticator or supplicant, if needed.
- A Guest VLAN can be either port-based or tagged.
- To use the Guest VLAN feature, the designated VLAN must already exist on the switch.
- The switch must be running in the user-configured VLAN mode to support 802.1x port-based network access control. The feature is not supported when the switch is running in a multiple VLAN mode. For further information, refer to "Selecting a VLAN Mode" on page 612.
- The AT-S62 management software only supports EAP-MD5 authentication for both authenticators and supplicants.

Here are guidelines that apply when adding VLAN assignments to supplicant accounts on a RADIUS server:

- The VLAN can be either port-based or tagged.
- The VLAN must already exist on the switch.
- A client can have only one VLAN associated with it on the RADIUS server.
- When a supplicant logs on, the switch port is moved as an untagged port to the designated VLAN.

Setting Port Roles

This procedure sets port roles. For an explanation of port roles, refer to “Port Roles” on page 646. You must set up the port roles before you enable port access control.

To set port roles, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **2** to select Port Access Control (802.1X).

   The Port Access Control (802.1X) menu is shown in Figure 229.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Marketing
User: Manager                            11:20:02 02-Mar-2006
                Port Access Control (802.1X)

1 - Port Access Control ............... Enabled
2 - Authentication Method ............. RADIUS EAP
3 - Configure Port Access Role
4 - Configure Authenticator
5 - Configure Supplicant
6 - Display Port Access Status
7 - Configure Accounting

R - Return to Previous Menu

Enter your selection?
```

Figure 229. Port Access Control (802.1X) Menu

3. From the Port Access Control menu, type **3** to select Configure Port Access Role.

   The following prompt is displayed:

   ```
   Enter port list ->
   ```

4. Enter the port whose role you want to change. You can specify one port or multiple ports.

The Configure Port Access Role menu is shown in Figure 230.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Marketing
User: Manager                            11:20:02 02-Mar-2006
                Configure Port Access Role

Configuring Port 3
1 - Port Role ......... None

R - Return to Previous Menu

Enter your selection?
```

Figure 230. Configure Port Access Role Menu

5. Type **1** to select Port Role.

   The following prompt is displayed:

   ```
   Enter new Port Role [N-None, A-Authenticator,
   S-Supplicant] ->
   ```

6. If you type **N** for None, the port does not participate in port access control. This is the default setting. If the port is connected to a supplicant, type **A** to set the port's role to Authenticator. If the port is connected to an authenticator, type **S** to set the port's roles to Supplicant.

7. Repeat this procedure starting with Step 3 to configure the role of the other ports on the switch.

   After you have set port roles, go to "Configuring Authenticator Port Parameters" on page 665 and "Configuring Supplicant Port Parameters" on page 671 to configure the port settings.

Enabling and Disabling 802.1x Port-based Network Access Control

This procedure explains how to enable and disable port-based access control on the switch. If you have not assigned port roles and configured the parameter settings, you should skip this procedure and go first to “Setting Port Roles” on page 662. To configure the port settings, refer to “Configuring Authenticator Port Parameters” on page 665 and “Configuring Supplicant Port Parameters” on page 671.

To enable or disable 802.1x Port-based Network Access Control, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **2** to select Port Access Control (802.1X).

   The Port Access Control (802.1X) menu is shown in Figure 229 on page 662.

3. From the Port Access Control menu, type **1** to select Port Access Control.

   The following prompt is displayed:

   ```
   Port Access Control (E-Enable, D-Disable):
   ```

4. Type **E** to enable port access control, or **D** to disable port access control.

   A change to the status of 802.1x port-based access control is immediately implemented on the switch.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring Authenticator Port Parameters

Note
A port must already be set to the authenticator role before you can configure its settings. For instructions on how to change the role of a port, refer to "Setting Port Roles" on page 662.

To configure the parameters of an authenticator port, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **2** to select Port Access Control (802.1X).

   The Port Access Control (802.1X) menu is shown in Figure 229 on page 662.

3. From the Port Access Control menu, type **4** to select Configure Authenticator.

   The Configure Authenticator menu is shown in Figure 231.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Marketing
User: Manager                              11:20:02 02-Mar-2006
                   Configure Authenticator

 1 - Configure Authenticator Port Access Parameters
 2 - Display Authenticator Port Access Parameters

 R - Return to Previous Menu

 Enter your selection?
```

Figure 231. Configure Authenticator Menu

4. From the Configure Authenticator menu, type **1** to select Configure Authenticator Port Access Parameters.

   The following prompt is displayed:

   ```
   Enter port list ->
   ```

5. Enter the authenticator port number whose parameters you want to change. You can configure more than one port at a time.

The Configure Authenticator Port Access Parameters menu is shown in Figure 232.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                         Marketing
User: Manager                              11:20:02 02-Mar-2006
      Configure Authenticator Port Access Parameters

 Configuring Port 3
0 - Authentication Mode ..... 802.1x
1 - Supplicant Mode ......... Single
2 - Port Control ............ Auto
3 - Quiet Period ............ 60 Seconds
4 - TX Period ............... 30 Seconds
5 - Reauth Enabled .......... Enabled
6 - Reauth Period ........... 3600 Seconds
7 - Supplicant Timeout ...... 30 Seconds
8 - Server Timeout .......... 30 Seconds
9 - Max Requests ............ 2
A - VLAN Assignment ......... Enabled
B - Secure VLAN ............. On
C - Control Direction ....... Both
D - Piggyback Mode .......... Disabled
E - Guest VLAN .............. None

R - Return to Previous Menu

Enter your selection?
```

Figure 232. Configure Authenticator Port Access Parameters Menu

6. Adjust the following parameters as necessary.

   0 - Authentication Mode
   This parameter can take the following values on an authenticator port:

   - **802.1x**: Specifies 802.1x username and password authentication. With this authentication method the supplicant must provide, either manually or automatically, a username and password. This authentication method requires 802.1x client software on the supplicant nodes.
   - **MAC based**: Specifies MAC address-based authentication. The authenticator port extracts the source MAC address from the initial frames received from a supplicant and automatically sends the address as both the username and password of the supplicant to the authentication server. Supplicant nodes must have 802.1x client software for this authentication method.

   For further information, refer to “Authentication Modes” on page 646.

1 - Supplicant Mode
This parameter can take the following values on an authenticator port:

- **Single**: Configures the authenticator port to accept only one authentication. This supplicant mode should be used together with the piggy-back mode. When an authenticator port is set to the Single mode and the piggy-back mode is disabled, only the one client who is authenticated can use the port. Packets from or to other clients on the port are discarded. If piggy-back mode is enabled, other clients can piggy-back onto another client's authentication and so be able to use the port.
- **Multiple**: Configures the authenticator port to accept up to 20 authentications. Every client using an authenticator port in this mode must have a username and password combination.

For addition information, refer to "Authenticator Ports with Single and Multiple Supplicants" on page 649.

2 - Port Control
The possible settings for this parameter 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 messages 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. This is the default setting.

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.

Note
A supplicant connected to an authenticator port set to force-authorized must have 802.1x client software if the port is configured for the 802.1x authentication mode. Though this setting precludes an authentication exchange, the supplicant must still have the client software. Supplicants without 802.1 client software cannot forward traffic through an authenticator port set to force-authorized.

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

3 - Quiet Period
The quiet period is 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.

4 - TX Period
This parameter 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.

5 - Reauth Enabled
Specifies if reauthentication should occur according to the reauthentication period. The options are Enabled or Disabled. If disabled, the supplicant is not require to reauthenticate after the initial authentication.

6 - Reauth Period
Specifies the time period in seconds between reauthentications of the client when the Reauth. Enabled option is set to Enabled. The default value is 3600 seconds. The range is 1 to 65,535 seconds.

7 - Supplicant Timeout
This parameter sets the switch-to-client retransmission time for the EAP-request frame. The default value for this parameter is 30 seconds. The range is 1 to 600 seconds.

8 - Server Timeout
This parameter sets the timer used by the switch to determine authentication server timeout conditions. The default value for this parameter is 30 seconds. The range is 1 to 65,535 seconds.

9 - Max Requests
This parameter specifies 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.

A - VLAN Assignment
This parameter controls whether an authenticator port uses the VLAN assignments returned by a RADIUS server. Options are:

- **Enabled**: Specifies that the authenticator port is to use the VLAN assignment returned by the RADIUS server when a supplicant logs on. This is the default setting. The port automatically moves to the designated VLAN after the supplicant successfully logs on.
- **Disabled**: Specifies that the authenticator port ignore any VLAN assignment information returned by the RADIUS server when a supplicant logs on. The authenticator port remains in its predefined VLAN assignment even if the RADIUS server returns a VLAN assignment when a supplicant logs on. This is the default setting.

For additional information, refer to "Supplicant and VLAN Associations" on page 655.

B - Secure VLAN
This parameter controls the action of an authenticator port to subsequent authentications after the initial authentication where VLAN assignments have been added to the user accounts on the RADIUS server. This parameter only applies when the port is operating in the Multiple operating mode. Possible settings are:

- **On**: Specifies that only those supplicants with the same VLAN assignment as the initial supplicant are authenticated. Supplicants with a different or no VLAN assignment are denied entry to the port. This is the default setting.
- **Off**: Specifies that all supplicants, regardless of their assigned VLANs, are authenticated. However, the port remains in the VLAN specified in the initial authentication, regardless of the VLAN assignments of subsequent authentications.

For further information, refer to "Supplicant and VLAN Associations" on page 655.

C - Control Direction
This parameter specifies how the port handles ingress and egress broadcast and multicast packets when in the unauthorized state. When a port is set to the authenticator role, it remains in the unauthorized state until a client logs on by providing a username and password combination. In the unauthorized state, the port only accepts EAP packets from the client. All other ingress packets that the port might receive from the client, including multicast and broadcast traffic, is discarded until the supplicant has logged in. The options are:

- **Ingress**: A port, when in the unauthorized state, discards all ingress broadcast and multicast packets from the client, but forwards all egress broadcast and multicast traffic to the same client.
- **Both**: A port, when in the unauthorized state, does not forward ingress or egress broadcast and multicast packets from or to the same client until the client logs in. This is the default.

Note
This parameter is only available when the authenticator's mode is set to Single. When set to Multiple, a port does not forward ingress or egress broadcast or multicast packets until at least one client has logged on.

D - Piggyback Mode
This parameter controls who can use the switch port in cases where there are multiple clients using the port (e.g., the switch port is connected to an Ethernet hub). If set to enabled, the port allows all clients on the port to piggy-back onto the initial client's authentication, forwarding all packets after one client is authenticated. If set to Disabled, the switch port forwards only those packets from the client who is authenticated and discards packets from all other users.

Note
This parameter is only available when the authenticator's mode is set to Single. For further information, refer to "Authenticator Ports with Single and Multiple Supplicants" on page 649.

E - Guest VLAN
This parameter specifies the name or VID of a Guest VLAN. The authenticator port is a member of a Guest VLAN when no supplicant is logged on. Clients do not log on to access a Guest VLAN. To remove a Guest VLAN without assigning a new one, enter "none".

For further information, refer to "Guest VLAN" on page 657.

7. Repeat this procedure starting with Step 4 to configure additional authenticator ports on the switch.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring Supplicant Port Parameters

To configure supplicant port parameters, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **2** to select Port Access Control (802.1X).

   The Port Access Control (802.1X) menu is shown in Figure 229 on page 662.

3. From the Port Access Control menu, type **5** to select Configure Supplicant.

   The Configure Supplicant menu is shown in Figure 231.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Marketing
User: Manager                              11:20:02 02-Mar-2006
                     Configure Supplicant

 1 - Configure Supplicant Port Access Parameters
 2 - Display Supplicant Port Access Parameters

 R - Return to Previous Menu

 Enter your selection?
```

Figure 233. Configure Supplicant Menu

4. From the Configure Supplicant menu, type **1** to select Configure Authenticator Port Access Parameters.

   The following prompt is displayed:

   ```
   Enter port list ->
   ```

5. Enter the supplicant port number whose parameters you want to change. You can specify one port or multiple ports.

   Note
   A port must already be configured as an supplicant before you can configure its settings. For instructions on how to change the role of a port, refer to "Setting Port Roles" on page 662.

The Configure Supplicant Port Access Parameters menu is shown in Figure 232.

```
    Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                      Marketing
User: Manager                          11:20:02 02-Mar-2006
       Configure Supplicant Port Access Parameters

 Configuring Port 5-8
 1 - Auth Period........... 30 Seconds
 2 - Held Period........... 60 Seconds
 3 - Max Start  ........... 3
 4 - Start Period.......... 30 Seconds
 5 - User Name: ...........
 6 - User Password: .......

 R - Return to Previous Menu

 Enter your selection?
```

Figure 234. Configure Supplicant Port Access Parameters Menu

6. Adjust the following parameters as necessary.

   1 - Auth Period
   This parameter specifies the period of time in seconds that the supplicant waits for a reply from the authenticator after sending an EAP-Response frame. The range is 1 to 60 seconds. The default is 30 seconds.

   2 - Held Period
   The held period specifies the amount of time in seconds the supplicant is to refrain from retrying to re-contact the authenticator in the event the end user provides an invalid username and/or password. After the time period has expired, the supplicant can attempt to log on again. The range is 0 to 65,535. The default value is 60.

   3 - Max Start
   Max start is the maximum number of times the supplicant sends EAPOL-Start frames before assuming that there is no authenticator present. The range is 1 to 10. The default is 3.

   4 - Start Period
   The start period is the time period in seconds between successive attempts by the supplicant to establish contact with an authenticator when there is no reply. The range is 1 to 60. The default is 30.

   5 - User Name
   The user name is the username for the switch port. The port sends the name to the authentication server for verification when the port logs on to the network. The username can be from 1 to 16 alphanumeric characters (A to Z, a to z, 1 to 9). Do not use spaces or special

characters, such as asterisks or exclamation points. The username is case sensitive.

6 - User Password
This parameter specifies the password for the switch port. The port sends the password to the authentication server for verification when the port logs on to the network. The password can be from 1 to 16 alphanumeric characters (A to Z, a to z, 1 to 9). Do not use spaces or special characters, such as asterisks or exclamation points. The password is case sensitive.

7. Repeat this procedure starting with Step 4 to configure additional supplicant ports on the switch.

8. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying the Port Access Parameters

To display the port access parameters for the ports on the switch, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **2** to select Port Access Control (802.1X).

   The Port Access Control (802.1X) menu is shown in Figure 229 on page 662.

3. From the Port Access Control menu, type **6** to select Display Port Access status.

   The Display Port Access Status menu is shown in Figure 235.

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                              Marketing
User: Manager                                      11:20:02 02-Mar-2006
                        Display Port Access Status

Port  PortRole        AuthMode      State           Additional Info
-------------------------------------------------------------------------
1     None                          ------          ------------------
2     Authenticator   802.1x        Connecting      ------------------
3     Authenticator   802.1x        Authenticated   00:a0:d2:18:1a:c8
4     Authenticator   MAC Based     Connecting      ------------------
5     None                          ------          ------------------
6     None                          ------          ------------------
7     None                          ------          ------------------
8     Supplicant                    Disabled        ------------------

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 235. Display Port Access Status Menu

The Display Port Access Status menu displays a table that contains the following columns of information:

Port
Port number.

AuthMode
The port's authentication mode: 802.1x or MAC Based. For further information, refer to "Authentication Modes" on page 646.

Port Role
Port access role configured for the port. The possible settings are None, Authenticator, or Supplicant.

State
State of the port. The state field is dependent on whether a port is configured as an authenticator or a supplicant.

When you configure a port with an Authenticator Role, the State field can have the following values:

Aborting
Authenticated
Authenticating
Connecting
Disconnected
Force_Auth
Force_Unauth
Held
Initialize

When you configure a port with a Supplicant role, the State field can have the following values:

Acquired
Authenticated
Authenticating
Connecting
Disconnected
Held
Logoff

Additional Info
When you assign a port the role of Authenticator and it has a status of Authenticated, this field also displays the MAC address of the Authenticator.

Configuring RADIUS Accounting

The AT-S62 management software supports RADIUS accounting for ports operating in the Authenticator role. The accounting information sent by the switch to a RADIUS server includes the date and time when clients log on and log off, as well as the number of packets sent and received by a switch port during a client session. For background information on this feature, refer to “RADIUS Accounting” on page 658. This feature is disabled by default on the switch.

To configure this feature, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **2** to select Port Access Control (802.1X).

   The Port Access Control (802.1X) menu is shown in Figure 229 on page 662.

3. From the Port Access Control (802.1X) menu, type **7** to select Configure Accounting.

   The RADIUS Accounting menu is shown in Figure 236.

```
      Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                           Marketing
User: Manager                              11:20:02 02-Mar-2006
                       RADIUS Accounting

1 - Status............ Disabled
2 - Port.............. 1813
3 - Type.............. Network
4 - Trigger Type...... Start_Stop
5 - Update Status..... Disabled
6 - Update Interval... 60

R - Return to Previous Menu

Enter your selection?
```

Figure 236. Radius Accounting Menu

4. Adjust the following parameters as necessary.

   1 - Status
   This parameter activates or deactivates RADIUS accounting on the switch. Select Enabled to activate the feature or Disabled to deactivate it. The default is Disabled.

   2 - Port
   This parameter specifies the UDP port for RADIUS accounting. The default is port 1813.

   3 - Type
   This parameter specifies the type of RADIUS accounting. The default is Network. This value cannot be changed.

   4 - Trigger Type
   This parameter specifies the action that causes the switch to send accounting information to the RADIUS server. The options are:

   Start_Stop
   The switch sends accounting information whenever a client logs on or logs off the network. This is the default.

   Stop
   The switch sends accounting information only when a client logs off.

   5 - Update Status
   This parameter controls whether the switch is to send interim accounting updates to the RADIUS server. The default is disabled. If you enable this feature, use the next option in the menu to specify the intervals at which the switch is to send the accounting updates.

   6 - Update Interval
   This parameter specifies the intervals at which the switch sends interim accounting updates to the RADIUS server. The range is 30 to 300 seconds. The default is 60 seconds.

5. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Section VII

Management Security

The chapters in this section explain the management security features of the AT-S62 software. The chapters include:

Chapter 30
Web Server

The chapter provides an overview of the web server feature and the procedure for configuring the server. It contains the following sections:

Web Server Overview

The AT-S62 management software comes with web server software so you can remotely manage a switch with a web browser from a management workstation on your network. (For instructions on how to manage a switch with a web browser, refer to the *AT-S62 Web Browser Interface User's Guide*.)

The web server can operate in two modes. The first is referred to as non-secure HTTP mode. In this mode, packets sent between the switch and the web browser during a management session are transmitted in plaintext. Anyone monitoring your network with a network analyzer, such as a sniffer, will be able to view the contents of the management packets.

The web server can also operate in the secure HTTPS mode where all communications between the switch and a web browser are encrypted. This feature uses the Secure Sockets Layer (SSL) protocol. It can help protect your switch from intruders who might be monitoring your network.

If you intend to use the secure HTTPS mode of the web server, there are several procedures you need to perform before you can configure the web server. You must create an encryption key, as explained in Chapter 31, "Encryption Keys" on page 687. You must also create a certificate and add the certificate to the certificate database. This latter part is explained in Chapter 32, "PKI Certificates and SSL" on page 705. For an overview to the procedures, refer to "General Steps to Configuring the Web Server for Encryption" on page 685. For an overview of all the steps, refer to "General Steps to Configuring the Web Server for Encryption" on page 685.

The default setting for the web server is enabled, with the non-secure HTTP mode as the default active mode.

Note
To use SSL in an enhanced stack, all switches in the stack must use SSL. For further information, refer to "SSL and Enhanced Stacking" on page 709.

Supported Protocols

The switch supports the following HTTP and HTTPs protocols:

- HTTP v1.0 and v1.1 protocols
- HTTPS v1.0 and v1.1 protocols running over SSL

The switch supports the following SSL protocols:

- SSL version 2.0
- SSL version 3.0
- TLS (Transmission Layer Security) version 1.0

Configuring the Web Server

This procedure explains how to enable and disable the web server and how to configure the HTTP and HTTPS settings from a local or Telnet management session. The default setting for the web server is enabled, with the non-secure HTTP mode as the active web server mode.

Before configuring the web server, note the following:

- You cannot make any changes to the HTTP or HTTPS settings while the web server is enabled. You must first disable the server before making changes.
- To configure the web server for the HTTPS secure mode, you must first create an encryption key and a certificate, and add the certificate to the certificate database. The management software will not allow you to configure the web server for the HTTPS secure mode until those steps have been completed. For instructions, refer to Chapter 31, "Encryption Keys" on page 687 and Chapter 32, "PKI Certificates and SSL" on page 705. For an overview of all the steps, refer to "General Steps to Configuring the Web Server for Encryption" on page 685.
- To make a change to an HTTP or HTTPS setting, you must perform the entire procedure. For instance, to change the port number for HTTP, you must first disable the web server and then reselect HTTP.

To configure the web server, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.
2. From the System Administration menu, type **4** to select Web Server Configuration.

   The Web Server Configuration Menu is shown in Figure 237.

```
      Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                        Production Switch
User: Manager                              11:20:02 02-Jan-2006
                     Web Server Configuration
1 - Status ................................. Enabled
2 - Mode ................................... HTTPS
3 - Port Number ............................ 443
4 - SSL Key ID ............................. 11

R - Return to Previous Menu

Enter your selection?
```

Figure 237. Web Server Configuration Menu

Menu option 4 is displayed only for HTTPS operation. The option is hidden for HTTP.

3. Type **1** to select Status to toggle the web server between enabled and disabled. To configure the web server, you must first disable it.

   Toggle between the following values:

   Enabled - Enables the web server. This is the default setting.

   Disabled - Disables the web server. (If you are making any changes to the web server settings, you must first disable it.)

4. Type **2** to select Mode to set the mode of the web server.

   The following prompt appears:

   ```
   Enter Web Server Mode (1 - HTTP, 2 - HTTPS):
   [1 to 2] ->
   ```

5. Choose one of the following:

   1 - HTTP to select the non-secure HTTP mode for the web server. This is the default value.

   2 - HTTPS to select the secure HTTPS mode. This setting activates the SSL protocol on the web server.

   If you are configuring HTTPS, the following prompt appears:

   ```
   Enter SSL Key ID ->
   ```

6. Enter an SSL Key ID.

   Enter the ID number of an encryption key on the switch. (To view the encryption key IDs, refer to “Creating an Encryption Key” on page 695.) You must have already created the encryption key and a certificate using the key. You must also have already added the certificate to the certificate database.

7. To change the protocol port number, type **3** to select Port Number. The following prompt appears:

   ```
   Enter Port Number [1 to 65535]-> 80
   ```

8. Enter the new protocol port number.

   The default port number for HTTP is 80. The default port number for HTTPS is 443.

9. To enable the web server, type **1** to toggle Status to **Enabled**.

10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

General Steps to Configuring the Web Server for Encryption

There are several procedures you need to perform in order to implement HTTPS and web browser encryption on the switch. This section is here to provide you with the general steps and the procedures for performing them. There is a section for configuring the web server with a self-signed certificate and another for a public or private CA certificate.

General Steps for a Self-signed Certificate

Below are the general steps to setting up the web server with a self-signed certificate.

1. Set the switch's date and time. You must do this before you create a self-signed certificate because the date and time are stamped in the digital document. For instructions, refer to "Setting the System Time" on page 61.
2. Create a key pair, as explained in "Creating an Encryption Key" on page 695.
3. Create a self-signed certificate using the key pair, as explained in "Creating a Self-signed Certificate" on page 718.
4. Add the certificate to the certificate database, as explained in "Adding a Certificate to the Database" on page 722.
5. Configure the web server on the switch by activating HTTPS and specifying the key pair used to create the certificate as the active key. This step is explained in "Configuring the Web Server" on page 683.

General Steps for a Public or Private CA Certificate

Below are the steps for setting up the web server with a public or private CA certificate. This requires generating an enrollment request.

1. Set the switch's date and time. You must do this before you create the enrollment request. The date and time are stamped in the request. The instructions for this are in "Setting the System Time" on page 61.
2. Create a key pair, as explained in "Creating an Encryption Key" on page 695.
3. Generate an enrollment request, as explained in "Generating an Enrollment Request" on page 730.
4. Upload the enrollment request from the AT-S62 file system onto your management workstation or a TFTP server, as explained in "Uploading a System File" on page 209.
5. Submit the enrollment request to the public or private CA.

6. Once you have received the appropriate certificates from the CA, download them into the AT-S62 file system from your management workstation or a TFTP server, as explained in “Downloading a System File” on page 202.

7. Add the certificates to the certificate database, as explained in “Adding a Certificate to the Database” on page 722.

8. Configure the web server on the switch by activating HTTPS and specifying the key pair used to create the enrollment request as the active key. This step is explained in “Configuring the Web Server” on page 683.

Chapter 31

Encryption Keys

This chapter describes how to improve the security of your switches with encryption keys. Because of the complexity of the feature, two overview sections are provided. The Basic Overview section offers a general review of the purpose of this feature along with relevant guidelines. For additional information, refer to the Technical Overview section. The sections in this chapter include:

For an overview of the procedures to configuring the switch’s web server for encryption, refer to “General Steps to Configuring the Web Server for Encryption” on page 685.

Basic Overview

Protecting your managed switches from unauthorized management access is an important role for a network manager. Network operations and security can be severely compromised should an intruder gain access to critical switch information, such as a manager's login username and password, and use that information to alter a switch's configuration settings.

One means by which an intruder could obtain critical switch information is by covertly monitoring network traffic with a network analyzer, such as a sniffer, and capturing management packets from remote Telnet or web browser management sessions. The payload in the packets exchanged during remote management sessions is transmitted in plaintext. The information garnered from the management packets could enable an intruder to access the management software on a switch.

One means of foiling this type assault is by encrypting the payload in the packets exchanged during a remote management session between a management workstation and a switch. Encryption makes the packets unintelligible to an outside agent. Only the remote workstation and the switch engaged in the management session are able to decode each other's packets.

The heart of encryption is the encryption key. The key converts plaintext into encrypted text, and vice versa. A key consists of two separate keys: a private key and a public key. Together they create a *key pair.*

The AT-S62 management software supports encryption for remote web browser management sessions using the Secure Sockets Layer (SSL) protocol. Adding encryption to your web browser management sessions involves creating one key pair and adding the public key of the key pair to a certificate, a digital document stored on the switch. You can have the switch create the certificate itself or you can have a public or private certificate authority (CA) create it for you. For an overview of the steps to adding encryption to your web browser management sessions, refer to "General Steps to Configuring the Web Server for Encryption" on page 685.

The Telnet application protocol does not support encryption. To add encryption when you remotely manage a switch using the menu interface, you must first obtain a Secure Shell (SSH) protocol application. SSH offers the same functionality as Telnet, but with encryption.

SSH encryption requires two key pairs on the switch— a server key pair and a host key pair. You then configure the Secure Shell protocol server software on the switch, as explained in Chapter 33, "Secure Shell (SSH) Protocol" on page 737, by specifying the keys as the host and server SSH keys.

Encryption Key Length

To create a key pair, you must specify its length. The length is given in bits. The range is 512 to 1,536 bits, in increments of 256 bits. The default is 512 bits.

The general rule on key lengths is that the longer the key, the more difficult it is for someone to break (decipher). If you are particularly concerned about the safety of your management sessions, you might go with a longer key length than the default, though in all likelihood the default will be more than sufficient.

It should be pointed out that creating a key is a very CPU intensive operation for a switch. The switch will not stop forwarding packets between the ports, but the process can impact the CPU's handling of network events, such as the processing of spanning tree BPDU packets. This can result in unexpected and unwanted switch behavior.

A key with the default length should take the switch less than a minute to create, while longer keys can take upwards of fifteen minutes. You should take this into account when creating a key so as not to impact the operations of your network. If you want a longer key, you might consider creating it before you connect the switch to the network, or during periods of low network traffic.

Encryption Key Guidelines

Below are guidelines to observe when creating an encryption key pair:

- Web browser encryption requires only one key pair.
- SSH encryption requires two key pairs. The keys must be of different lengths of at least one increment (256 bits) apart. The recommended size for the server key is 768 bits and the recommended size for the host key is 1024 bits.
- An AT-8500 Series switch can only use those key pairs it has generated itself. The switch cannot use a key created on another system and imported onto the switch.
- The AT-S62 management software does not allow you to copy or export a private key from a switch. However, you can export a public key.
- The AT-S62 management software uses the RSA public key algorithm.
- Web browser and SSH encryption can share a key pair.

Technical Overview

The encryption feature provides the following data security services:

- data encryption
- data authentication
- key exchange algorithms
- key creation and storage

Data Encryption

Data encryption for switches is driven by the need for organizations to keep sensitive data private and secure. Data encryption operates by applying an encryption algorithm and key to the original data (the plaintext) to convert it into an encrypted form (the ciphertext). The ciphertext produced by encryption is a function of the algorithm used and the key. Since it is easy to discover what type of algorithm is being used, the security of an encryption system relies on the secrecy of its key information. When the ciphertext is received by the remote router, the decryption algorithm and key are used to recover the original plaintext. Often, a checksum is added to the data before encryption. The checksum allows the validity of the data to be checked on decryption.

There are two main classes of encryption algorithm in use: symmetrical encryption and asymmetrical encryption.

Symmetrical Encryption

Symmetrical encryption refers to algorithms in which a single key is used for both the encryption and decryption processes. Anyone who has access to the key used to encrypt the plaintext can decrypt the ciphertext. Because the encryption key must be kept secret to protect the data, these algorithms are also called private, or secret key algorithms. The key can be any value of the appropriate length.

DES Encryption Algorithms

The most common symmetrical encryption system is the *Data Encryption Standard* (DES) algorithm (FIPS PUB 46). The DES algorithm has withstood the test of time and proved itself to be a highly secure encryption algorithm. To fully conform to the DES standard, the actual data encryption operations must be carried out in hardware. Software implementations can only be DES-compatible, not DES-compliant. The DES algorithm has a key length of 56 bits and operates on 64-bit blocks of data. DES can be used in the following modes:

- **Electronic Code Book (ECB)** is the fundamental DES function. Plaintext is divided into 64-bit blocks which are encrypted with the DES

algorithm and key. For a given input block of plaintext ECB always produces the same block of ciphertext.

- **Cipher Block Chaining (CBC)** is the most popular form of DES encryption. CBC also operates on 64-bit blocks of data, but includes a feedback step which chains consecutive blocks so that repetitive plaintext data, such as ASCII blanks, does not yield identical ciphertext. CBC also introduces a dependency between data blocks which protects against fraudulent data insertion and replay attacks. The feedback for the first block of data is provided by a 64-bit Initialization Vector (IV). This is the DES mode used for the switch's data encryption process.
- **Cipher FeedBack (CFB)** is an additive-stream-cipher method which uses DES to generate a pseudo-random binary stream that is combined with the plaintext to produce the ciphertext. The ciphertext is then fed back to form a portion of the next DES input block.
- **Output FeedBack (OFB)** combines the first IV DES algorithms with the plaintext to form ciphertext. The ciphertext is then used as the next IV.

The DES algorithm has been optimized to produce very high speed hardware implementations, making it ideal for networks where high throughput and low latency are essential.

Triple DES Encryption Algorithms

The Triple DES (3DES) encryption algorithm is a simple variant on the DES CBC algorithm. The DES function is replaced by three rounds of that function, an encryption followed by a decryption followed by an encryption. This can be done by using either two DES keys (112-bit key) or three DES keys (168-bit key).

The two-key algorithm encrypts the data with the first key, decrypts it with the second key and then encrypts the data again with the first key. The three-key algorithm uses a different key for each step. The three-key algorithm is the most secure algorithm due to the long key length.

There are several modes in which Triple DES encryption can be performed. The two most common modes are:

- **Inner CBC mode** encrypts the entire packet in CBC mode three times and requires three different initial is at ion vectors (IV's).
- **Outer CBC mode** triple encrypts each 8-byte block of a packet in CBC mode three times and requires one IV.

Asymmetrical (Public Key) Encryption

Asymmetrical encryption algorithms use two keys—one for encryption and one for decryption. The encryption key is called the public key because it cannot be used to decrypt a message and therefore does not have to be kept secret. Only the decryption, or private key, needs to be kept secret. The other name for this type of algorithm is public key encryption. The

public and private key pair cannot be randomly assigned, but must be generated together. In a typical scenario, a decryption station generates a key pair and then distributes the public key to encrypting stations. This distribution does not need to be kept secret, but it must be protected against the substitution of the public key by a malicious third party. Another use for asymmetrical encryption is as a digital signature. The signature station publishes its public key, and then signs its messages by encrypting them with its private key. To verify the source of a message, the receiver decrypts the messages with the published public key. If the message that results is valid, then the signing station is authenticated as the source of the message.

The most common asymmetrical encryption algorithm is RSA. This algorithm uses mathematical operations which are relatively easy to calculate in one direction, but which have no known reverse solution. The security of RSA relies on the difficulty of factoring the modulus of the RSA key. Because key lengths of 512 bits or greater are used in public key encryption systems, decrypting RSA encrypted messages is almost impossible using current technology. The AT-S62 software uses the RSA algorithm.

Asymmetrical encryption algorithms require enormous computational resources, making them very slow when compared to symmetrical algorithms. For this reason they are normally only used on small blocks of data (for example, exchanging symmetrical algorithm keys), and not for entire data streams.

Data Authentication

Data authentication for switches is driven by the need for organizations to verify that sensitive data has not been altered.

Data authentication operates by calculating a message authentication code (MAC), commonly referred to as a *hash*, of the original data and appending it to the message. The MAC produced is a function of the algorithm used and the key. Since it is easy to discover what type of algorithm is being used, the security of an authentication system relies on the secrecy of its key information. When the message is received by the remote switch, another MAC is calculated and checked against the MAC appended to the message. If the two MACs are identical, the message is authentic.

Typically a MAC is calculated using a keyed one-way hash algorithm. A keyed one-way hash function operates on an arbitrary-length message and a key. It returns a fixed length hash. The properties which make the hash function one-way are:

- it is easy to calculate the hash from the message and the key
- it is very hard to compute the message and the key from the hash
- it is very hard to find another message and key which give the same hash

The two most commonly used one-way hash algorithms are MD5 (Message Digest 5, defined in RFC 1321) and SHA-1 (Secure Hash Algorithm, defined in FIPS-180-1). MD5 returns a 128-bit hash and SHA-1 returns a 160-bit hash. MD5 is faster in software than SHA-1, but SHA-1 is generally regarded to be slightly more secure.

HMAC is a mechanism for calculating a keyed Message Authentication Code which can use any one-way hash function. It allows for keys to be handled the same way for all hash functions and it allows for different sized hashes to be returned.

Another method of calculating a MAC is to use a symmetric block cypher such as DES in CBC mode. This is done by encrypting the message and using the last encrypted block as the MAC and appending this to the original message (plain-text). Using CBC mode ensures that the whole message affects the resulting MAC.

Key Exchange Algorithms

Key exchange algorithms are used by switches to securely generate and exchange encryption and authentication keys with other switches. Without key exchange algorithms, encryption and authentication session keys must be manually changed by the system administrator. Often, it is not practical to change the session keys manually. Key exchange algorithms enable switches to re-generate session keys automatically and on a frequent basis.

The most important property of any key exchange algorithm is that only the negotiating parties are able to decode, or generate, the shared secret. Because of this requirement, public key cryptography plays an important role in key exchange algorithms. Public key cryptography provides a method of encrypting a message which can only be decrypted by one party. A switch can generate a session key, encrypt the key using public key cryptography, transmit the key over an insecure channel, and be certain that the key can only be decrypted by the intended recipient. Symmetrical encryption algorithms can also be used for key exchange, but commonly require an initial shared secret to be manually entered into all switches in the secure network.

The *Diffie-Hellman* algorithm, which is used by the AT-S62 management software, is one of the more commonly used key exchange algorithms. It is not an encryption algorithm because messages cannot be encrypted using Diffie-Hellman. Instead, it provides a method for two parties to generate the same shared secret with the knowledge that no other party can generate that same value. It uses public key cryptography and is commonly known as the first public key algorithm. Its security is based on the difficulty of solving the *discrete logarithm problem*, which can be compared to the difficulty of factoring very large integers.

A Diffie-Hellman algorithm requires more processing overhead than RSA-based key exchange schemes, but it does not need the initial exchange of public keys. Instead, it uses published and well tested public key values.

The security of the Diffie-Hellman algorithm depends on these values. Public key values less than 768 bits in length are considered to be insecure.

A Diffie-Hellman exchange starts with both parties generating a large random number. These values are kept secret, while the result of a public key operation on the random number is transmitted to the other party. A second public key operation, this time using the random number and the exchanged value, results in the shared secret. As long as no other party knows either of the random values, the secret is safe.

Creating an Encryption Key

This section contains the procedure for creating an encryption key pair.

Caution
Key generation is a CPU-intensive process. Because this process may affect switch behavior, Allied Telesyn recommends performing it when the switch is not connected to a network or during periods of low network activity.

To create an encryption key pair, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **7** to select Keys/ Certificates Configuration.

   The Keys/Certificates Configuration menu is shown in Figure 238.

```
Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                Production Switch
User: Manager                  11:20:02 02-Jan-2006
          Keys/Certificates Configuration
1 - Switch Distinguished Name (DN)
2 - Key Management
3 - Public Key Infrastructure (PKI) Configuration

R - Return to Previous Menu

Enter your selection?
```

Figure 238. Keys/Certificate Configuration Menu

3. Type **2** to select Key Management.

The Key Management menu is shown in Figure 239.

```
          Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                          Production Switch
User: Manager                                    11:20:02 02-Jan-2006
                             Key Management
ID        Algorithm    Length   Digest       Description
----------------------------------------------------------------
1         RSA-Private  512     642C6FC8      Production Switch key 1
2         RSA-Private  512     5333E64F      Production Switch key 2

1 - Create Key
2 - Delete Key
3 - Modify Key
4 - Export Key To File
5 - Import Key To File

N - Next Page
U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 239. Key Management Menu

This menu lists the key pairs already existing on the switch. The fields in the menu are described below:

ID
The identification number of the key.

Algorithm
The algorithm used in creating the encryption. This is always RSA - Private.

Length
The length of the key in bits.

Digest
The CRC32 value of the MD5 digest of the public key.

Description
The key's description.

4. To create a new encryption key pair, type **1** to select Create Key.

The Create Key menu is shown in Figure 240.

```
    Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                    Production Switch
User: Manager                          11:20:02 02-Jan-2006
                          Create Key
1 - Key ID ..............  0
2 - Key Type ............  RSA-Private
3 - Key Length ..........  512
4 - Key Description ....
5 - Generate Key

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 240. Create Key Menu

5. Type **1** to select Key ID.

   The following prompt is displayed:

   ```
   Enter Key Id -> [0 to 65535] -> 0
   ```

6. Enter an identification number for the key. This number can be from 0 to 65,535. This number is used only for identification purposes and not in generating the actual encryption key. The ID for each key on the switch must be unique.

   Note
   You cannot change the value for option 2 - Key Type. This value is always RSA - Private.

7. Type **3** to select Key Length.

   The following message is displayed:

   ```
   Enter Key Length ->[512 to 1536] -> 512
   ```

8. Enter a key length. The range is 512 to 1,536 bits, in increments of 256 bits (for example, 512, 768, 1024, etc). Before selecting a key length, note the following

- ❒ For an encryption key for SSL and web browser encryption, key length can be any valid value within the range.
- ❒ For SSH host and server key pairs, the two keys must be created separately and be of different lengths of at least one increment (256 bits) apart. The recommended length for the server key is 768 bits and the recommended length for the host key is 1024 bits.

9. Type **4** to create a key description.

   The following prompt is displayed:

   ```
   Enter new Description ->
   ```

10. Enter a description for the key. For instance, the description could reflect the key's function (for example, Sales switch SSL key). You can enter up to 40 alphanumeric characters including spaces.

11. Type **5** to generate the key.

    The following message is displayed:

    ```
    Key generation will take some time. Please wait...
    ```

    The management software begins to create the key. This process can take from less than a minute to more than fifteen minutes, depending on key length. Once the key is created, you will see this message:

    ```
    Press any key to continue ...
    ```

12. Press any key.

    The new key is added to the list of keys in the Key Management menu.

    Returning to the Main Menu to save the new key is not necessary with this procedure. This type of change is automatically saved by the management software.

    To create a self-signed SSL certificate using the new encryption key, go to "Creating a Self-signed Certificate" on page 718. To create an enrollment request for submission to a CA, go to "Generating an Enrollment Request" on page 730.

    If you created server and host keys for SSH encryption, go to "Configuring the SSH Server" on page 742 to configure the SSH server software on the switch.

Deleting an Encryption Key

This section contains the procedure for deleting an encryption key pair from the switch. Note the following before performing this procedure.

- Deleting a key pair from the key management database also deletes the key's corresponding ".UKF" file from the AT-S62 file system.
- You cannot delete a key pair if it is being used by SSL or SSH. You must first either disable the SSL or SSH server software or reconfigure the software by specifying another key.
- Deleting a key pair used in creating an SSL certificate voids the certificate.

This procedure starts from the Key Management menu. If you are unsure how to display the menu, perform steps 1 to 3 in "Creating an Encryption Key" on page 695.

To delete a key pair, do the following:

1. From the Key Management menu, type **2** to select Delete Key.
2. When prompted, enter the ID number of the key you want to delete.

   The key pair is deleted from the key database.

   Returning to the Main Menu to save your changes is not necessary with this procedure. This type of change is automatically saved by the management software.

Modifying an Encryption Key

The Key Management menu has a selection for modifying the description of an encryption key. This is the only item of a key you can modify.

This procedure starts from the Key Management menu. If you are unsure how to display the menu, perform steps 1 to 3 in “Creating an Encryption Key” on page 695.

To change the description of a key, perform the following procedure:

1. From the Key Management menu, type **3** to select Modify Key.

   The following prompt is displayed:

   ```
   Enter Key Id to modify -> [0 to 65535] -> 0
   ```

2. Enter the ID of the key whose description you want to modify.

   The following message is displayed.

   ```
   Enter new Description ->
   ```

3. Enter the new description for the key. The description can be up to 40 alphanumeric characters including spaces. To help identify the key, you might make the description the name of the web server the key will be used to protect (for example, Production switch web server).

   The following message is displayed:

   ```
   Press any key to continue ...
   ```

   The key has been modified.

4. Press any key to return to the Key Management menu.

   Returning to the Main Menu to save your changes is not necessary with this procedure. This type of change is automatically saved by the management software.

Exporting an Encryption Key

The following procedure exports the public key of a key pair into the AT-S62 file system. (The management software does not allow you to export a private key.) Before performing this procedure, please note the following:

- The only circumstance in which you are likely to perform this procedure is if you are using an SSH client that does not upload the key automatically when you start an SSH management session. You can use this procedure to export the SSH client key from the key database into the AT-S62 file system, from where you can download it onto the SSH management session and incorporate into your SSH client software.
- You should not use this procedure to export an SSL public key. Typically, an SSL public key only has value when incorporated into a certificate or enrollment request.

This procedure starts from the Key Management menu. If you are unsure how to display the menu, perform steps 1 to 3 in "Creating an Encryption Key" on page 695.

To export a public key into the file system, perform the following procedure:

1. From the Key Management Menu, type **4** to select Export Key to File.

   The Export Key to File Menu is shown in Figure 241.

```
    Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                    Production Switch
User: Manager                          11:20:02 02-Jan-2006
                  Export Key to File Menu
1 - Key ID ............ 0
2 - Key Type ........... RSA-Public
3 - Key File Format ... HEX
4 - Key File Name
5 - Export Key To File

R - Return to Previous Menu

Enter your selection?
```

Figure 241. Export Key to File Menu

2. Type **1** to select Key ID and, when prompted, enter the key ID of the public key you want to export into the file system.

Note
Key Type is a read-only field. You cannot change this value.

3. Type **3** to toggle Key File Format to specify the format of the key. Possible settings are:

   HEX - Indicates an internal format for storing files. Select this value for SSL configuration. This is the default.

   SSH - Indicates a format for a SSH1 environment. This is the correct setting for a key intended for an SSH1 client.

   SH2 - Indicates a format for a SSH2 environment. This is the correct setting for a key intended for an SSH2 client.

4. Type **4** to select Key File Name and specify a filename for the key. The filename can be from one to eight alphanumeric characters, not including the extension. Spaces are allowed. You must include the extension ".key".

5. Type **5** to select Export Key to File to export the key to a file.

   The following message is displayed:

   ```
   Key Export in Progress. Please wait...Done
   ```

6. Press any key to return to the Key Management menu.

   To view the public key in the switch's file system, refer to "Displaying System Files" on page 185.

   You do not need to return to the Main Menu to save your changes for this procedure. This type of change is automatically saved by the management software.

Importing an Encryption Key

Use the following procedure to import a public key from the AT-S62 file system into the key management database. If a file contains both public and private keys, only the public key is imported. The private key is ignored.

> **Note**
> It is unlikely you will ever have reason to perform this procedure. The switch can only use those keys it has generated itself.

This procedure starts from the Key Management menu. If you are unsure how to display the menu, perform steps 1 to 3 in "Creating an Encryption Key" on page 695.

To import a public key, perform the following procedure:

1. From the Key Management Menu, type **5** to select Import Key From File to import a RSA - Public key.

   The Import Key From File Menu is shown in Figure 242.

```
        Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006
                      Import Key From File Menu
1 - Key ID ............ 0
2 - Key Type ........... RSA-Public
3 - Key File Format ... HEX
4 - Key File Name .....
5 - Import Key From File

R - Return to Previous Menu

Enter your selection?
```

Figure 242. Import Key From File Menu

2. Type **1** to select Key ID and, when prompted, enter a unique key ID for the public key to import from the file system into the key management database. This must be an unused key ID. It cannot match any of the key IDs already in use on the switch.

> **Note**
> Option 2 - Key Type cannot be changed.

3. Type **3** to select Key File Format to choose the format of the key. Selections are:

   HEX - Indicates an internal format for storing files. Select this value for SSL configuration. This is the default.

   SSH - Indicates a format for a SSH1 environment. This is the correct setting for a key intended for an SSH1 client.

   SH2 - Indicates a format for a SSH2 environment. This is the correct setting for a key intended for an SSH2 client.

4. Type **4** to select Key File Name and, when prompted, specify the file name of the key.

   The key filename must include the ".key" extension. If you are unsure of the filename, display the files in the switch's file system by referring to "Displaying System Files" on page 185.

5. Type **5** to select Import Key From File to import a key to the switch from an external file.

   The following message is displayed:

   ```
   Key Import in Progress. Please wait...Done
   ```

   After you receive this message, the key is added to the Key Management database. See the Key Management Menu in Figure 239 on page 696.

   You do not need to return to the Main Menu to save your changes for this procedure. This type of change is automatically saved by the management software.

Chapter 32

PKI Certificates and SSL

This chapter contains the procedures for creating Public Key Infrastructure (PKI) certificates for web server security. Because of the complexity of this feature, two overview sections are provided. The Basic Overview section offers a general review of the purpose of certificates along with relevant guidelines. For additional information, refer to the Technical Overview section. This chapter contains the following sections:

Basic Overview

This chapter explains how to implement encryption for your web browser management sessions. Encryption can protect your managed switches from unauthorized access by making it impossible for an intruder monitoring network traffic to decipher the contents of the management packets exchanged between your workstation and a switch during a web browser management session.

Web browser encryption involves an encryption key pair and a digital document called a certificate. The key, as explained in Chapter 31, “Encryption Keys” on page 687, consists of two parts, a private key and a public key. The private key always remains on the switch. The public key is incorporated into a certificate. Your web browser downloads the certificate from the switch when you begin a management session.

Web browser encryption is provided by the Secure Sockets Layer (SSL) protocol. SSL was originally designed to offer security in Internet commerce and other web transactions, so as to provide Internet users a means of protecting their information from prying eyes as it crosses the Internet.

Of course, managing a switch with a web browser cannot be characterized as Internet commerce. But the sensitive nature of the information contained within the management packets makes protecting the packets a critical component of network security, and SSL provides the means for doing just that.

Types of Certificates

The AT-S62 management software supports two types of certificates. The first is called a *self-signed certificate*. This is the quickest and easiest to create because the switch creates it itself. For small to medium sized networks, this might be the way to go. The procedure for creating this kind of certificate is found in “Creating a Self-signed Certificate” on page 718. To review all the steps to configuring the web server on the switch for this type of certificate, refer to “General Steps for a Self-signed Certificate” on page 685.

The second type of certificate is a *CA certificate*. Here, you create the encryption key pair on the switch but someone else issues the certificate, which you then load onto the switch. That person, group, or organization that issues the certificate is called a *certification authority* (CA).

There are two kinds of CAs: public and private. A public CA issues certificates for other companies and organizations. A well known example is Verisign. A public CA will require proof of the identify of the company or organization that wants a certificate before it will issue it.

Public CAs issue certificates typically intended for use by the general public. Since a certificate for an AT-8500 Series switch is not intended for

general use, but will only be used by you and other network managers, you might decide that the switch's certificate need not be issued by this type of CA.

Some large companies have private CAs. This is a person or group within the company with the responsibility of issuing certificates for the company's network equipment. The value of a private CA is that the company can keep track of the certificates and control access to various network devices.

If your company is large enough, it might have a private CA and you might want that group to issue any AT-8500 Series certificates, if for no other reason than to follow company policy.

To obtain a CA certificate you have to create a key pair. You then need to generate an digital document called an *enrollment request*. The request will contain the public key, along with other information you want the CA to use to create the certificate.

Before you send an enrollment request to a CA, you should first contact the CA to determine what other documents or procedures might be required in order for the CA to create the certificate. This is particularly important with public CAs, which typically have strict guidelines on issuing certificates.

Distinguished Names

Part of the task to creating a self-signed certificate or enrollment request is selecting a *distinguished name*. A distinguished name is integrated into a certificate along with the key. A distinguished name can have up to five parts. The parts are:

- cn - common name

  This can be the name of the person who will use the certificate.

- ou - organizational unit

  This is the name of a department, such as Network Support or IT.

- o - organization

  This is the name of the company.

- st - state

  This is the state.

- c - country

  This is the country

A certificate name does not have to contain all of these parts. You can use as many or as few as you want. You separate the parts with a comma. You can use alphanumeric characters, as well as spaces in the name strings.

You cannot use quotation marks. To use the following special characters {=,+<>#;\<CR>}, type a “\” before the character

Here are a few examples. This distinguished name contains only one part, the name of the switch:

cn=Production Switch

This distinguished name omits the common name, but includes everything else:

ou=Network Support,o=XYZ Inc.,st=CA,c=US

So what would be a good distinguished name for a certificate for an AT-8500 Series switch? If the switch has an IP address, such as a master switch, you could use its address as the name. The following example is a distinguished name for a certificate for a master switch with the IP address 149.11.11.11:

cn=149.11.11.11

If your network has a Domain Name System and you mapped a name to the IP address of a switch, you can specify the switch’s name instead of the IP address as the distinguished name.

For those switches that do not have an IP address, such as slave switches, you could assign their certificates a distinguished name using the IP address of the master switch of the enhanced stack.

There is a benefit to giving a certificate a distinguished name equivalent to a master switch’s IP address or domain name. It relates to what happens when you start a web browser management session with a switch using SSL. The web browser on your workstation will check to see if the name to whom the certificate was issued matches the name of the web site. In the case of a master or slave AT-8500 Series switch, the web site’s name is the master switch’s IP address or domain name. If the names do not match, the web browser displays a security warning. Of course, even if you see the security warning, you can simply close the warning prompt. The management session will still use encryption.

Note
If the certificate will be issued by a private or public CA, you should check with the CA to see if they have any rules or guidelines on distinguished names for the certificates they issue.

SSL and Enhanced Stacking

Secure Sockets Layer (SSL) is supported in an enhanced stack, but only when all switches in the stack are using the feature.

A web server can operate in one of two modes -- HTTP or HTTPS. When a switch's web server is operating in HTTP, management packets are transmitted in plaintext. When it operates in HTTPS, management packets are sent encrypted.

The web server on an AT-8500 Series switch, and also an AT-8400 Series switch, can operate in either mode. Enhanced stacking switches that do not support SSL, such as the AT-8000 Series switches, use HTTP exclusively.

A web browser management session of the switches in an enhanced stack cannot change its security mode during a session. The management session assumes that the web server mode that the master switch is using is the same for all the switches in the stack.

As an example, if the master switch is using HTTPS, a web browser management session assumes that all the other switches in the stack are also using HTTPS, and it will not allow you to manage any switches running HTTP.

For those networks that consist of enhanced stacking switches where some switches support SSL and others do not, there are two approaches you can take. One is to create different enhanced stacks for the different switches. You could create one enhanced stack for those switches that support SSL and another stack for those that do not. You create different enhanced stacks by assigning switches to different Management VLANs, as explained in "Specifying a Management VLAN" on page 579.

Another approach is to leave the switches in one enhanced stack, but designate two master switches. One master switch could be using HTTP and the other HTTPS. When you want to use your web browser to manage those switches that support SSL, you would start the management session on the master switch whose server mode is set to HTTPS. To manage those switch not supporting SSL, you would start the management session on the master switch whose web server is set to HTTP.

In order to implement SSL in an enhanced stack, each switch in the stack must be given its own encryption key pair and certificate. Switches cannot share keys and certificates. When you start a web browser management session on the master switch of an enhanced stack, the management session uses the certificate and key pair on the master switch. When you change to another switch in the stack, the management session starts to use the certificate and key pair on that switch, and so forth.

Guidelines

Here are guidelines to creating certificates:

- A certificate can have only one public key.
- A switch can only use those certificates with keys it generated itself.
- You can create multiple certificates on a switch, but the device will only use the certificate whose key pair has been designated as the active key pair for the switch's web server.
- Most web browsers support both unsecured (plaintext) and secured (encrypted) operation. These modes are referred to as HTTP and HTTPS, respectively. If you choose to use encryption when you manage a switch, the web browser you use must support HTTPS.

Technical Overview

The Secure Sockets Layer (SSL) feature is a security protocol that provides a secure and private TCP connection between a client and server.

SSL can be used with many higher layer protocols including HTTP, File Transfer Protocol (FTP) and Net News Transfer Protocol (NNTP). Most web browsers and servers support SSL, and its most common deployment is for secure connections between a client and server over the Internet.

The switch supports SSL versions 2.0 (client hello only) and 3.0 which were developed by Netscape, and the Internet Engineering Task Force (IETF) standard for SSL, known as SSL version 3.1 or Transport Layer Security (TLS).

Within the Ethernet protocol stack, SSL is a layer 4 protocol that is in between the HTTP and TCP protocol layers. HTTP communicates with SSL in the same way as with TCP. In other words, TCP processes SSL requests like any other protocol requesting its services.

SSL provides a secure connection over which web pages can be accessed from an HTTP server. The operation of SSL is transparent to the end user who is accessing a web site with the following exceptions:

- the site's URL changes from HTTP to HTTPS
- the browser indicates that it is a secured connection by displaying an icon, such as a padlock icon

By default, HTTP and HTTPS use the separate well-known ports 80 and 443, respectively. Secure connections over the Internet are important when transmitting confidential data such as credit card details or passwords. SSL allows the client to verify the server's identity before either side sends any sensitive information. SSL also prevents a third party from interfering with the message because only trusted devices have access to the unprotected data.

SSL Encryption

SSL uses *encryption* to ensure the security of data transmission. Encryption is a process that uses an algorithm to encode data so it can only be accessed by a trusted device. An encrypted message remains confidential.

All application data messages are authenticated by SSL with a *message authentication code* (MAC). The MAC is a checksum that is created by the sender and is sent as part of the encrypted message. The recipient re-calculates the MAC, and if the values match, the sender's identity is verified. The MAC also ensures that the message has not been tampered with by a third party because any change to the message changes the MAC.

SSL uses *asymmetrical (Public Key)* encryption to establish a connection between client and server, and *symmetrical (Secret Key)* encryption for the data transfer phase.

User Verification

An SSL connection has two phases: *handshake* and *data transfer*. The *handshake* initiates the SSL *session*, during which data is securely transmitted between a client and server. During the handshake, the following occurs:

1. The client and server establish the SSL version they are to use.
2. The client and server negotiate the *cipher suite* for the session, which includes encryption, authentication, and key exchange algorithms.
3. The *symmetrical key* is exchanged.
4. The client authenticates the server (optionally, the server authenticates the client).

SSL messages are encapsulated by the *Record Layer* before being passed to TCP for transmission. Four types of SSL messages exist, they are:

- ❒ Handshake
- ❒ Change Cipher Spec
- ❒ Alert
- ❒ Application data (HTTP, FTP or NNTP)

As discussed previously, the *Handshake* message initiates the SSL session.

The *Change Cipher Spec* message informs the receiving party that all subsequent messages are encrypted using previously negotiated security options. The parties use the strongest cryptographic systems that they both support.

The *Alert* message is used if the client or server detects an error. Alert messages also inform the other end that the session is about to close. In addition, the Alert message contains a severity rating and a description of the alert. For example, an alert message is sent if either party receives an invalid certificate or an unexpected message.

The *Application data* message encapsulates the encrypted application data.

Authentication

Authentication is the process of ensuring both the web site and the end user are genuine. In other words, they are not imposters. Both the server and an individual users need to be authenticated. This is especially important when transmitting secure data over the Internet.

To verify the authenticity of a server, the server has a public and private key. The public key is given to the user.

SSL uses *certificates* for authentication. A certificate binds a public key to a server name. A Certification Authority issues certificates after checking that a public key belongs to its claimed owner. There are several agencies that are trusted to issue certificates. Individual browsers have approved Root CAs that are built in to the browser.

Public Key Infrastructure

The Public Key Infrastructure (PKI) feature is part of the switch's suite of security modules, and consists of a set of tools for managing and using certificates. The tools that make up the PKI allow the switch to securely exchange public keys, while being sure of the identity of the key holder.

The switch acts as an End Entity (EE) in a certificate-based PKI. More specifically, the switch can communicate with Certification Authorities (CAs) and Certificate Repositories to request, retrieve and verify certificates.The switch allows protocols running on the switch, such as ISAKMP, access to these certificates. The following sections of this chapter summarize these concepts and describe the switch's implementation of them.

Public Keys

Public key encryption involves the generation of two keys for each user, one private and one public. Material encrypted with a private key can only be decrypted with the corresponding public key, and vice versa. An individual's private key must be kept secret, but the public key may be distributed as widely as desired, because it is impossible to calculate the private key from the public key. The advantage of public key encryption is that the private key need never be exchanged, and so can be kept secure more easily than a shared secret key.

Message Encryption

One of the two main services provided by public key encryption is the exchange of encrypted messages. For example, user 1 can send a secure message to user 2 by encrypting it with user 2's public key. Only user 2 can decrypt it, because only user 2 has access to the corresponding private key.

Digital Signatures

The second main service provided by public key encryption is digital signing. Digital signatures both confirm the identity of the message's supposed sender and protect the message from tampering. Therefore they provide message authentication and non-repudiation. It is very difficult for the signer of a message to claim that the message was corrupted, or to deny that it was sent.

Both the exchange of encrypted messages and digital signatures are secure only if the public key used for encryption or decryption belongs to the message's expected recipient. If a public key is insecurely distributed, it is possible a malicious agent could intercept it and replace it with the malicious agent's public key (the Man-in-the-Middle attack). To prevent

this, and other attacks, PKI provides a means for secure transfer of public keys by linking an identity and that identity's public key in a secure certificate.

Warning
While a certificate binds a public key to a subject to ensure the public key's security, it does not guarantee that the security of the associated private key has not been breached. A secure system is dependent upon private keys being kept secret, by protecting them from malicious physical and virtual access.

Certificates

A *certificate* is an electronic identity document. To create a certificate for a subject, a trusted third party (known as the Certification Authority) verifies the subject's identity, binds a public key to that identity, and digitally signs the certificate. A person receiving a copy of the certificate can verify the Certification Authority's digital signature and be sure that the public key is owned by the identity in it.

The switch can generate a self-signed certificate but this should only be used with an SSL enabled HTTP server, or where third party trust is not required.

X.509 Certificates

The X.509 specification specifies a format for certificates. Almost all certificates use the X.509 version 3 format, described in RFC 2459, *Internet X.509 Public Key Infrastructure Certificate and CRL Profile*. This is the format which is supported by the switch.

An X.509 v3 certificate consists of:

- A serial number, which distinguishes the certificate from all others issued by that issuer. This serial number is used to identify the certificate in a Certificate Revocation List, if necessary.
- The owner's identity details, such as name, company and address.
- The owner's public key, and information about the algorithm with which it was produced.
- The identity details of the organization which issued the certificate.
- The issuer's digital signature and the algorithm used to produce it.
- The period for which the certificate is valid.
- Optional information is included, such as the type of application with which the certificate is intended to be used.

The issuing organization's digital signature is included in order to authenticate the certificate. As a result, if a certificate is tampered with during transmission, the tampering is detected.

Elements of a Public Key Infrastructure

A Public Key Infrastructure is a set of applications which manage the creation, retrieval, validation and storage of certificates. A PKI consists of the following key elements:

- At least one Certification Authority (CA), which issues and revokes certificates.
- At least one publicly accessible repository, which stores certificates and Certificate Revocation Lists.
- At least one End Entity (EE), which retrieves certificates from the repository, validates them and uses them.

End Entities (EE)

End Entities own public keys and may use them for encryption and digital signing. An entity which uses its private key to digitally sign certificates is not considered to be an End Entity, but is a Certification Authority.

The switch acts as an End Entity.

Certification Authorities

A Certification Authority is an entity which issues, updates, revokes and otherwise manages public keys and their certificates. A CA receives requests for certification, validates the requester's identity according to the CA's requirements, and issues the certificate, signed with one of the CA's keys. CAs may also perform the functions of End Entities, in that they may make use of other CAs' certificates for message encryption and verification of digital signatures.

An organization may own a Certification Authority and issue certificates for use within its own networks. In addition, an organization's certificates may be accepted by another network, after an exchange of certificates has validated a certificate for use by both parties. As an alternative, an outside CA may be used. The switch can interact with the CA, whether a CA is part of the organization or not, by sending the CA requests for certification.

The usefulness of certificates depends on how much you trust the source of the certificate. You must be able to trust the issuing CA to verify identities reliably. The level of verification required in a given situation depends on the organization's security needs.

Certificate Validation

To validate a certificate, the End Entity verifies the signature in the certificate, using the public key of the CA who issued the certificate.

CA Hierarchies and Certificate Chains

It may not be practical for every individual certificate in an organization to be signed by one Certification Authority. A certification hierarchy may be formed, in which one CA (for example, national headquarters) is declared

to be the root CA. This CA issues certificates to the next level down in the hierarchy (for example, regional headquarters), who become subordinate CAs and issue certificates to the next level down, and so on. A hierarchy may have as many levels as needed.

Certificate hierarchies allow validation of certificates through certificate chains and cross-certification. If a switch X, which holds a certificate signed by CA X, wishes to communicate securely with a switch Y, which holds a certificate signed by CA Y, there are two ways in which the switches can validate each other's certificates. Cross-certification occurs when switch X validates switch Y's CA (CA Y) by obtaining a certificate for switch Y's CA which has been issued by its own CA (CA X). A certificate chain is formed if both CA X and CA Y hold a certificate signed by a root CA Z, which the switches have verified out of band. Switch X can validate switch Y's certificate (and vice versa) by following the chain up to CA Z.

Root CA Certificates

A root CA must sign its own certificate. The root CA is the most critical link in the certification chain, because the validity of all certificates issued by any CA in the hierarchy depends on the root CA's validity. Therefore, every device which uses the root CA's certificate must verify it out-of-band.

Out-of-band verification involves both the owner of a certificate and the user who wishes to verify that certificate generating a one-way hash (a fingerprint) of the certificate. These two hashes must then be compared using at least one non-network-based communication method. Examples of suitable communication methods are mail, telephone, fax, or transfer by hand from a storage device such as a smartcard or floppy disk. If the two hashes are the same, the certificate can be considered valid.

Certificate Revocation Lists (CRLs)

A certificate may become invalid because some of the details in it change (for example, the address changes), because the relationship between the Certification Authority (CA) and the subject changes (for example, an employee leaves a company) or because the associated private key is compromised. Every CA is required to keep a publicly accessible list of its certificates which have been revoked.

PKI Implementation

The following sections discuss Allied Telesyn's implementation of PKI for the AT-S62 management software. The following topics are covered:

- PKI Standards
- Certificate Retrieval and Storage
- Certificate Validation
- Root CA Certificates

PKI Standards

The following standards are supported by the switch:

- draft-ietf-pkix-roadmap-05 — *PKIX Roadmap*
- RFC 1779 — *A String Representation of Distinguished Names*
- RFC 2459 — *PKIX Certificate and CRL Profile*
- RFC 2511 — *PKIX Certificate Request Message Format*
- PKCS #10 v1.7 — *Certification Request Syntax Standard*

Certificate Retrieval and Storage

Certificates are stored by CAs in publicly accessible repositories for retrieval by end entities. The following repositories used in PKI are commonly accessed via the following protocols: *Hypertext Transfer Protocol* (HTTP), *File Transfer Protocol* (FTP).

Before the switch can use a certificate, it must be retrieved and manually added to the switch's Certificate Database, which is stored in RAM memory. The switch attempts to validate the certificate, and if validation is successful the certificate's public key is available for use.

Root CA Certificate Validation

Root CA certificates are verified out of band by comparing the certificate's *fingerprint* (the encrypted one-way hash with which the issuing CA signs the certificate) with the fingerprint which the CA has supplied by a non-network-based method. To view a certificate's fingerprint, use the procedure described in "Viewing a Certificate" on page 728.

Creating a Self-signed Certificate

This section contains the procedure for creating a self-signed certificate. Please review the following before you perform the procedure:

- For a general review of all the steps to configuring the switch for a self-signed certificate, refer to “General Steps for a Self-signed Certificate” on page 685.)
- The switch’s time and date must be set before creating a self-signed certificate. You can set this manually or you can configure the switch to obtain the date and time from an SNTP server on your network or the Internet. For instructions, refer to “Setting the System Time” on page 61.
- You must generate an encryption key pair before you create a certificate. For instructions, refer to “Creating an Encryption Key” on page 695.
- During this procedure you are prompted to enter the ID number of the encryption key pair to be used to create the certificate. If you do not know the ID number, refer to “Creating an Encryption Key” on page 695 to view key ID numbers.

To create a self-signed certificate, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.

   The Security and Services menu is shown in Figure 103 on page 313.

2. From the Security and Services menu, type **7** to select Keys/Certificates Configuration.

   The Keys/Certificates Configuration menu is shown in Figure 238 on page 695.

   Note
   The certificate must have a distinguished name. You can specify the distinguished name for the certificate from this menu by selecting option 1 - Distinguished Name in the Keys/Certificates Configuration menu and entering the name. Or, you can wait and specify the distinguished name later in this procedure. For information about distinguished names, refer to “Distinguished Names” on page 707.

3. From the Keys/Certificate menu, select **3** to select Public Key Infrastructure (PKI) Configuration.

The Public Key Infrastructure (PKI) Configuration menu is shown in Figure 243.

```
          Allied Telesyn Ethernet Switch AT-8524M - ATS62
                         Production Switch

User: Manager                                11:20:02 02-Jan-2006
            Public Key Infrastructure (PKI) Configuration

1 - Maximum Number of Certificates....... 256
2 - X509 Certificate Management
3 - Generate Enrollment Request

R - Return to Previous Menu

Enter your selection?
```

Figure 243. Public Key Infrastructure (PKI) Configuration Menu

4. Type **2** to select X509 Certificate Management.

   The X509 Certificate Management menu is shown in Figure 244.

```
          Allied Telesyn Ethernet Switch AT-8524M - ATS62
                         Production Switch
User: Manager                                  11:20:02 02-Jan-2006
                     X509 Certificate Management
Certificate Database:
Name                 State         MTrust   Type   Source
-------------------------------------------------------------
Switch43cert         Trusted        False    Self   Command

1 - Create Self-Signed Certificate
2 - Add Certificate
3 - Delete Certificate
4 - Modify Certificate
5 - View Certificate Details

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 244. X509 Certificate Management Menu

The Certificate Database portion of the window lists the certificates currently in the database. These could be certificates that you created or had a CA create. The switch's web server can only use a certificate if it is in the database.

Note
In the X509 Certificate Management Menu, MTrust means manually trusted. This field indicates that you verified the certificate. The Source field indicates the certificate was generated on the switch.

5. Type **1** to select Create Self-Signed Certificate.

   The Create Self-Signed Certificate menu is shown in Figure 245.

```
  Allied Telesyn Ethernet Switch AT-8524M - ATS62
                Production Switch
User: Manager                   11:20:02 02-Jan-2006
           Create Self-Signed Certificate
1 - Certificate Name.............
2 - Key Pair ID.................. 0
3 - Format....................... DER
4 - Serial Number................ 0
5 - Subject DN...................
6 - Create Self-Signed Certificate

R - Return to Previous Menu

Enter your selection?
```

Figure 245. Create Self-Signed Certificate Menu

6. Type **1** to select Certificate Name to enter a filename for the certificate.

   The following message is displayed:

```
Enter certificate name (24 chars max) ->
```

7. Enter a filename for the certificate. This is the filename under which the certificate will be stored as in the AT-S62 file system. The name can be up to 24 alphanumeric characters. Spaces are allowed.

Note
The management software automatically adds a ".cer" extension to the filename.

8. Type **2** to select Key Pair ID.

   The following message is displayed:

```
Enter certificate Key Pair ID -> [0 to 65535] ->
```

9. Enter the ID number of the encryption key you want to use to create the certificate. The encryption key must already exist on the switch. (If you have forgotten the key ID number, return to the Key Management menu to view the keys on the switch.) The value can be from 0 to 65,535.

10. Type **3** to select Format to choose the encoding format for the certificate. Possible settings are:

    DER - Indicates the certificate contents are in a binary format. This is the default.

    PEM - Indicates the certificate are in the Privacy Enhanced Mail (PEM) format which is an ASCII format.

11. Type **4** to select Serial Number.

    The following message is displayed:

    ```
    Enter certificate serial number -> [0 to 2147483647] -> 0
    ```

12. Enter a value between 0 and 2,147,483,647.

    Self-signed certificates are usually assigned a serial number of 0.

13. Type **5** to select Subject DN and enter a distinguished name for the certificate. (Do not enclose the distinguished name in quotes.)

    Note
    If you did not enter a distinguished name back in Step 2, you need to enter one here. A certificate must have a distinguished name. For further information, refer to "Distinguished Names" on page 707. If you enter a name both here and in Step 2, the certificate will contain the name entered here.

14. Type **6** to create the certificate.

    The following message is displayed:

    ```
    Please wait while certificate is generated...Done!
    ```

15. Press any key.

    The X509 Certificate Management menu is displayed again.

    The certificate is automatically saved in the AT-S62 file system. You do not need to return to the Main Menu to permanently save the new certificate.

16. Go to the next procedure to add the certificate to the certificate database.

Adding a Certificate to the Database

Once you have created a certificate or received a certificate from a public or private CA, you need to add it into the certificate database to make it available for use by the switch's web server. After you add a certificate to the certificate database, it appears in the X509 Certificate Management menu.

During the procedure you are asked to specify the certificate's filename. If you have forgotten the certificate's filename, refer to "Displaying System Files" on page 185.

To add a certificate to the certificate database, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **7** to select Keys/ Certificates Configuration.
3. From the Keys/Certificate menu, type **3** to select Public Key Infrastructure (PKI) Configuration.
4. From the Public Key Infrastructure (PKI) Configuration menu, type **2** to select X509 Certificate Management.

   The X509 Certificate Management menu is shown in Figure 244 on page 719.
5. From the X509 Certificate Management menu, type **2** to select Add Certificate.

   The Add Certificate Menu is shown in Figure 246.

```
 Allied Telesyn Ethernet Switch AT-8524M - ATS62
                Production Switch
User: Manager                   11:20:02 02-Jan-2006
              Add Certificate Menu
1 - Certificate Name .............
2 - State ........................ Trusted
3 - Type ......................... EE
4 - File Name .....................
5 - Add Certificate

R - Return to Previous Menu

Enter your selection?
```

Figure 246. Add Certificate Menu

6. Type **1** to select Certificate Name and enter a name for the certificate.

   This is the name for the certificate as it will appear in the certificate database list. You can enter up to 24 alphanumeric characters. Spaces are allowed. No extension is needed.

   You might want to include in the name the filename of the certificate in the file system. This will make it easier for you to match a certificate in the database with its corresponding file in the file system. Here is an example:

   ```
   Switch 12 - sw12.cer
   ```

7. Type **2** to set the certificate state. Possible settings are:

   Trusted - This value indicates you have verified that the certificate is from a trusted CA. This is the default.

   Untrusted - This value indicates the certificate is from an untrusted CA either because you have not verified the CA or you have verified the CA is untrusted.

   Note
   This parameter has no affect on the operation of a certificate. The parameter is included only for informational purposes when the certificate is displayed in the certificate database.

8. Type **3** to specify the type of certificate. There are 3 types to choose from:

   EE - Indicates the certificate was issued by a public or private CA. This is the default.

   CA - Indicates the certificate belongs to a public or private CA.

   Self - This value is a self-signed certificate. Use this value for a self-signed certificate. The switch treats this type of certificate as its own.

   Note
   This parameter has no affect on the operation of a certificate. The parameter is included only for informational purposes when the certificate is displayed in the certificate database.

9. Type **4** to select File Name and specify the filename of the certificate.

   This is the filename of the certificate in the AT-S62 file system, with the ".cer" extension. For example, if you created a self-signed certificate and gave it the name "webserver127", the filename of the certificate would be "webserver127.cer". If you have forgotten the filename of the certificate, refer to "Displaying System Files" on page 185.

10. Type **5** to select Add Certificate to add the certificate to the certificate database.

    The management software adds the certificate to the database, a process that requires only a few seconds.

11. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying a Certificate

The procedure in this section modifies a certificate. (The certificate to be modified must be in the certificate database.) Here are the certificate items you can modify:

- State - trusted or untrusted
- Type - EE, CA, or Self

Note
These parameters have no affect on the operation of a certificate. They are only included for informational purposes when the certificate is displayed in the certificate database.

This procedure starts from the X509 Certificate Management menu. If you are unsure how to access the menu, perform steps 1 to 4 in the procedure "Adding a Certificate to the Database" on page 722.

To modify a certificate, perform the following procedure:

1. From the X509 Certificate Management menu, type **4** to select Modify Certificate. The following message is displayed:

   ```
   Enter a certificate name ->
   ```

2. Enter the name of the certificate you want to modify. (This field is case-sensitive.)

   The Modify Certificate Menu is shown in Figure 247.

```
     Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                     Production Switch
User: Manager                          11:20:02 02-Jan-2006
                   Modify Certificate Menu
1 - Certificate Name................. Switch12
2 - State ........................... Trusted
3 - Type ............................ Self
4 - Modify Certificate

R - Return to Previous Menu

Enter your selection?
```

Figure 247. Modify Certificate Menu

Note
Option 1 - Certificate Name cannot be changed.

3. Type **2** to select State and specify if a certificate is trusted or untrusted.

   Trusted - This value indicates you have verified that the certificate is from a trusted CA. This is the default.

   Untrusted - This value indicates the certificate is from an untrusted CA either because you have not verified the CA or you have verified the CA is untrusted.

4. Type **3** to specify the type assigned to the certificate. There are 3 types to choose from:

   EE - This value indicates the End Entity type. When you specify this type, the switch tags the certificate to indicate that it belongs to another end entity. This is the default.

   CA - Use this value for a certificate issued by a public or private CA.

   Self - Use this value for a self-signed certificate. This type of certificate is created by the switch itself. The switch treats this type of certificate as its own.

5. Type **4** to select Modify Certificate.

   Your changes are implement in the certificate.

   The following message is displayed:

   ```
   Please wait while certificate is updated...Done.
   ```

6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting a Certificate

The procedure in this section deletes a certificate from the certificate database. Please note the following before performing this procedure:

- Deleting a certificate from the database does not delete it from the switch. It continues to reside in the AT-S62 file system. To completely remove a certificate from the switch, you must also delete it from the file system. For instructions, refer to "Copying, Renaming, and Deleting System Files" on page 183.
- You cannot delete a certificate from the database if you specified its corresponding encryption key as the active key in the web server configuration. The switch will consider the certificate as in use and will not allow you to delete it. You must first configure the web server with another encryption key pair for a different certificate. For instructions, refer to "Configuring the Web Server" on page 683.

This procedure starts from the X509 Certificate Management menu. If you are unsure how to access the menu, perform steps 1 to 4 in the procedure "Adding a Certificate to the Database" on page 722.

To delete a certificate from the certificate database, perform the following procedure:

1. From the X509 Certificate Management menu, type **3** to delete a certificate.

   The following message is displayed:

   ```
   Enter certificate name (ALL - delete all) ->
   ```

2. Enter the name of the certificate you want to delete. (This field is case-sensitive.) To delete all the certificates, enter ALL.

3. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Viewing a Certificate

This procedure displays information about a certificate, such as its distinguished name and serial number.

This procedure starts from the X509 Certificate Management menu. If you are unsure how to access the menu, perform steps 1 to 4 in the procedure “Adding a Certificate to the Database” on page 722.

To view the details of a certificate, perform the following procedure:

1. From the X509 Certificate Management menu, type **5** to select View Certificate Details.

   The following message is displayed:

   ```
   Enter certificate name ->
   ```

2. Enter a name of the certificate you want to view. (This field is case-sensitive.)

   The View Certificate Details menu (page 1) is shown in Figure 248.

```
          Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                          Production Switch
User: Manager                                      11:20:02 02-Jan-2006
                       View Certificate Details
Certificate Details:
   Name ............... Switch12
   State ............... Trusted
   Manually Trusted ... True
   Type ............... Self
   Source ............. Command

   Version ............ V3 (0X2)
   Serial Number ...... 0 (0X0)
   Signature Alg ...... md5WithRSAEncryption
   Public Key Alg ..... rsaEncryption
   Not Valid Before ... Jan 9 01:28:18 2004 GMT
   Not Valid After .... Jan 8 01:28:18 2006 GMT

N - Next Page
R - Return to Previous Menu

Enter your selection?
```

Figure 248. View Certificate Details Menu (page 1)

3. Type **N** to see the second page of certificate details.

The View Certificate Details menu (page 2) is shown in Figure 249.

```
            Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                            Production Switch
User: Manager                                          11:20:02 02-Jan-2006
                          View Certificate Details

 Subject ......... CN=149.44.44.44
 Issuer .......... CN=149.44.44.44
 MD5 Fingerprint...4E:76:06:FA:F6:C1:DA:FF:4D:E9:76:02:1D:8F:DA:CB
 SHA1 Fingerprint..F8:43:CB:E2:0A:BF:4A:02:CA:C6:B0:47:DF:74:1E:D3:A8:A3:F0:00

N - Previous Page
R - Return to Previous Menu

Enter your selection?
```

Figure 249. View Certificate Details Menu (page 2)

The fields are defined below:

Name - lists the name of the certificate.

State - Indicates the certificate is Trusted or Untrusted.

Manually Trusted - Indicates you verified the certificate is from a trusted or untrusted authority.

Type - Indicates the type of the certificate. The options are EE, SELF, and CA.

Source - Indicates the certificate was created on the switch.

Version - Indicates the version number of the software.

Serial Number - Indicates the serial number of the certificate.

Signature Alg - Indicates the signature algorithm of the certificate.

Public Key Alg - Indicates the public key algorithm.

Not Valid Before - Indicates the date the certificate became active.

Not Valid After - Indicates the date the certificate expires. Self-signed certificates are valid for two years.

Subject - Lists the Subject Distinguished Name.

Issuer - Lists the Distinguished Name of the issuer of the certificate.

MD5 Fingerprint - Indicates the MD5 algorithm. This value provides a unique sequence for each certificate consisting of 16 bytes.

SHA1 Fingerprint - Indicates the Secure Hash Algorithm. This value provides a unique sequence for each certificate consisting of 20 bytes.

Generating an Enrollment Request

To request a certificate from a public or private CA, you need to generate an enrollment request. The request contains the public key for the certificate, a distinguished name, and other information. The request is stored as a file with a ".csr" extension in the AT-S62 file system, from where you can upload it onto your management workstation or FTP server for submission to the CA.

Please review the following before performing the procedure:

- You must generate an encryption key pair before creating an enrollment request. For instructions, refer to "Creating an Encryption Key" on page 695.
- During this procedure you are prompted to enter the ID number of the encryption key pair to be used to create the enrollment request. If you do not know the ID number, refer to "Creating an Encryption Key" on page 695 to view the key ID numbers.
- For a review of all the steps to creating an enrollment request and downloading a CA certificate onto a switch, refer to "General Steps for a Public or Private CA Certificate" on page 685.

To generate an enrollment request, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **7** to select the Keys/Certificates Configuration menu.
3. From the Keys/Certificates Configuration menu, type **1** to select Switch Distinguished Name (DN) and, when prompted, enter a name. An enrollment request must have a distinguished name. For information, refer to "Distinguished Names" on page 707.
4. From the Keys/Certificates Configuration menu, type **3** to select Public Key Infrastructure (PKI) Configuration.

   The Public Key Infrastructure (PKI) Configuration menu is shown in Figure 243 on page 719.

5. From the Public Key Infrastructure (PKI) Configuration Menu, type **3** to generate an enrollment request. The Generate Enrollment Request Menu is shown in Figure 250.

```
        Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                        Production Switch
 User: Manager                              11:20:02 02-Jan-2006
                  Generate Enrollment Request Menu
1 - Request Name.....................
2 - KeyPair ID ..................... 0
3 - Format ......................... PEM
4 - Type ........................... PKCS10
5 - Generate Enrollment Request

R - Return to Previous Menu

Enter your selection?
```

Figure 250. Generate Enrollment Request Menu

6. Type **1** to select Request Name.

   The following message is displayed:

   ```
   Enter enrollment request name (24 chars max) ->
   ```

7. Enter a name of up to 24 alphanumeric characters for the enrollment request. Spaces are allowed.

   The name is used to create the filename of the enrollment request when it is stored in the AT-S62 file system. The full filename consists of the enrollment request name followed by ".csr" extension, which the management software adds automatically. For example, if you enter "certificate75" as the enrollment request name, the enrollment request's filename will be "certificate75.csr".

8. Type **2** to select KeyPair ID.

   The following message is displayed:

   ```
   Enter keypair ID -> [0 to 65535] -> 0
   ```

9. Enter the ID number of the encryption key you want to use to create the enrollment request. The encryption key must already exist on the switch. (If you have forgotten the key ID number, return to the Key Management menu to view the keys on the switch.) The value can be from 0 to 65,535.

10. Type **3** to toggle the Format option between the following values:

    DER - Creates the certificate in binary format. This is the default.

PEM - Creates the certificate in the Privacy Enhanced Mail (PEM) format, which is an ASCII format.

Note
Option 4, Type, cannot be changed. The PKCS10 value indicates the internal format of an enrollment request.

11. Type **5** to select Generate Enrollment Request.

    Once the switch has finished generating the request, you will see a message similar to the following.

    ```
    Enrollment request is being generated. Please wait
    ...Done.
    Enrollment Request available in file [Switch 12.csr].
    Press any key to continue ...
    ```

    The enrollment request is now stored in the AT-S62 file system. To see the file, refer to “Displaying System Files” on page 185.

12. Press any key to return to the Public Key Infrastructure (PKI) Configuration menu.

13. To submit the request to a CA, upload the enrollment request from the file system on the switch to your management workstation or to an FTP server on your network. For instructions, refer to “Uploading a System File” on page 209. After receiving the certificates from the CA, refer to “Installing CA Certificates onto a Switch” on page 733 for an overview of the procedures to loading the certificates onto the switch.

    When submitting an enrollment request, be sure to follow the rules and guidelines of the CA. Failure to follow their guidelines may delay the issuing of the certificate.

Installing CA Certificates onto a Switch

This section lists the procedures for installing a certificate created by a public or private CA onto the switch. It should be noted that a CA generated certificate will consist of several certificates, with a minimum of two. All the certificates from the CA must be installed on the switch.

Note
A CA certificate can only be used on the switch where you created its certificate request. The certificate will not work on any other switch.

To install CA certificates on a switch, perform the following procedure:

1. Download the certificates from your management workstation or FTP server to the AT-S62 file system on the switch. For instructions, refer to "Downloading a System File" on page 202.

2. Load the certificates into the certificate database. For instructions, refer to "Adding a Certificate to the Database" on page 722.

3. Activate HTTPS on the switch by configuring the web server and specifying the key pair used to create the enrollment request as the active key pair. For instructions, refer to "Configuring the Web Server" on page 683.

Configuring PKI

Option 1 - Maximum Number of Certificates in the Public Key Infrastructure (PKI) Configuration menu controls the maximum number of certificates you can add to the certificate database. The range is 12 to 256. The default value is 256. There should be little cause or need for you to adjust this value. To display the Public Key Infrastructure (PKI) Configuration menu, perform steps 1 to 3 of the procedure “Creating a Self-signed Certificate” on page 718.

Configuring SSL

To configure the SSL protocol, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **9** to select Secure Socket Layer (SSL).

   The Secure Socket Layer (SSL) menu is shown in Figure 251.

```
        Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006
                      Secure Socket Layer (SSL)
1 - Maximum Number of Sessions......... 50
2 - Session Cache Timeout.............. 300 seconds

R - Return to Previous Menu

Enter your selection?
```

Figure 251. Secure Socket Layer (SSL) Menu

3. Select **1** - Maximum number of Sessions to increase the number of sessions.

   Enter a value from 1 to 100. The maximum number of sessions is used to speed up a connection. By increasing the number of sessions, you increase HTTPS performance. However, increasing the number of sessions also increases the memory requirements. The default is 50.

4. Select **2** - Session Cache Timeout to increase or decrease the timer that determines when the session cache times out.
5. Enter a value, in seconds, from 1 to 600. The default is 300 seconds.
6. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Chapter 33
Secure Shell (SSH) Protocol

The chapter contains overview information about the Secure Shell (SSH) protocol and the procedure for configuring this protocol on a switch from a local or Telnet management session. It contains the following sections:

SSH Overview

Secure management is increasingly important in modern networks, as the ability to easily and effectively manage switches and the requirement for security are two universal requirements. Switches are often remotely managed using remote sessions via the Telnet protocol. This method, however, has a serious security problem—it is only protected by plaintext usernames and passwords which are vulnerable to wiretapping and password guessing.

The Secure Shell (SSH) protocol provides encrypted and strongly authenticated remote login sessions, similar to the Telnet and rlogin protocols, between a host running a Secure Shell server and a machine with a Secure Shell client.

The AT-S62 management software features Secure Shell server software to enable network managers to securely manage the switch over an insecure network. It offers the benefit of cryptographic authentication and encryption. Secure Shell can replace Telnet for remote management sessions.

Support for SSH

The AT-S62 management software implementation of the SSH protocol is compliant with SSH1 (versions 1.3 and 1.5) and SSH2 (version 2.0).

In addition, the following SSH options and features are supported:

- Inbound SSH connections (server mode) is supported.
- The following security algorithms are supported:
  - 128-bit Advanced Encryption Standard (AES), 192-bit AES, and 256-bit AES
  - Arcfour (RC4) security algorithm is supported.
  - Triple-DES (3DES) encryption for SSH sessions is supported.
- RSA public keys with lengths of 512 to 2048 bits are supported. Keys are stored in a format compatible with other Secure Shell implementations, and mechanisms are provided to copy public keys to and from the switch.
- Compression of SSH traffic.

The following SSH options and features are not supported:

- IDEA or Blowfish encryption
- Nonencrypted Secure Shell sessions

- Tunnelling of TCP/IP traffic

> **Note**
> Non-encrypted Secure Shell sessions serve no purpose.

SSH Server

The AT-S62 management software includes SSH server software. When the SSH server is activated, your remote management sessions of the switch from a management station that has SSH client software will be encrypted.

> **Note**
> If your switch is in a network protected by a firewall, you may need to configure the firewall to permit SSH connections.

An SSH management session uses the same usernames and passwords as the other types of switch management sessions. You can log in using the default manager or operator login account, or as a user configured with the RADIUS and TACACS+ protocols, as explained in Chapter 34, "TACACS+ and RADIUS Authentication Protocols" on page 747.

The Secure Shell server requires two encryption key pairs. The first, called the *host key*, is the switch's own RSA key. The recommended length of this key is 1024 bits. The second key, the *server key,* is used by the SSH server software on the switch. If desired, you can configure the switch to periodically re-generate this key. The two keys cannot be of the same length. For the procedure for creating an encryption key, see "Creating an Encryption Key" on page 695.

For information on how to create an encryption key, see "Creating an Encryption Key" on page 695.

SSH Clients

The SSH protocol provides a secure connection between the switch and SSH clients. Once you have configured the SSH server, you need to install SSH client software on your management workstation. The AT-S62 software supports both SSH1 and SSH2 clients.

You can download client software from the Internet. Two popular SSH clients are PuTTY and CYGWIN. To install SSH client software, follow the directions from the vendor.

After you have installed the SSH client software on your workstation and configured the server software on the switch, you can use the client software to login to the switch for an encrypted SSH management session. The SSH server can support up to one manager session and eight operator sessions at one time.

SSH and Enhanced Stacking

The AT-S62 management software allows for encrypted SSH management sessions between a management workstation and a master switch of an enhanced stack, but not with slave switches, as explained in this section.

When you remotely manage a slave switch, all management communications are conducted through the master switch using the enhanced stacking feature. Management packets from your workstation are first directed to the master switch before being forwarded to the slave switch. The reverse is true as well. Management packets from a slave switch first pass through the master switch before reaching your management workstation.

Enhanced stacking uses a proprietary protocol that does not support encryption between a master switch and a slave switch. SSH encryption only occurs between your workstation and the master switch, not between your workstation and a slave switch.

This is illustrated in Figure 252. The figure shows an SSH management session between a remote management workstation and a slave switch of an enhanced stack. The packets exchanged between the slave switch and the master switch are transmitted in plaintext and those exchanged between the master switch and the SSH management workstation are encrypted.

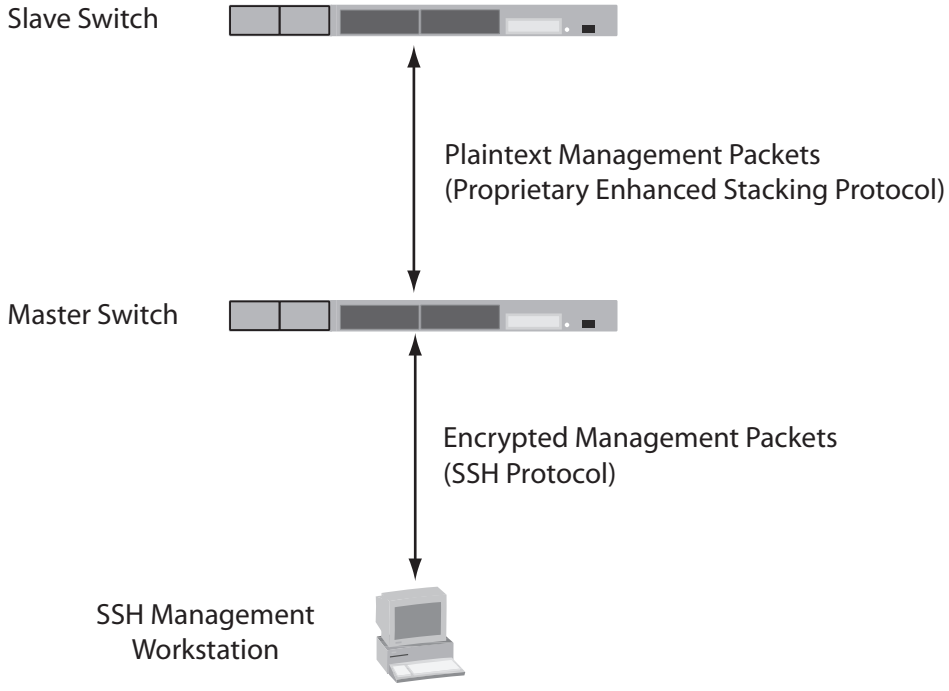

Figure 252. SSH Remote Management of a Slave Switch

Since enhanced stacking does not allow for SSH encrypted management sessions between a management station and a slave switch, you configure SSH only on the master switch of a stack. Activating SSH on a slave switch has no affect.

Guidelines

Below are the guidelines to observe when configuring SSH:

- SSH requires two encryption key pairs. One key pair will function as the host key and the other the server key. For instructions on creating keys, refer to "Creating an Encryption Key" on page 695.
- The two encryption key pairs must be of different lengths of at least one increment (256 bits) apart. The recommended bit size for a server key is 768 bits. The recommended size for the host key is 1024 bits.
- You activate and configure SSH on the master switch of an enhanced stack, not on slave switches.
- The AT-S62 software uses well-known port 22 as the SSH default port.

General Steps to Configuring SSH

Configuring the SSH server involves several procedures. This section lists the procedures you need to complete to configure the SSH feature.

1. Create two encryption key pairs on the master switch of the enhanced switch. One pair will function as the host key and the other the server key.

2. Configure and activate the Secure Shell server on the switch by specifying the two encryption keys in the server software.

   For instructions, see "Configuring the SSH Server" on page 742.

3. Install SSH client software on your management workstation.

   Follow the directions provided with the client software. You can download SSH client software from the Internet. Two popular SSH clients are PuTTY and CYGWIN.

4. Disable the Telnet server.

   Although the switch allows the SSH and Telnet servers to be enabled simultaneously, allowing Telnet to be enabled negates the security of the SSH feature. To disable the Telnet server, see "Enabling or Disabling the Telnet Server" on page 67.

5. Logon to the switch from your SSH management workstation.

Configuring the SSH Server

This section describes how to configure the SSH server software on the switch. For a description of all the steps required to configure an SSH server, see “General Steps to Configuring SSH” on page 741.

This procedure assumes that you have already created the two key pairs. If you have not created the keys, go to “Creating an Encryption Key” on page 695.

While you are configuring the SSH feature, you must disable the SSH server. When you have completed your configuration changes, enable the SSH server to permit SSH client connections.

Note
Allied Telesyn recommends disabling the Telnet server before activating SSH. Otherwise, the security functions provided by SSH are lost. See “Enabling or Disabling the Telnet Server” on page 67.

To configure the SSH server software on the switch, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **8** to select Secure Shell (SSH).

   The Secure Shell (SSH) Menu is shown in Figure 253.

```
        Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                          Production Switch
User: Manager                                  11:20:02 02-Jan-2006
                         Secure Shell (SSH)
1 - SSH Server Status ....... Disabled
2 - Host Key ID.............. <Not Defined>
3 - Server Key ID ........... <Not Defined>
4 - Server Key Expiry Time .. 0 hours
5 - Login Timeout ........... 180 seconds
6 - Show Server Information

R - Return to Previous Menu

Enter your selection?
```

Figure 253. Secure Shell (SSH) Menu

3. Select **1** - SSH Server Status to enable or disable the SSH server.

4. Choose from one of the following:

   Disabled - While you are configuring SSH, you must set this field to Disabled. This is the default.

   Enabled - Select this value to enable the SSH server. Select this value after you have finished configuring SSH and want to log on to the server.

   Note
   You cannot disable the SSH server when there is an active SSH connection. Otherwise, you receive a warning message.

5. Type **2** to select Host Key ID and, when prompted, enter the key ID of the key pair which will act as the SSH host key. If you have forgotten the key ID, refer to "Creating an Encryption Key" on page 695.

6. Type **3** to select Server Key ID and enter the ID of the key pair which will act as the SSH server key. If you have forgotten the key ID, refer to "Creating an Encryption Key" on page 695.

7. Type **4** to select Server Key Expiry Time to set the time, in hours, for the server key to expire.

   This timer determines how often the switch generates a new server key. A server key is regenerated for security purposes. A server key is only valid for the time period configured in the Server Key Expiry (Expiration) Time timer. Allied Telesyn recommends you set this field to 1. With this setting, a new key is generated every hour.

   The default is 0 hours which means the server key never expires. The range is 0 to 5 hours.

8. Select **5** and enter a value for Login Timeout.

   This is the time it takes to release the SSH server from an incomplete SSH client connection. Enter a time in seconds. The default is 180 seconds (3 minutes). The range is 60 to 600 seconds.

9. Select **1** to toggle SSH Server Status to **Enable**.

   Note
   Allied Telesyn recommends disabling the Telnet server before you enable SSH. Otherwise, the security provided by SSH is lost.

10. After making changes, type **R** to until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying SSH Information

To display SSH server information, perform the following procedure:

1. From the Main Menu, type **7** to select Security and Services.
2. From the Security and Services menu, type **8** to select Secure Shell (SSH).

   The Secure Shell (SSH) Menu is shown in Figure 253 on page 742.
3. From the Secure Shell (SSH) menu, select **6** - Show Server information to display the SSH Server data.

   The Show Server Information Menu is shown in Figure 254.

```
            Allied Telesyn Ethernet Switch - AT-8524M - AT-S62
                           Production Switch
User: Manager                                          11:20:02 02-Jan-2006
                        Show Server Information Menu
Versions Supported ....... 1.3, 1.5, 2.0
Server Status ............ Enabled
Server Port .............. 22
Host Key ID .............. 200
Host Key Bits ............ 1024
Server Key ID ............ 250
Server Key Bits .......... 768
Server Key Expiry ........ 0 hours
Login Timeout ............ 180 seconds
Authentication Available . Password
Ciphers Available ........ 3DES, 128 bit AES, 192 bit AES,256 bit AES,
                           Arcfour (RC4)
MACs Available ........... hmac-sha1, hmac-md5
Data Compression ......... Available

R - Return to Previous Menu

Enter your selection?
```

Figure 254. Show Server Information Menu

The following information is displayed:

- Versions Supported: Indicates the versions of SSH which are supported by the AT-S62 software.
- Server Status: Indicates whether or not the SSH server is enabled or disabled.
- Server Port: Indicates the well-known port for SSH. The default is port 22.

- Host Key ID: Indicates the host key ID defined for SSH.
- Host Key Bits: Indicates the number of bits in the host key.
- Server Key ID: Indicates the server key ID defined for SSH.
- Server Key Bits: Indicates the number of bits in the server key.
- Server Key Expiry: Indicates the length of time, in hours, until the server key is regenerated. The default is 0 hours which means the server key is not regenerated.
- Login Timeout: Indicates the time, in seconds, until a SSH server is released from an incomplete connection with a SSH client.
- Authentication Available: Indicates the authentication method available. Currently, password authentication is the only supported method.
- Ciphers Available: Indicates the SSH ciphers that are available on the switch.
- MACs Available: Indicates the Message Authorization Code (MAC) that is used to validate incoming SSH messages to the server. Two algorithms are supported.
- Data Compression: Indicates whether or not data compression is available on the switch. Data compression is useful for networks that have a slow throughput speed.

Chapter 34

TACACS+ and RADIUS Authentication Protocols

This chapter explains how to configure the parameter settings for the two authentication protocols TACACS+ and RADIUS. Sections in the chapter include:

TACACS+ and RADIUS Overview

TACACS+ and RADIUS are authentication protocols for enhancing the security of your network. (TACACS+ is an acronym for Terminal Access Controller Access Control System. RADIUS is an acronym for Remote Authentication Dial In User Services.) In general terms, these authentication protocols transfer the task of authenticating network access from a network device to an authentication protocol server.

The AT-S62 software comes with TACACS+ and RADIUS client software. You can use the client software to add two security features to the switch. The first feature, described in this chapter, involves creating new manager accounts for controlling who can log onto a switch to change its parameter settings. The second feature is 802.1x Port-based Access Control, explained in Chapter 29, “802.1x Port-based Network Access Control” on page 643, which controls which end users and end nodes can send packets through the switch.

This chapter explains the manager accounts feature. The AT-S62 software has two standard manager login accounts: Manager and Operator. The Manager account lets you change a switch’s parameter settings while the Operator account lets you view the settings, but not change them. Each account has its own password.

For networks managed by just one or two network managers, the standard accounts may be all you need. However, for larger networks managed by several network managers, you might want to give each manager his or her own management login account rather than have them share an account.

This is where TACACS+ and RADIUS can be useful. You can use them to create additional manager accounts and transfer the task of validating management access from the switch to an authentication protocol server. You use the protocols to create a series of username and password combinations that define who can manage an AT-8500 Series switch.

There are three basic functions an authentication protocol provides:

- Authentication
- Authorization
- Accounting

When a network manager logs in to a switch to manage the device, the switch passes the username and password entered by the manager to the authentication protocol server. The server checks to see if the username and password are valid for that switch. This is referred to as authentication.

If the combination is valid, the authentication protocol server notifies the switch and the switch completes the login process, allowing the manager to manage the switch.

If the username and password are invalid, the authentication protocol server notifies the switch and the switch cancels the login.

Authorization defines what a manager can do once logged in to a switch. You assign an authorization level to each username and password combination that you create on the server software. The access level can either Manager or Operator.

The final function of an authentication protocol is accounting, which is used to keep track of user activity on network devices. The AT-S62 management software does not support RADIUS or TACACS+ accounting as part of new manager accounts. However, it does support RADIUS accounting with the 802.1x port-based access control feature, explained in Chapter 29, "802.1x Port-based Network Access Control" on page 643.

Note
The AT-S62 management software does not support the two earlier versions of the TACACS+ protocol, TACACS and XTACACS.

Guidelines

Here are the main points to using the RADIUS and TACACS+ protocols.

- First, you need to install TACACS+ or RADIUS server software on one or more of your network servers or management stations. Authentication protocol server software is not available from Allied Telesyn.

  Note
  The switch communicates with the authentication server via the switch's management VLAN. Consequently, the node functioning as the authentication server must be communicating with a switch through a port that is a member of that VLAN. The default management VLAN is Default_VLAN. For further information, refer to "Specifying a Management VLAN" on page 579.

- The authentication protocol server can be on the same subnet or a different subnet as the AT-8500 Series switch. If the server and switch are on different subnets, be sure to specify a default gateway in the Administration Menu so that the switch and server can communicate with each other.
- You need to configure the TACACS+ or RADIUS software on the authentication server. This involves the following:
  - Specifying the username and password combinations. A username can contain up to 30 alphanumeric characters and a

password up to 16 characters. Spaces are allowed in both a username and password, but special characters, such as asterisks and exclamation points, should be avoided.

– Assigning each combination an authorization level. How this is achieved differs depending on the server software you are using. TACACS+ controls this through the sixteen (0 to 15) different levels of the Privilege attribute. A privilege level of “0” gives the combination Operator status. Any value from 1 to 15 gives the combination Manager status.

  For RADIUS, management level is controlled by the Service Type attribute. This attribute has 11 different values, of which only two apply to the AT-S62 management software. A value of Administrative for this attribute gives the username and password combination Manager access. A value of NAS Prompt assigns the combination Operator status.

Note
This manual does not explain how to configure TACACS+ or RADIUS server software. For that you need to refer to the documentation that came with the software.

❒ You must activate the TACACS+ or RADIUS client software on the switch using the AT-S62 software and configure the settings, which includes the IP addresses of up to three authentication server. The procedure for this step is found in this chapter.

By default, authentication protocol is disabled in the AT-S62 software. Before activating it, you need to provide the following information:

❒ Which authentication protocol, TACACS+ or RADIUS, you want to use. Only one authentication protocol can be active on a switch at a time.

❒ IP addresses of up to three authentication servers.

❒ The encryption key used by the authentication servers.

You can specify up to three TACACS+ or RADIUS servers. Specifying multiple servers adds redundancy to your network. For example, removing an authentication server from the network for maintenance will not prevent network managers from logging into switches if there are one or two other authentication servers on the network.

When a switch receives a username and password combination from a network manager, it sends the combination to the first authentication server in its list. If the server fails to respond, the switch sends the combination to the next server in the list, and so on.

If no authentication server responds or if no servers have been defined and you are managing the switch locally, the management software

defaults to the standard manager and operator accounts.

Note
For more information on TACACS+, refer to the RFC 1492 standard. For more information on RADIUS, refer to the RFC 2865 standard.

Configuring TACACS+ Authentication Protocol Settings

To configure the TACACS+ settings on the switch, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **6** to select Authentication Configuration.

   The Authentication Configuration menu is shown in Figure 255.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                  Production Switch
User: Manager                      11:20:02 02-Jan-2006
                  Authentication Menu
1 - Server-based Authentication ..... Disabled
2 - Authentication Method ........... TACACS+
3 - TACACS+ Configuration
4 - RADIUS Configuration
5 - Passwords Configuration

R - Return to Previous Menu

Enter your selection?
```

Figure 255. Authentication Configuration Menu

Note
Option 1 - Server-based Authentication in the menu applies only to the manager accounts feature described in this chapter. This menu selection has no affect on the 802.1x port-based access control feature described in Chapter 29, “802.1x Port-based Network Access Control” on page 643. When Option 1 is set to disabled, the default setting, the switch uses the default manager and operator accounts. When set to enabled, the switch seeks its manager accounts on a TACACS+ or RADIUS authentication server.

3. To select the active authentication protocol, type **2** to select Authentication Method. The following prompt is displayed:

   ```
   Enter T-TACACS+, R-RADIUS ->
   ```

4. Type **T** to select TACACS+. The default is TACACS+. Only one protocol can be active on the switch at a time.

5. To configure the TACACS+ parameters, type **3** to select TACACS+ Configuration.

   The TACACS+ Client Configuration menu is shown in Figure 256.

```
   Allied Telesyn Ethernet Switch AT-8524M - AT-S62
User: Manager                      11:20:02 02-Jan-2006
              TACACS+ Client Configuration

1 - TAC Server 1 ................... 0.0.0.0
2 - TAC Server 2 ................... 0.0.0.0
3 - TAC Server 3 ................... 0.0.0.0
4 - TAC Global Secret ..............
5 - TAC Timeout ..................... 10 seconds

R - Return to Previous Menu

Enter your selection?
```

Figure 256. TACACS+ Client Configuration Menu

6. Configure the settings as needed. The settings are described below:

   1 - TAC Server 1
   2 - TAC Server 2
   3 - TAC Server 3
   Use these parameters to specify the IP addresses of up to three network servers containing TACACS+ server software. After you have entered an IP address, you will see the following prompt:

   ```
   Use per-server secret [Y/N] ->
   ```

   If you will be specifying more than one TACACS+ server and if all of the servers use the same encryption secret, you can answer No to this prompt and enter the encryption secret using the TAC Global Secret parameter.

   However, if you are specifying only one TACACS+ server or if the servers have difference encryption secrets, then respond with Yes to this prompt. You will see:

   ```
   Enter per-server secret [max 40 characters] ->
   ```

   Use this prompt to enter the encryption secret for the TACACS+ server whose IP address you are specifying.

   4 - TAC Global Secret
   If all of the TACACS+ servers have the same encryption secret, rather then entering the same secret when you enter the IP addresses, you can use this option to enter the secret just once.

5 - TAC Timeout
This parameter specifies the maximum amount of time the switch waits for a response from a TACACS+ server before assuming the server will not respond. If the timeout expires and the server has not responded, the switch queries the next TACACS+ server in the list. If there are not any more servers, the switch defaults to the standard Manager and Operator accounts. The default is 10 seconds. The range is 1 to 60 seconds.

7. After configuring the parameters in the TACACS+ Client Configuration menu, type **R** to return to the Authentication Menu, shown in Figure 255 on page 752.

Note
Steps 8 and 9 activate the feature. Do not activate the feature if you have not defined the new manager username and password accounts on the TACACS+ authentication server. Otherwise, you will not be able to initiate future management sessions with the switch.

8. From the Authentication Menu, type **1** to select Server-based Authentication. The following prompt is displayed:

```
Server Based User Authentication (E-Enabled, D-Disabled)
->
```

9. Type **E** to enable server-based authentication on the switch.

   After enabling the feature, you must use the new manager username and password combinations that you defined on the TACACS+ authentication server when you start future management sessions on the switch.

10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Configuring RADIUS Authentication Protocol Settings

To configure the RADIUS settings on the switch, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **6** to select Authentication Configuration.

   The Authentication Menu is shown in Figure 255 on page 752.

   Note
   Option 1 - Server-based Authentication in the menu applies only to the manager accounts feature described in this chapter. This menu selection has no affect on the 802.1x port-based access control feature described in Chapter 29, "802.1x Port-based Network Access Control" on page 643. When Option 1 is set to disabled, the default setting, the switch uses the default manager and operator accounts. When set to enabled, the switch seeks its manager accounts on a TACACS+ or RADIUS authentication server.

3. To select the active authentication protocol, type **2** to select Authentication Method. The following prompt is displayed:

   ```
   Enter T-TACACS+, R-RADIUS ->
   ```

4. Type **R** for RADIUS. Only one protocol can be active on the switch at a time.

5. To configure the RADIUS parameters, type **4** to select RADIUS Configuration.

The RADIUS Client Configuration menu is shown in Figure 257.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                    Production Switch
User: Manager                         11:20:02 02-Jan-2006
                RADIUS Client Configuration

1 - Global Encryption Key ............. ATI
2 - Global Server Timeout period....... 10 second(s)
3 - RADIUS Server 1 Configuration ..... 0.0.0.0
4 - RADIUS Server 2 Configuration ..... 0.0.0.0
5 - RADIUS Server 3 Configuration ..... 0.0.0.0
6 - Show Status

R - Return to Previous Menu

Enter your selection?
```

Figure 257. RADIUS Client Configuration

6. Configure the parameters as needed. The parameters are defined here:

   Global Encryption Key
   This parameter specifies the encryption key for the RADIUS servers. This option is useful if you will be entering more than one RADIUS server and all the servers share the same encryption key. The default is ATI.

   Global Server Timeout period
   This parameter specifies the maximum amount of time the switch will wait for a response from a RADIUS server before assuming that the server will not respond. If the timeout expires and the server has not responded, the switch queries the next RADIUS server in the list. If there are not any more servers, than the switch will default to the standard Manager and Operator accounts. The default is 10 seconds. The range is 1 to 60 seconds.

   3 - RADIUS Server 1 Configuration
   4 - RADIUS Server 1 Configuration
   5 - RADIUS Server 1 Configuration
   Use these parameters to specify the IP addresses of up to three network servers containing the RADIUS server software. Selecting one of the options displays the RADIUS Server Configuration menu, shown in Figure 258 on page 757.

```
       Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Production Switch
User: Manager                                11:20:02 02-Jan-2006
                  RADIUS Server 1 Configuration
1 - Server IP Address ................. 0.0.0.0
2 - Server Authentication UDP Port .... 1812
3 - Server Encryption Key ............. <Not Defined>

R - Return to Previous Menu

Enter your selection?
```

Figure 258. RADIUS Server Configuration

The parameters are described below:

1 - Server IP Address
Use this option to specify the IP address of the RADIUS server.

2 - Server Authentication UDP Port
Use this option to specify the UDP port of the RADIUS protocol.

3 - Server Encryption Key
Use this option to specify the encryption key for the RADIUS server.

7. After you finish configuring the parameters in the RADIUS Client Configuration menu, type **R** to return to the Authentication Menu, shown in Figure 255 on page 752.

> **Note**
> Steps 8 and 9 activate the feature. Do not enable the feature if you have not defined the new manager username and password accounts on the RADIUS authentication server. Otherwise, you will not be able to initiate future management sessions with the switch.

8. From the Authentication Menu, type **1** to select Server-based Authentication. The following prompt is displayed:

```
Server Based User Authentication (E-Enabled, D-Disabled) ->
```

9. Type **E** to enable server-based authentication on the switch.

   After enabling the feature, you must use the new manager username and password combinations that you defined on the RADIUS authentication server when you start future management sessions on the switch.

10. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Displaying RADIUS Status and Settings

The RADIUS Client Configuration menu shown in Figure 257 on page 756 has a selection that displays the RADIUS client software settings. The selection, 6 - Show Status, displays the Show Status menu, as shown in Figure 259

```
            Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                          Production Switch
User: Manager                                      11:20:02 02-Jan-2006
                                Show Status
Global Configuration
--------------------
Encryption Key       : ATI
Server Timeout       : 10 second(s)

Server IP Address    Auth Port    Encryption Key   Auth Req   Auth Resp
-----------------------------------------------------------------------
149.11.11.11         1812         WRRT             100        96
149.22.22.22         1812         LLST             4          4
149.22.22.22         1812         OORT             0          0

U - Update Display
R - Return to Previous Menu

Enter your selection?
```

Figure 259. Show Status Menu

The information in this menu is for viewing purposes only. Most of the columns are self-explanatory, with the possible exceptions of “Auth Req” and “Auth Resp.” The “Auth Req” column displays the number of authentication requests the switch has made to the RADIUS server. The “Auth Resp” shows the number of responses that the switch has received back from the server.

Chapter 35

Management Access Control List

This chapter explains how to create an access control list (ACL) to restrict Telnet and web browser management access to the switch. Sections in this chapter include:

Management ACL Security Overview

This chapter explains how to restrict remote management access to a switch by creating a management access control list (management ACL). This feature controls which management stations can remotely manage the device using the Telnet application protocol or a web browser.

The switch uses the management ACL to filter the management packets that it receives, accepting and processing only those management packets that meet the criteria stated in the ACL. Those management packets that do not meet the criteria are discarded.

The benefit of this feature is that you can prevent unauthorized access to the switch by controlling which workstations are to have remote management access. You can even control which method, Telnet or web browser, that a remote manager can use.

For example, you can create a management ACL that allows the switch to accept management packets only from the management stations in one subnet or from just one or two specific management stations.

An access control list (ACL) is a list of one or more statements that define which management packets the switch accepts. Each statement, referred to as an access control entry (ACE), contains criteria that the switch uses in making the determination.

An ACE in a management ACL is an implicit “permit” statement. A management packet that meets the criteria of an ACE is processed by the switch. Consequently, the ACEs that you enter into the management ACL should specify which management packets you want the switch to process. Packets that do not meet any of the ACEs in the management ACL are discarded.

Parts of a Management ACE

An ACE has the following three parts:

- IP address
- Subnet mask
- Application

IP Address

You can specify the IP address of a specific management station or a subnet.

Mask

You need to enter a mask that indicates the parts of the IP address the switch should filter on. A binary “1” indicates the switch should filter on the

corresponding bit of the address, while a "0" indicates that it should not. If you are filtering on a specific IP address, use the mask 255.255.255.255. If you are filtering on a subnet, the mask would depend on the address. For example, to allow all management stations in the subnet 149.11.11.0 to manage the switch, you would enter the mask 255.255.255.0.

Application

The application parameter allows you control whether the remote management station can manage the switch using Telnet, a web browser, or both. For example, you might create an ACE that states that a particular remote management station can only use a web browser to manage the switch. You can also use this option to control whether the management station can ping the switch.

Management ACL Guidelines

Below are guidelines to observe when you create a management ACL:

- The default setting for this feature is disabled.
- A switch has only one management ACL.
- A management ACL can have up to 256 ACEs.
- An ACE must have an IP address and mask.
- All management ACEs are implicit "permit" statements. A management packet that meets the criteria of an ACE is accepted by the switch. Consequently, the ACEs you enter into the management ACL should specify which management packets you want the switch to process. Management packets that do not meet any of the ACEs in the management ACL are discarded.
- A management packet that meets an ACE is immediately processed by the switch and is not compared against any remaining ACEs in the management ACL.
- The ACEs are performed in the order of their identification number, starting with 1.
- The management ACL does not control local management or remote SSH or SNMP management of a switch.
- Activating this feature without specifying any ACEs prohibits you from managing the switch remotely using a Telnet application or web browser because the switch discards all Telnet and web browser management packets.
- You can apply management ACLs to both master and slave switches in an enhanced stack. A management ACL on a master switch filters management packets intended for the master switch as well as those intended for any slave switches that you manage through the master switch. A management ACL applied to a slave switch filters only those management packets directed to the slave switch.

Examples

Following are several examples of ACEs.

This ACE allows the management station with the IP address 149.11.11.11 to remotely manage the switch using either the Telnet application protocol or a web browser, and to ping the device:

IP Address:	149.11.11.11
Mask:	255.255.255.255
Application Type:	All

If the management ACL contained only the above ACE, then only that management station would be allowed to remotely manage the switch.

This ACE allows all management stations in the subnet 149.11.11.0 to remotely manage the switch using either the Telnet application or a web browser, and to ping the device:

IP Address:	149.11.11.0
Mask:	255.255.255.0
Application Type:	All

This ACE allows all management stations in the subnet 149.11.11.0 to remotely manage the switch using a web browser and to ping the device, but prevents management with the Telnet application:

IP Address:	149.11.11.0
Mask:	255.255.255.0
Application Type:	Web, Ping

A management ACL can contain multiple ACEs. The two ACEs in this ACL allow all management packets from the subnets 149.11.11.0 and 149.22.22.0 to manage the switch using just the Telnet application. They cannot use a web browser and they cannot ping the device:

ACE #1

IP Address:	149.11.11.0
Subnet Mask:	255.255.255.0
Application Type:	Telnet

ACE #2

IP Address:	149.22.22.0
Subnet Mask:	255.255.255.0
Application Type:	Telnet

The two ACEs in this management ACL permit remote management from the management station with the IP address 149.11.11.11 and all management stations in the subnet 149.22.22.0:

ACE #1

IP Address:	149.11.11.11
Mask:	255.255.255.255
Application Type:	All

ACE #2

IP Address:	149.22.22.0
Mask:	255.255.255.0
Application Type:	All

This example allows the management station with the IP address 149.11.11.4 to ping the device, but not manage it:

IP Address:	149.11.11.4
Mask:	255.255.255.255
Application Type:	Ping

Enabling or Disabling the Management ACL

This procedure enables and disables the management ACL. When enabled, only those management stations specified in the ACL are allowed to manage the switch remotely using the Telnet application protocol or a web browser. When the feature is disabled, the management software on the switch can be accessed remotely from any management workstation.

> **Note**
> Do not activate the management ACL from a remote management session until you have specified the access control entries (ACEs). Otherwise, the switch will discard all remote management packets, making it impossible for you to remotely manage the unit from a Telnet or web browser management session. For instructions on how to add ACEs, refer to “Creating an ACE” on page 766.

To enable or disable the Management ACL, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **7** to select Management ACL.

   The Management ACL Configuration menu is shown in Figure 260.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Marketing
User: Manager                           11:20:02 02-Mar-2006
                Management ACL Configuration

Configuring Management ACL
1 - Management ACL Status ..... Disabled
2 - Create Management ACL Entry
3 - Modify Management ACL Entry
4 - Delete Management ACL Entry
5 - Display All Management ACL Entries

R - Return to Previous Menu

Enter your selection?
```

Figure 260. Management ACL Configuration Menu

3. Type **1** to select Management ACL Status and toggle the selection to either Enabled or Disabled. The default setting is disabled.

A change to the status of the management ACL is immediately activated on the switch.

Note
If you activate this feature from a Telnet management session, your management session will end and you will not be able to reestablish it should the management ACL not contain an ACE with the IP address or subnet address of your management workstation.

4. After making changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Creating an ACE

To create a new ACE in the management ACL, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **7** to select Management ACL.

   The Management ACL Configuration menu is shown in Figure 260 on page 764.

3. From the Management ACL Configuration menu, type **2** to select Create Management ACL Entry.

   The following prompt is displayed:

   ```
   Enter the entry ID :  [1 to 256] -> 1
   ```

4. Enter an identification number for the access control entry. Every ACE must have a unique number. The range is 1 to 256.

   The following prompt is displayed:

   ```
   Enter the IP address:
   ```

5. Enter the IP address of a specific management station (for example, 149.11.11.11) or a subnet (for example, 149.11.11.0).

   The following prompt is displayed:

   ```
   Enter the Mask:
   ```

6. Enter a mask that indicates the parts of the IP address the switch should filter on. A binary “1” indicates the switch should filter on the corresponding bit of the address, while a “0” indicates it should not. If you are filtering on a specific IP address, enter the mask 255.255.255.255. If you are filtering on a subnet, the mask will depend on the subnet. For example, to allow all management stations in the subnet 149.11.11.0 to manage the switch, you would enter the mask 255.255.255.0.

   The following prompt is displayed:

   ```
   Enter the Application Type [TELNET, WEB, PING, ALL]:
   ```

7. Specify which application you want the management station to be able to use when managing the switch. The options are:

- Telnet - Permits Telnet management.
- Web - Permits web browser management.
- Ping - Permits the management workstation to ping the switch.
- All - Permits all of the above.

You can specify more than one by separating the selections with a comma (for example, "Telnet,Ping").

The new ACE is added to the ACL.

8. After making your changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Modifying an ACE

To modify an ACE, you need to know its identification number. To view the identification numbers of the ACEs, refer to “Displaying the ACEs” on page 771.

To modify an ACE, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **7** to select Management ACL.

   The Management ACL Configuration menu is shown in Figure 260 on page 764.

3. From the Management ACL Configuration menu, type **3** to select Modify Management ACL Entry.

   The following prompt is displayed:

   ```
   Enter the entry ID :  [1 to 256] -> 1
   ```

4. Enter the identification number of the ACE to modify. You can modify one ACE at a time.

   The specifications of the selected ACE are displayed in the Modify Management ACL Entry window. An example of the window is shown in Figure 261.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Marketing
User: Manager                          11:20:02 02-Mar-2006
                Modify Management ACL Entry
Configuring Management ACL
1 - ID ........................... 11
2 - IP Address ................... 149.44.44.44
3 - Network Mask ................. 255.255.255.255
4 - Application(s) ............... Telnet,Ping

M - Modify Management ACL Entry

R - Return to Previous Menu

Enter your selection?
```

Figure 261. Modify Management ACL Entry

5. Make the desired changes to the entry by selecting the corresponding option and entering a new value. You cannot change an entry's ID number. For information on an entry's IP address, network mask, and applications, refer to steps 5, 6, and 7 in the procedure "Creating an ACE" on page 766.

6. After entering your changes, type **M** to select Modify Management ACL Entry.

   Your changes are immediately implemented on the switch.

7. After making your changes, type **R** until you return to the Main Menu. Then type **S** to select Save Configuration Changes.

Deleting an ACE

To delete an ACE, you need to know its identification number. To view the identification numbers of the ACEs, refer to “Displaying the ACEs” on page 771.

Note
If you are managing the switch from a Telnet management session and the management ACL is active, your management session will end and you will not be able to reestablish it if you delete the ACE that specifies your management workstation.

To delete an ACE, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **7** to select Management ACL.

   The Management ACL Configuration menu is shown in Figure 260 on page 764.

3. From the Management ACL Configuration menu, type **4** to select Delete Management ACL Entry.

   The following prompt is displayed:

   ```
   Enter the entry ID :  [1 to 256] -> 1
   ```

4. Enter the identification number of the ACE to be deleted.

   The ACE is immediately deleted from the management ACL.

5. If desired, repeat this procedure starting with Step 3 to delete more ACEs from the Management ACL.

Displaying the ACEs

To display the ACEs in the management ACL, perform the following procedure:

1. From the Main Menu, type **5** to select System Administration.

   The System Administration menu is shown in Figure 4 on page 52.

2. From the System Administration menu, type **7** to select Management ACL.

   The Management ACL Configuration menu is shown in Figure 260 on page 764.

3. From the Management ACL Configuration menu, type **5** to select Display All Management ACL Entries.

   The Display All Management ACL Entries menu is shown in Figure 262.

```
     Allied Telesyn Ethernet Switch AT-8524M - AT-S62
                        Marketing
User: Manager                           11:20:02 02-Mar-2006
             Display All Management ACL Entries
ID    IP Address         Mask                 Application
-----------------------------------------------------------
1     133.22.145.18      255.255.255.255      All
2     133.22.146.0       255.255.255.0        Web

U - Update display
R - Return to Previous Menu

Enter your selection?
```

Figure 262. Display All Management ACL Entries Menu

The menu provides the following information about the ACEs:

ID
The entry's identification number.

IP Address
The IP address of a management station or a subnet.

Mask
The parts of the IP address the switch is filtering on.

Application
The application that the management station is permitted to use to manage the switch. The options are Telnet, Web, Ping and All.

Appendix A
AT-S62 Default Settings

This appendix lists the AT-S62 factory default settings. It contains the following sections:

Basic Switch Default Settings

This section lists the default settings for basic switch parameters. The following topics are covered:

Boot Configuration File Default Setting

The following table lists the File Menu default setting.

File Menu Setting	Default
Default Configuration File	boot.cfg

Management Access Default Settings

The following table lists the management access default settings.

Remote Management Access Setting	Default
Telnet	Enabled
SNMP	Disabled
TFTP	Enabled
Web Server	Enabled

Management Interface Default Settings

The following table lists the management interface default settings.

Management Interface Setting	Default
Manager Login Name	manager
Manager Password	friend
Operator Login Name	operator
Operator Password	operator
Console Disconnect Timer Interval	10 minutes

Note
Login names and passwords are case-sensitive.

RS-232 Port Default Settings

The following table lists the RS-232 Terminal Port default settings.

RS-232 Port Setting	Default
Data Bits	8
Stop Bits	1
Parity	None
Flow Control	None
Baud Rate	9600 bps

SNTP Default Settings

The following table lists the SNTP default settings.

SNTP Setting	Default
System Time	00:00:00 on January 1, 1980
SNTP Status	Disabled
SNTP Server	0.0.0.0
UTC Offset	+0
Daylight Savings Time (DST)	Enabled
Poll Interval	600 seconds

Switch Administration Default Settings

The following table describes the switch administration default settings.

Administration Setting	Default
IP Address	0.0.0.0
Subnet Mask	0.0.0.0
Gateway Address	0.0.0.0
System Name	None
Administrator	None
Comments	None
BOOTP	Disabled

Administration Setting	Default
DHCP	Disabled
MAC Address Aging Time	300 seconds

System Software Default Settings

The following table lists the system software default settings.

System Software Setting	Default
Console Startup Mode	Command line

AT-8524POE Fan Control Default Setting

The following table lists the default setting for the fan control feature on the AT-8524POE switch.

System Software Setting	Default
Fan Control	Off

Denial of Service Defense Default Settings

The following table lists the default settings for the Denial of Service defense feature.

Denial of Service Defense Setting	Default
IP Address	0.0.0.0
Subnet Mask	0.0.0.0
Uplink Port	26
SYN Flood Defense	Disabled
Smurf Defense	Disabled
Land Defense	Disabled
Teardrop Defense	Disabled
Ping of Death Defense	Disabled
IP Options Defense	Disabled

Enhanced Stacking Default Setting

The following table lists the enhanced stacking default setting.

Enhanced Stacking Setting	Default
Switch State	Slave

Event Log Default Settings

The following table lists the event log default settings.

Event Log Setting	Default
Status	Enabled
Full Log Action	Wrap

GVRP Default Settings

This section provides the default settings for GVRP.

GVRP Setting	Default
Status	Disabled
GIP Status	Enabled
Join Timer	20 centiseconds
Leave Timer	60 centiseconds
Leave All Timer	1000 centiseconds
Port Mode	Normal

IGMP Snooping Default Settings

The following table lists the IGMP Snooping default settings.

IGMP Snooping Setting	Default
IGMP Snooping Status	Disabled
Multicast Host Topology	Single Host/ Port (Edge)
Host/Router Timeout Interval	260 seconds
Maximum Multicast Groups	64
Multicast Router Ports Mode	Auto Detect

MAC Address-based Security Default Settings

The following table lists the MAC address security default settings.

MAC Address Security Setting	Default
Security Mode	Automatic (no security)
Intrusion Action	Discard
Participating	No
MAC Limit	No Limit

Management Access Control List Default Setting

The following table lists the default setting for the Management Access Control List.

Management ACL Setting	Default
Status	Disabled

PKI Default Settings

The following table lists the PKI default settings, including the generate enrollment request settings.

PKI Setting	Default
Switch Distinguished Name	None
Maximum Number of Certificates	256
Request Name	None
Key Pair ID	0
Format	PEM
Type	PKCS10

Port Configuration Default Settings

The following table lists the port configuration default settings.

Port Configuration Setting	Default
Status	Enabled
Speed	Auto-Negotiation
Duplex Mode	Auto-Negotiation
MDI/MDI-X	Auto-MDI/MDIX
Broadcast, Multicast, and Unknown Unicast Packet Rate Limiting	Disabled
Unknown Multicast Packet Filtering	Disabled
Broadcast Packet Filter	Disabled
Head of Line Blocking Threshold	61,440 cells
Backpressure Status	Disabled
Backpressure Threshold	57,344 cells
Flow Control	Disabled
Flow Control Threshold	57,344 cells
Port Priority	0
Override Priority	No

802.1x Port-Based Network Access Control Default Settings

The following table describes the 802.1x Port Access Control default settings.

802.1x Port Access Control Setting	Default
Port Access Control	Disabled
Authentication Method	RADIUS EAP
Port Role	None

The following table lists the default settings for RADIUS accounting.

RADIUS Accounting Setting	Default
Status	Disabled
Port	1813
Type	Network
Trigger Type	Start_Stop
Update Status	Disabled
Update Interval	60

The following table lists the default settings for an authenticator port.

Authenticator Port Setting	Default
Authentication Mode	802.1x
Supplicant Mode	Single
Port Control	Auto
Quiet Period	60 seconds
TX Period	30 seconds
Reauth Enabled	Enabled
Reauth Period	3600 seconds
Supplicant Timeout	30 seconds
Server Timeout	30 seconds
Max Requests	2

Authenticator Port Setting	Default
VLAN Assignment	Enabled
Secure VLAN	On
Control Direction	Both
Piggyback Mode	Disabled
Guest VLAN	None

The following table lists the default settings for a supplicant port.

Supplicant Port Setting	Default
Auth Period	30 seconds
Held Period	60 seconds
Max Start	3
Start Period	30 seconds
User Name	(none)
User Password	(none)

Power Over Ethernet

The following table describes the Power over Ethernet (PoE) default settings. This feature only applies to the AT-8524POE switch.

PoE Setting	Default
PoE Status	Enabled
Port PoE Status	Enabled
Maximum Port Power	15.4 W
Port Priority	Low
PoE Threshold	95%

Class of Service

The following table lists the default mappings of IEEE 802.1p priority levels to egress port priority queues.

IEEE 802.1p Priority Level	Port Priority Queue
0	Q1
1	Q0
2	Q0
3	Q1
4	Q2
5	Q2
6	Q3
7	Q3

Server-Based Authentication Default Settings

This section describes the server-based authentication, RADIUS, and TACACS+ client default settings.

Server-Based Authentication Default Settings

The following table describes the server-based authentication default settings.

Server-based Authentication Setting	Default
Server-based Authentication	Disabled
Active Authentication Method	TACACS+

RADIUS Default Settings

The following table lists the RADIUS configuration default settings.

RADIUS Configuration Setting	Default
Global Encryption Key	ATI
Global Server Timeout Period	30 seconds
RADIUS Server 1 Configuration	0.0.0.0
RADIUS Server 2 Configuration	0.0.0.0
RADIUS Server 3 Configuration	0.0.0.0
Auth Port	1812
Encryption Key	Not Defined

TACACS+ Client Default Settings

The following table lists the TACACS+ client configuration default settings.

TACACS+ Client Configuration Setting	Default
TAC Server 1	0.0.0.0
TAC Server 2	0.0.0.0
TAC Server 3	0.0.0.0
TAC Global Secret	None
TAC Timeout	30 seconds

SNMP Default Settings

The following table describes the SNMPv1 and SNMPv2c default settings.

SNMP Communities Setting	Default
SNMP Status	Disabled
Authentication Failure Trap Status	Disabled
Community Name	public (Read only)
Community Name	private (Read|Write)
Status (public)	Enabled
Status (private)	Enabled
Open Status (public	Yes
Open Status (private)	Yes

STP, RSTP, and MSTP Default Settings

This section provides the spanning tree, STP RSTP, and MSTP, default settings.

Spanning Tree Switch Settings

The following table describes the Spanning Tree Protocol default settings for the switch.

STP Switch Setting	Default
Spanning Tree Status	Disabled
Active Protocol Version	RSTP

STP Default Settings

The following table describes the STP default settings.

STP Setting	Default
Bridge Priority	32768
Bridge Hello Time	2
Bridge Forwarding	15
Bridge Max Age	20
Port Cost	Automatic -Update
Port Priority	128

RSTP Default Settings

The following table describes the RSTP default settings.

RSTP Setting	Default
Force Version	RSTP
Bridge Priority	32768
Bridge Hello Time	2
Bridge Forwarding	15
Bridge Max Age	20
Edge Port	Yes
Point-to-Point	Auto Detect
Port Cost	Automatic Update
Port Priority	128

MSTP Default Settings

The following table lists the MSTP default settings.

MSTP Setting	Default
Status	Disabled
Force Version	MSTP
Bridge Hello Time	2
Bridge Forwarding Delay	15
Bridge Max Age	20
Maximum Hops	20
Configuration Name	null
Revision Level	0
CIST Priority	Increment 8 (32768)
Port Priority	Increment 8 (128)
Port Internal Path Cost	Auto Update
Port External Path Cost	200,000
Point-to-Point	Auto Detect
Edge Port	Yes

SSH Default Settings

The following table lists the SSH default settings.

SSH Setting	Default
Status	Disabled
Host Key ID	Not Defined
Server Key ID	Not Defined
Server Key Expiry Time	0 hours
Login Timeout	180 seconds

SSL Default Settings

The following table lists the SSL default settings.

SSL Setting	Default
Maximum Number of Sessions	50
Session Cache Timeout	300 seconds

VLAN Default Settings

This section provides VLAN default settings.

VLAN Setting	Default
Default VLAN Name	Default_VLAN (all ports)
Management VLAN ID	1 (Default_VLAN)
VLAN Mode	User Configured
Uplink Port	None
Ingress Filtering	Disabled

Web Server Default Settings

The following table lists the web server default settings.

Web Server Configuration Setting	Default
Status	Enabled
Mode	HTTP
Port Number	80
SSL Key ID	None

Appendix B
SNMPv3 Configuration Examples

This appendix provides two examples of SNMPv3 configuration using the SNMPv3 Table menus and a worksheet to use as an aid when configuring the SNMPv3 protocol. It includes the following sections:

SNMPv3 Configuration Examples

- This appendix provides SNMPv3 configuration examples for the following types of users:
- Manager
- Operator

In addition an SNMPv3 Configuration Table is provided to record your SNMPv3 configuration.

For more information about the SNMPv3 protocol, see Chapter 21, “SNMPv3” on page 375.

SNMPv3 Manager Configuration

This section provides a sample configuration for a Manager with a User Name of systemadmin24. Each table is listed with its parameters.

Configure SNMPv3 User Table Menu

```
User Name: systemadmin24
Authentication Protocol: MD5
Privacy Protocol: DES
Storage Type: NonVolatile
```

Configure SNMPv3 View Table Menu

```
View Name: internet
View Subtree OID: internet (or 1.3.6.1)
Subtree Mask:
View Type: Included
Storage Type: NonVolatile
```

Configure SNMPv3 Access Table

```
Group Name: Managers
Security Model: SNMPv3
Security Level: P-Authentication and Privacy
Read View Name: internet
Write View Name: internet
Notify View Name: internet
Storage Type: NonVolatile
```

Configure SNMPv3 SecurityToGroup Table

```
User Name:systemadmin24
Security Model:v3
Group Name: Managers
Storage Type: NonVolatile
```

Configure SNMPv3 Notify Table

```
Notify Name: sysadminTrap
Notify Tag: sysadminTag
Notify Type: Trap
Storage Type: NonVolatile
```

Configure SNMPv3 Target Address Table

```
Target Address Name: host451
Target IP Address: 198.35.11.1
UDP Port#: 162
Timeout: 1500
Retries: 3
Tag List: sysadminTag
Target Parms Name: SNMPmanagerPC
Storage Type: NonVolatile
```

Configure SNMPv3 Target Parameters Table

```
Target Parameters Name:SNMPmanagerPC
User Name:systemadmin24
Security Model: v3
Security Level: P-Authentication and Privacy
Storage Type: NonVolatile
```

SNMPv3 Operator Configuration

This section provides a sample configuration for an Operator with a User Name of nikoeng73. Because this user will only send messages to a group and not an SNMP host, you do not need to configure message notification for this user.

Configure SNMPv3 User Table Menu

```
User Name: nikoeng73
Authentication Protocol: MD5
Privacy Protocol: None
Storage Type: NonVolatile
```

Configure SNMPv3 View Table Menu

```
View Name: internet
View Subtree OID: 1.3.6.1 (or internet)
Subtree Mask:
View Type: Included
Storage Type: NonVolatile
```

Configure SNMPv3 Access Table

```
Group Name: Operators
Security Model: SNMPv3
Security Level: Authentication
Read View Name: internet
Write View Name:
Notify View Name:
```

SNMPv3 Worksheet

This section supplies a table that you can use a worksheet when configuring SNMPv3. Each SNMPv3 Table is listed with its associated parameters.

SNMPv3 Parameters	
SNMPv3 User Table	
User Name	
Authentication Protocol	
Authentication Password	
Privacy Protocol	
Privacy Password	
Storage Type	
SNMPv3 View Table Menu	
View Name	
View Subtree OID	
Subtree Mask	
View Type	
Storage Type	
SNMPv3 Access Table Menu	
Group Name	

SNMPv3 Parameters (Continued)	
Security Model	
Security Level	
Read View Name	
Write View Name	
Notify View Name	
Storage Type	
SNMPv3 SecurityToGroup Table	
User Name	
Security Model	
Group Name	
Storage Type	
SNMPv3 Notify Table	
Notify Name	
Notify Tag	
Notify Type	
Storage Type	
SNMPv3 Target Address Table	
Target Address Name	
Target IP Address	
UDP Port	
Timeout	
Retries	
Tag List	
Target Parms Name	
Storage Type	
SNMPv3 Target Parameters Table	
Target Parameters Name	
User (Security) Name	

SNMPv3 Parameters (Continued)	
Security Model	
Security Level	
Storage Type	

Appendix C

Standards and Features

10/100Base-TX Twisted Pair Ports

IEEE 802.1d	Bridging
IEEE 802.3	10Base-T
IEEE 802.3u	100Base-TX
IEEE 802.3u	Auto-Negotiation
IEEE 802.3x	10/100 Mbps Flow Control / Backpressure
—	Auto-MDI/MDIX
—	Head of Line Blocking
—	Four Egress Queues Per Port

Fiber Optic Ports (AT-8516F/SC Switch)

IEEE 802.1d	Bridging
IEEE 802.3u	100Base-SX
—	Head of Line Blocking
—	Four Egress Queues Per Port

Traffic Control

RFC 2386	Quality of Service featuring: — Layer 2, 3, and 4 criteria — Flow Groups, Traffic Classes, and Policies — DSCP Replacement — 802.1q Priority Replacement

	— Type of Service Replacement
	— Type of Service to 802.1q Priority Replacement
	— 802.1q Priority to Type of Service Replacement
	— Maximum Bandwidth Control
	— Burst Size Control
	— Support for Ingress and Egress Ports
IEEE 802.1p	Class of Service with Strict and Weighted Round Robin Scheduling
—	Port Access Control Lists
—	Ingress Packet Rate Limiting

Spanning Tree Protocols

IEEE 802.1D	Spanning Tree Protocol
IEEE 802.1w	Rapid Spanning Tree Protocol
IEEE 802.1s	Multiple Spanning Tree Protocol

Port Trunks

IEEE 802.3ad	Link Aggregation Control Protocol (LACP)
—	Static Port Trunks

Virtual LANs

IEEE 802.1Q	Tagged VLANs
—	Port-based VLANs
—	Compliant and Non-compliant 802.1Q Multiple VLAN Modes
—	Protected Ports VLANs
—	Selectable Management VLAN

IEEE 802.3ac	VLAN Tag Frame Extension
IEEE 802.1P	GARP VLAN Registration Protocol

IP Multicast

RFC 1112	IGMP Snooping (Ver. 1.0)
RFC 2236	IGMP Snooping (Ver. 2.0)
RFC 3376	IGMP Snooping (Ver. 3.0)

Port Security

IEEE 802.1x	Port-based Network Access Control with Multiple Supplicants Per Port
RFC 2865	RADIUS Client
RFC 2866	RADIUS Accounting
—	MAC Address-based security

Management Access and Security

RFC 1157	SNMPv1
RFC 1901	SNMPv2
RFC 3411	SNMPv3
RFC 1492	TACACS+ Client
RFC 2865	RADIUS Client
RFC 2068	HTTP
RFC 2616	HTTPS
RFC 1866	HTML
RFC 854	Telnet Server
—	Secure Sockets Layer (SSL)
RFC 4325 (X.509)	Public Key Infrastructure (PKI)

—	Encryption Keys
—	Secure Shell (SSH) (Vers. 1.3, 1.5, 2.0)
—	Management Access Control List

Management MIBs

RFC 1213	MIB-II
RFC 1215	TRAP MIB
RFC 1493	Bridge MIB
RFC 2863	Interface Group MIB
RFC 1643	Ethernet-like MIB
RFC 2674	IEEE 802.1Q MIB
RFC 1757	RMON 4 groups
—	Allied Telesyn Private MIBs

System Monitoring

RFC 3195	Syslog client
—	Temporary Event Log (4,000 events maximum)
—	Port Statistics
RFC 1757	RMON groups 1, 2, 3, and 9

Additional Features

RFC 2131	DHCP client
RFC 951, 1542	BOOTP client
RFC 1350	TFTP client
RFC 2030	SNTP
—	Port Mirroring

—	MAC address table with a storage capacity of 8K entries
—	2 megabyte file system

Denial of Service Defenses

Smurf

SYN Flood

Teardrop

Land

IP Option

Ping of Death

Management Access Methods

Enhanced Stacking™

Out-of-band management (serial port)

In-band management (over the network) using Telnet, SSH, web browser, or SNMP

Management Interfaces

Menus

Command Line

Web Browser

SNMP v1, v2, & v3

Index

C

D

E

F

G

M

Q

R

V

W

X