Summit24e2 Installation and User's Guide

Extreme Networks, Inc.
3585 Monroe Street
Santa Clara, California 95051
1 (408) 579-2800
http://www.extremenetworks.com

Published: Octoberl 2001

Table of Contents

B Appendix B - Runtime Switching Software Defaults

1 Summit24e2 Overview

Features

The Summit24e2 Switch was designed for easy installation and high performance in an environment where traffic on the network and the number of users increase continuously.

The Summit24e2 is a 24 10/100 Mbps port Ethernet switch with two Gigabit Ethernet ports. The Summit24e2 supports auto-negotiation between 10 Mbps and 100 Mbps at full or half-duplex operation on all 24 ports simultaneously.

Switch features include the following:

Ports

- 24 high performance, auto-negotiating ports all operating at 10/100 Mbps for connecting to end stations, servers and hubs.
- All ports can auto-negotiate between 10 Mbps/100 Mbps, half-duplex or full duplex and flow control for half-duplex ports.
- RS-232 DTE Diagnostic port (console port) for setting up and managing the switch via a connection to a console terminal or PC using a terminal emulation program.
- Two 1000Base-T or two GBIC Ethernet ports.

Performance Features

- 8.8 Gbps switching fabric capacity
- Store and forward switching scheme.
- Full and half-duplex for both 10 Mbps and 100 Mbps connections. The Gigabit Ethernet ports operate at full-duplex only. Full-duplex allows the port to simultaneously transmit and receive data, and requires full-duplex end stations and switches. Connections to hubs must take place at half-duplex.
- Supports IEEE 802.3x flow control for full-duplex mode ports.
- Supports back-pressure flow control for half-duplex mode ports.
- Auto-polarity detection and correction of incorrect polarity on the transmit and receive twisted-pair at each port.
- IEEE 802.3z compliant for all Gigabit ports.
- IEEE 802.3x compliant Flow Control support for all Gigabit ports.
- IEEE 802.3ab compliant for 1000BASE-TX (Copper) Gigabit ports.
- Data forwarding rate 14,880 pps per port at 100% of wire-speed for 10Mbps speed.
- Data forwarding rate 148,800 pps per port at 100% of wire-speed for 100Mbps speed.
- Data filtering rate eliminates all error packets, runts, etc. at 14,880 pps per port at 100% of wire-speed for 10 Mbps speed.
- Data filtering rate eliminates all error packets, runts, etc. at 148,800 pps per port at 100% of wire-speed for 100 Mbps speed.
- 8K active MAC address entry table per device with automatic learning and aging (10 to 1,000,000 seconds).
- 16 MB packet buffer per device.
- Broadcast and Multicast storm filtering.
- Supports Port Mirroring.
- Supports Port Trunking – up to six trunk groups (each consisting of up to eight ports) can be set up.
- 802.1D Spanning Tree support.
- 802.1Q Tagged VLAN support – up to 63 User-defined VLANs per device (one VLAN is reserved for internal use).
- 802.1p Priority support with 4 priority queues.
- IGMP Snooping support.

Management Features

- RS-232 console port for out-of-band network management via a console terminal or PC.
- Spanning Tree Algorithm Protocol for creation of alternative backup paths and prevention of network loops.
- SNMP v.1 agent.
- Fully configurable either in-band or out-of-band control via SNMP based software.
- Flash memory for software upgrades. This can be done in-band via TFTP or out-of-band via the console.
- Built-in SNMP management:
  - Bridge MIB (RFC 1493).
  - MIB-II (RFC 1213).
  - Mini-RMON MIB (RFC 1757) - 4 groups.
  - CIDR MIB (RFC 2096), except IP Forwarding Table.
  - 802.1p MIB (RFC 2674).
  - RIP MIB v2 (RFC 1724).
- Supports Web-based management.
- TFTP support.
- BOOTP support.
- BOOTP relay agent.
- DCHP client support.
- DCHP relay agent.
- DNS relay agent.
- Password enabled.

Switching Technology

Another key development pushing the limits of Ethernet technology is in the field of switching technology. A switch bridges Ethernet packets at the MAC address level of the Ethernet protocol, transmitting among connected Ethernet or fast Ethernet LAN segments.

Switching is a cost-effective way of increasing the total network capacity available to users on a local area network. A switch increases capacity and decreases network loading by making it possible for a local area network to be divided into different *segments* which don't compete with each other for network transmission capacity, giving a decreased load on each.

The switch acts as a high-speed selective bridge between the individual segments. Traffic that needs to go from one segment to another (from one port to another) is automatically forwarded by the switch, without interfering with any other segments (ports). This allows the total network capacity to be multiplied, while still maintaining the same network cabling and adapter cards.

For Fast Ethernet or Gigabit Ethernet networks, a switch is an effective way of eliminating problems of chaining hubs beyond the "two-repeater limit." For example, a switch can be used to split parts of the network into different collision domains, making it possible to expand your Fast Ethernet network beyond the 205 meters network diameter limit for 100BASE TX networks. Switches supporting both traditional 10Mbps Ethernet and 100Mbps Fast Ethernet are also ideal for bridging between existing 10Mbps networks and new 100Mbps networks.

Switching LAN technology is a marked improvement over the previous generation of network bridges, which were characterized by higher latencies. Routers have also been used to segment local area networks, but the cost of a router and the setup and maintenance required can make routers impractical in some situations. Today's switches are an ideal solution for many kinds of local area congestion problems.

Fast Ethernet Technology

100 Mbps Fast Ethernet (or 100BASE-T) is a standard specified by the IEEE 802.3 LAN committee. It is an extension of the 10Mbps Ethernet standard with the ability to transmit and receive data at 100Mbps, while maintaining the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Ethernet protocol.

Gigabit Ethernet Technology

Gigabit Ethernet is an extension of IEEE 802.3 Ethernet utilizing the same packet structure, format, and support for CSMA/CD protocol, full duplex, flow control, and management objects, but with a tenfold increase in theoretical throughput over 100Mbps Fast Ethernet and a one hundred-fold increase over 10Mbps Ethernet. Since it is compatible with all 10Mbps and 100Mbps Ethernet environments, Gigabit Ethernet

provides a straightforward upgrade without wasting a company's existing investment in hardware, software, and trained personnel.

Gigabit Ethernet enables fast optical fiber connections and Unshielded Twisted Pair connections to support video conferencing, complex imaging, and similar data-intensive applications. Likewise, since data transfers occur 10 times faster than Fast Ethernet, servers outfitted with Gigabit Ethernet NIC's are able to perform 10 times the number of operations in the same amount of time.

2 Unpacking and Setup

This chapter provides unpacking and set-up information for the switch.

- Unpacking
- Installation
- Power On

Unpacking

Open the shipping carton of the switch and carefully unpack its contents. The carton should contain the following items:

- One Summit24e2 24-port Ethernet Switch
- Mounting kit: 2 mounting brackets and screws
- Four rubber feet with adhesive backing
- One AC power cord
- One Extreme Networks Documentation CD
- User Registration
- License Agreement

If any item is found missing or damaged, please contact your local Extreme Networks reseller for replacement.

Installation

Use the following guidelines when choosing a place to install the Switch:

- The surface must support at least 3 kg.
- The power outlet should be within 1.82 meters (6 feet) of the device.
- Visually inspect the power cord and see that it is secured to the AC power connector.
- Make sure that there is proper heat dissipation from and adequate ventilation around the switch. Do not place heavy objects on the switch.

Desktop or Shelf Installation

When installing the Switch on a desktop or shelf, the rubber feet included with the device should first be attached. Attach these cushioning feet on the bottom at each corner of the device, as shown in Figure 2-1. Allow adequate space for ventilation between the device and the objects around it.

Summit24e2

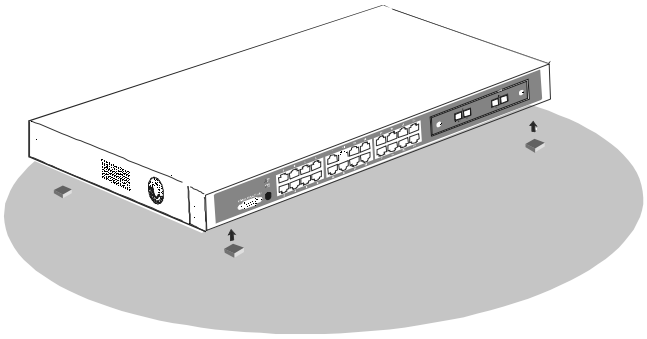

Figure 2-1: Installing Rubber Feet for Desktop Installation

Rack Installation

The Summit24e2 can be mounted in an EIA standard-sized, 19-inch rack, which can be placed in a wiring closet with other equipment. To install:

Attach the mounting brackets on the switch's side panels (one on each side), and secure them with the screws provided, as shown in Figure 2-2.

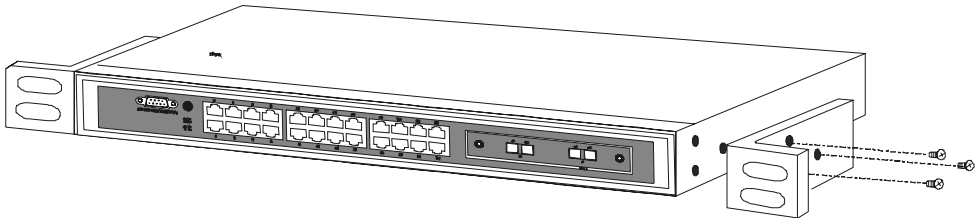

Figure 2-2: Attaching the Mounting Brackets to the Switch

Use the screws provided with the equipment rack to mount the switch on the rack as shown in Figure 2-3.

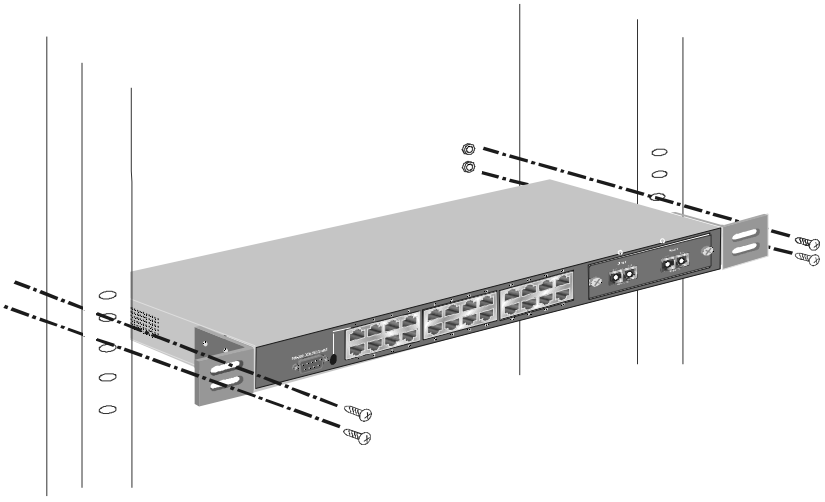

Figure 2-3: Installing the Switch on an Equipment Rack

Power On

The Summit24e2 switch can be used with AC power supply 100 - 240 VAC, 50 - 60 Hz. The switch's power supply will adjust to the local power source automatically and may be turned on without having any or all LAN segment cables connected.

After the power switch is turned on, the LED indicators should respond as follows:

- Power and Status LED indicators will momentarily blink. This blinking of the LED indicators represents a reset of the system.
- The Status LED indicator will blink while the switch loads onboard software and performs a self-test. After approximately 20 seconds, the LED will light again to indicate the switch is in a ready state.
- The console LED indicator will remain ON (glowing green) if there is a connection at the RS-232 port; otherwise this LED indicator is OFF.
- The 100M LED indicator may remain ON (glowing green) or OFF depending on the transmission speed.

Power Failure

As a precaution, in the event of a power failure, unplug the switch. When the power supply is restored, plug the switch back in.

3 Identifying External Components

This chapter describes the front panel, rear panel, and LED indicators of the Summit24e2.

- Front Panel
- Rear Panel
- Side Panels
- LED Indicators

Front Panel

The front panel of the Switch consists of:

- LED indicators
- An RS-232 communication port
- Two Gigabit Ethernet ports, either GBIC or 1000Base-T
- Twenty-four 10/100 Mbps Ethernet/Fast Ethernet ports

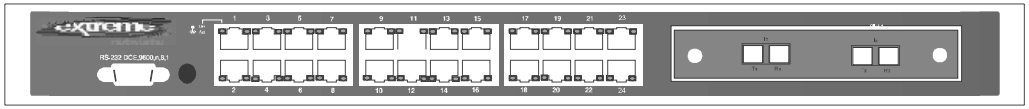

Figure 3-1: Front Panel View of the Switch

- Comprehensive LED indicators display the status of the switch and the network. A description of these LED indicators follows (see “LED Indicators” on page -5).
- An RS-232 DTE console port for setting up and managing the switch via a connection to a console terminal or PC using a terminal emulation program.
- 24 high-performance Ethernet ports all of which operate at 10/100 Mbps for connections to end stations, servers and hubs. All ports can auto-negotiate between 10 Mbps or 100 Mbps, full- or half-duplex, and flow control.

Rear Panel

The following displays the rear panel of the switch.

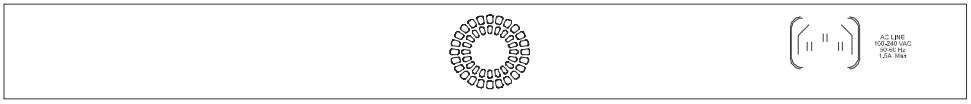

Figure 3-2: Rear Panel View of the Switch

- The AC power connector is a standard three-pronged connector that supports the power cord. Plug in the female connector of the provided power cord into this socket, and the male side of the cord into a power outlet. Supported input voltages range from 100 ~ 240 VAC at 50 ~ 60 Hz.

Side Panels

The right side panel of the Switch contains two system fans (see the top part of the diagram below). The left side panel contains heat vents.

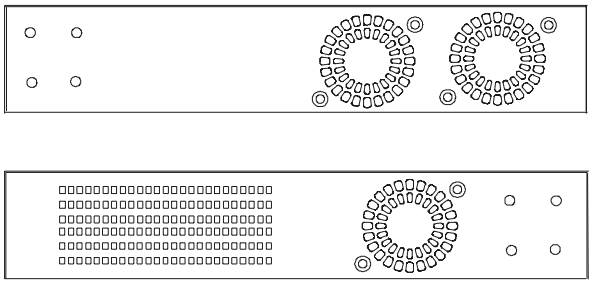

Figure 3-3: Side Panel Views of the Switch

The system fans are used to dissipate heat. The sides of the system also provide heat vents to serve the same purpose. Do not block these openings, and leave at least 6 inches of space at the rear and sides of the switch for proper ventilation. Remember that without proper heat dissipation and air circulation, system components might overheat, which could lead to system failure.

Gigabit Ethernet Ports

The Summit24e2 24-port Ethernet Switch includes two Gigabit Ethernet ports. The Gigabit Ethernet ports support either fiber optic cable (via a GBIC interface) or 1000BASE-TX (using Cat. 5 copper cable).

GBIC Slots

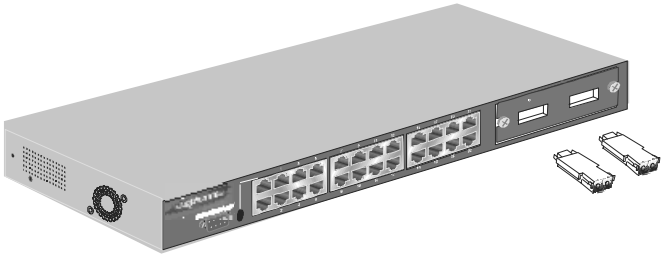

Figure 3-4: GBIC Two-Port Module

- Connects to GBIC devices at full duplex only.
- Allows multi-mode fiber optic connections of up to 550 m (SX and LX) and single-mode fiber optic connections of up to 5 km (LX only). GBIC modules are available in -SX and -LX fiber optic media.

Table 3-1: Fiber Cable Lengths for GBIC

	62.5µm	62.5µm	50µm	50µm
Modal Bandwidth	160MHz*km	200	400	500
Operating Distance	220 meters	275	500	550
Channel Insertion Loss	2.33dB	2.53	3.25	3.43

1000BASE-TX Ports

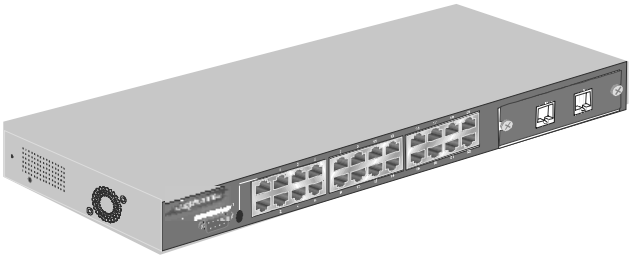

Figure 3-5: 1000BASE-TX ports

- Allows connection to 1000BASE-TX devices using Cat. 5e copper cable.
- 100/1000 Mbps Auto-Negotiation.
- 1000 Mbps connection is full-duplex only.
- Automatic MDI/MDIX uplink connection allowing the use of either a straight-through or a cross-over cable.

LED Indicators

The Switch's LED indicators include:

- Power
- Status
- Link/Act

The following shows the LED indicators and provides an explanation of each indicator.

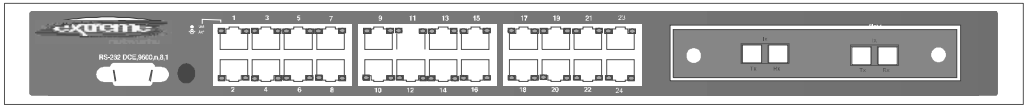

Figure 3-6: The Switch LED Indicators

Power — this indicator on the front panel should be lit during the Power-On Self Test (POST). It will light green approximately 2 seconds after the switch is powered on to indicate the ready state of the device.

Status — this indicator is lit green when the switch is able to be managed via out-of-band/local console management through the RS-232 console port using a straight-through serial cable.

Link/Act — these indicators are located to the left and right of each port. They are lit when there is a secure connection (or link) to a device at any of the ports. The LEDs blink whenever there is reception or transmission (i.e. Activity--Act) of data occurring at a port.

4 Connecting the Switch

This chapter describes how to connect the Summit24e2 to your Fast Ethernet network.

- Switch to End Node
- Switch to Hub or Switch

Switch to End Node

End nodes include PCs outfitted with a 10, 100, 10/100, 1000 or 100/1000 Mbps RJ-45 Ethernet/Fast Ethernet Network Interface Card (NIC) and most routers. Summit24e2 switches supplied with GBIC Gigabit Ethernet ports can be connected to other fiber-optic switch ports using the SC-type connector.

An end node can be connected to the Switch via a two-pair Category 3, 4, 5 UTP/STP straight cable (be sure to use Category 5e UTP or STP cabling for 100 and 1000 Mbps Fast Ethernet or Gigabit Ethernet connections). End nodes can be connected to any of the twenty-four 10/100 Mbps ports (1 - 24) of the Summit24e2, or to either of the two Gigabit ports on the front panel.

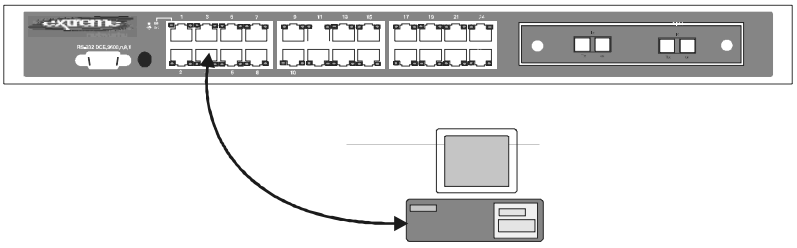

Figure 4-1: Switch Connected to an End Node

The LED indicators for the port that the end node is connected to are lit according to the capabilities of the NIC. If LED indicators are not illuminated after making a proper connection, check the PC's LAN card, the cable, switch conditions, and connections.

The following LED indicator states are possible for an end node to switch connection:

- The 100M LED indicator glows green for a 100 Mbps and stays OFF for 10 Mbps.
- The Link/Act LED indicator glows green upon hooking up a PC that is powered on and flashes to indicate that a working link has been established (there is activity — packets being received or transmitted — across the link).

Switch to Hub or Switch

These connections can be accomplished in a number of ways. The most important consideration is that when using a normal, straight-through cable, the connection should be made between a normal crossed port (Port 2, 3, etc.) and an Uplink (MDI-II) port. If you are using a crossover cable, the connection must be made from Uplink to Uplink (port 1 on the Summit24e2), or from a crossed port to another crossed port.

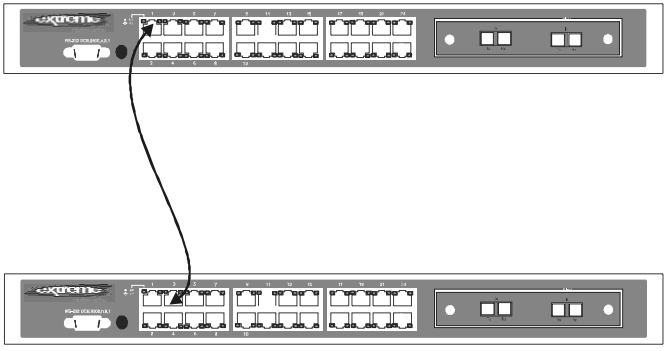

Figure 4-2: Switch Connected to Another Switch

- A 10BASE-T hub or switch can be connected to the Switch via a two-pair Category 3, 4 or 5 UTP/STP straight cable.
- A 100BASE-TX hub or switch can be connected to the Switch via a two-pair Category 5e UTP/STP straight cable.

If the other switch or hub contains an unused Uplink port, we suggest connecting the other device's Uplink (MDI-II) port to any of the switch's (MDI-X) ports (2 - 23, or one of the Gigabit module ports) using a normal straight-through cable — for the 1000BASE-TX module or a pair of fiber optic cables for the GBIC module, as shown in Figure 4-3, below.

- A 1000BASE-TX switch can be connected to the switch via a two-pair Category 5e UTP/STP straight cable.
- A GBIC switch can be connected to the switch via a pair of fiber optic cables, with the Tx port on one switch connected to the Rx port on the second switch, and vice-versa.
- GBIC connections can be made in full-duplex, only.
- GBIC allows multi-mode fiber optic connections of up to 550 m.

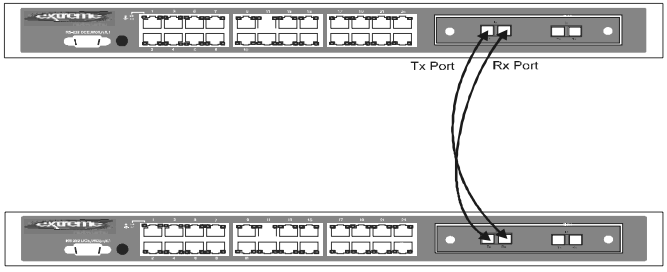

Figure 4-3: GBIC Port Connected to Another Switch's GBIC Port

5 Switch Management Concepts

This chapter discusses many of the features used to manage the switch and explains many concepts and important points regarding these features. Configuring the switch to implement these concepts is discussed in detail in the next chapters.

- Local Console Management
- IP Addresses and SNMP Community Names
- Traps
- MIBs
- SNMP
- Packet Forwarding
- Packet Filtering
- Spanning Tree Algorithm
- Port Aggregation
- VLANs
- Broadcast Storms

Local Console Management

A local console is a terminal or a workstation running a terminal emulation program that is connected directly to the switch via the RS-232 console port on the front of the switch. A console connection is referred to as an 'Out-of-Band' connection, meaning

that console is connected to the switch using a different circuit than that used for normal network communications. So, the console can be used to set up and manage the switch even if the network is down.

Local console management uses the terminal connection to operate the console program built-in to the switch (see Chapter 6 - Using the Console Interface). A network administrator can manage, control and monitor the switch from the console program.

The Summit24e2 switch contains a CPU, memory for data storage, flash memory for configuration data, operational programs, and SNMP agent firmware. These components allow the switch to be actively managed and monitored from either the console port or the network itself (out-of-band, or in-band).

Diagnostic (Console) Port (RS-232 DTE)

Out-of-band management requires connecting a terminal, such as a VT-100 or a PC running a terminal emulation program (such as HyperTerminal, which is provided with many versions of Microsoft Windows) to the RS-232 DTE console port of the switch. Switch management using the RS-232 DTE console port is called *Local Console Management* to differentiate it from management done via management platforms, such as HP OpenView, etc.

The console port is set for the following configuration:

- Baud rate = 9,600
- Data width = 8 bits
- Parity = none
- Stop bits = 1
- Flow Control = None

Make sure the terminal or PC you are using to make this connection is configured to match these settings.

If you are having problems making this connection on a PC, make sure the emulation is set to VT-100 or ANSI. If you still don't see anything, try typing **Ctrl + R** to refresh the screen.

The pin-out assignments for both the DB-9 (9 pin) and DB-25 (25 pin) RS-232 DTE cable connectors are shown in Figure 5-1, for convenience. The Summit24e2 switch's console port uses a Male DB-9 connector, but many data terminals (or PCs) may have a different

serial port connector. The console port uses a straight-through serial cable (not a null-modem cable).

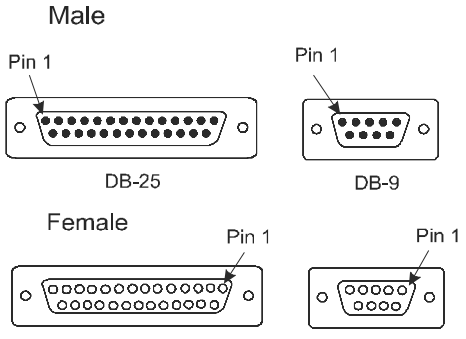

Console Port DB-9 Pin Number	Signal Line Description		Terminal Port DB-9 Pin Number
(1)	DCD (Data Carrier Detect)	In	(1)
(2)	RXD (Receive Data)	In	(2)
(3)	TXD (Tranmit Data)	Out	(3)
(4)	DTR (Data Terminal Ready)	Out	(4)
(5)	GND (Ground)	-	(5)
(6)	DSR (Data Set Ready)	In	(6)
(7)	RTS (Request to Send)	Out	(7)
(8)	CTS (Clear to Send)	In	(8)
(9)	Not Connected	-	(9)

Console Port DB-9 Pin Number	Signal Line Description		Terminal Port DB-25 Pin Number
(1)	DCD (Data Carrier Detect)	In	(8)
(2)	RXD (Receive Data)	In	(3)
(3)	TXD (Tranmit Data)	Out	(2)
(4)	DTR (Data Terminal Ready)	Out	(20)
(5)	GND (Ground)	-	(7)
(6)	DSR (Data Set Ready)	In	(6)
(7)	RTS (Request to Send)	Out	(4)
(8)	CTS (Clear to Send)	In	(5)
(9)	Not Connected (Ring)	-	(22)

Note: For DB-25 serial connectors, the following pins are not used - 1, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 21, 23, 24, and 25

Figure 5-1: RS-232 DTE Connector Pin-out and Cable End Connection Diagram

IP Addresses and SNMP Community Names

Each switch must be assigned its own IP address, which is used for communication with an SNMP network manager or other TCP/IP application (for example BOOTP, TFTP). The switch's default IP address is 0.0.0.0. You can change the default switch IP address to meet the specification of your networking address scheme.

The switch is also assigned a unique MAC address by the factory. This MAC address cannot be changed, and can be found from the initial boot console screen, as shown in Figure 5-2 below.

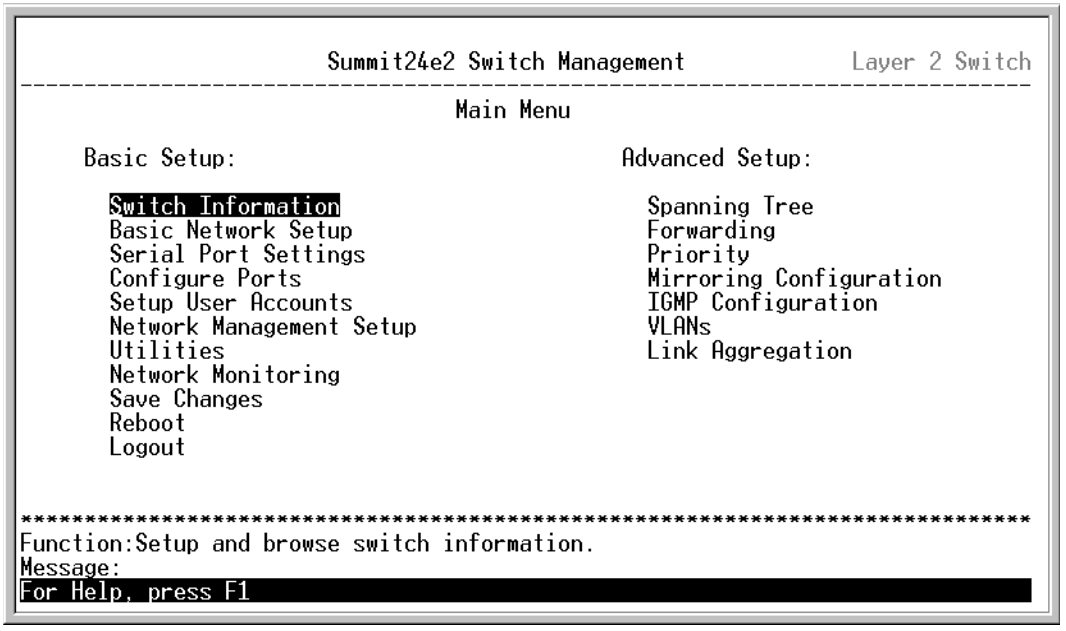

Figure 5-2: Boot Screen

The switch's MAC address can also be found from the console program under the Switch Information menu item, as shown in Figure 5-3 below.

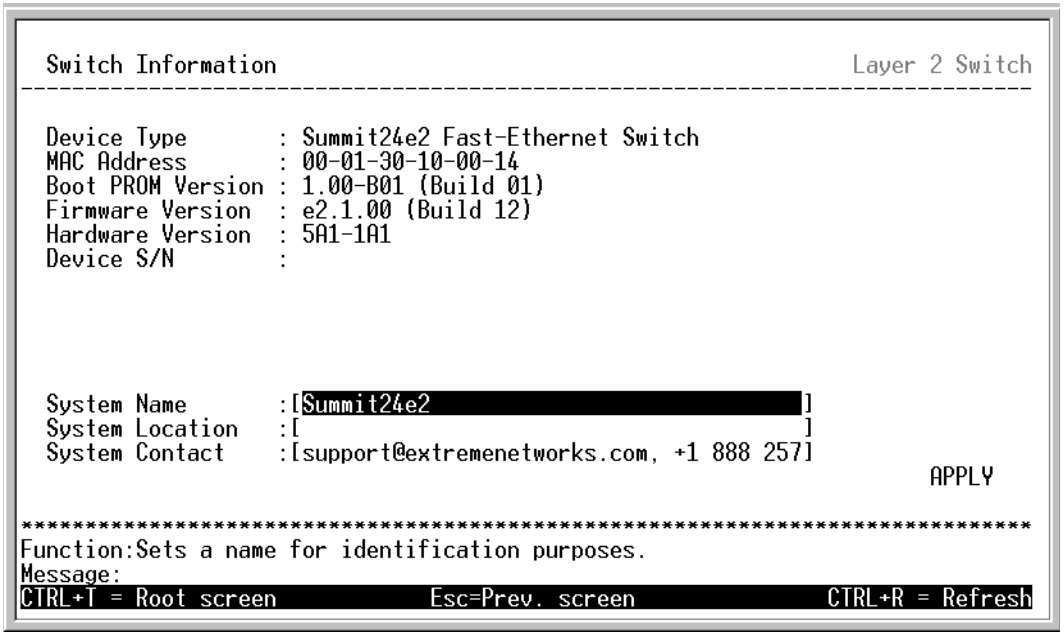

Figure 5-3: Switch Information Screen

In addition, you can also enter an IP address for a gateway router. This becomes necessary when the network management station is located on a different IP network from the switch, making it necessary for management packets to go through a router to reach the network manager, and vice-versa.

For security, you can set in the switch a list of IP Addresses of the network managers that you allow to manage the switch. You can also change the default SNMP Community Strings in the switch and set the access rights of these Community Strings. In addition, a VLAN may be designated as a Management VLAN.

Traps

Traps are messages that alert you of events that occur on the switch. The events can be as serious as a reboot (someone accidentally turned off the switch), or less serious (a port status change). The switch generates traps and sends them to the network manager (trap recipient).

Trap recipients are special users of the network who are given certain rights and access in overseeing the maintenance of the network. Trap recipients will receive traps sent from the switch; they must immediately take certain actions to avoid future failure or breakdown of the network. Trap recipients are configured using the Remote Management Setup menu, as shown in Figure 5-4 below.

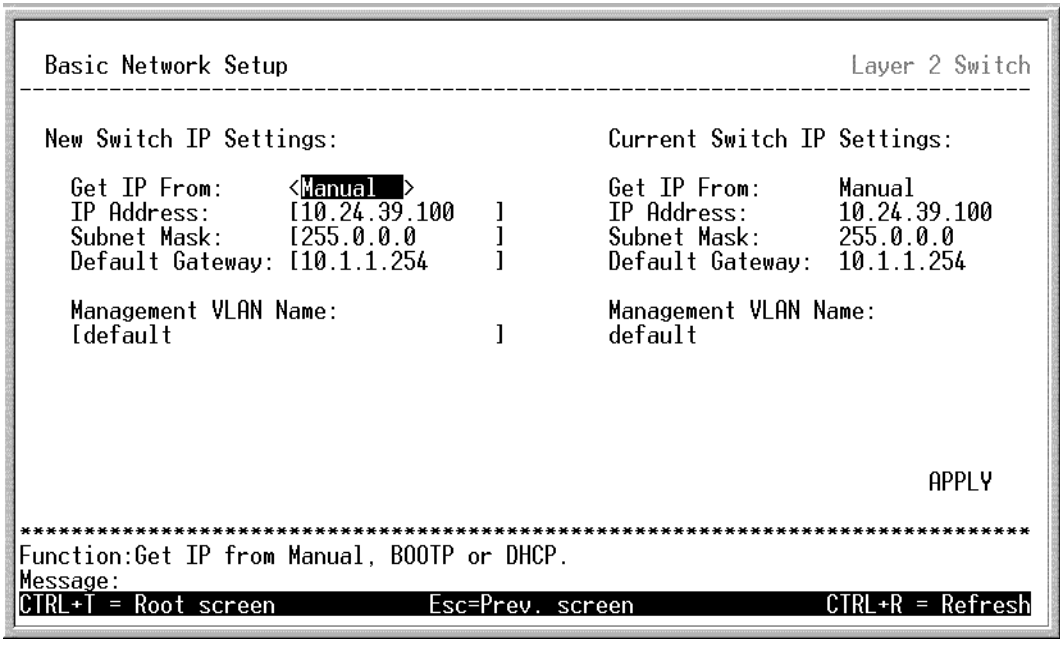

Figure 5-4: Remote Management Setup Menu

You can also specify which network managers may receive traps from the switch by entering a list of the IP addresses of authorized network managers. Up to four trap recipient IP addresses, and four corresponding SNMP community strings can be entered.

SNMP community strings function like passwords in that the community string entered for a given IP address must be used in the management station software, or a trap will be sent. The following are trap types a trap manager will receive:

Cold Start This trap signifies that the switch has been powered up and initialized such that software settings are reconfigured and hardware systems are rebooted.

Warm Start	This trap signifies that the switch has been rebooted, however the POST (Power On Self-Test) is skipped.
Authentication Failure	This trap signifies that someone has tried to logon to the switch using an invalid SNMP community string. The switch automatically stores the source IP address of the unauthorized user.
New Root	This trap indicates that the switch has become the new root of the Spanning Tree, the trap is sent by the switch soon after its election as the new root. This implies that upon expiration of the Topology Change Timer the new root trap is sent out immediately after the switch's election as the new root.
Topology Change	A Topology Change trap is sent by the switch when any of its configured ports transitions from the Learning state to the Forwarding state, or from the Forwarding state to the Blocking state. The trap is not sent if a new root trap is sent for the same transition.
Link Up	This trap is sent whenever the link of a port changes from link down to link up.
Link Down	This trap is sent whenever the link of a port changes from link up to link down.

MIBs

Management and counter information are stored in the switch in the Management Information Base (MIB). The switch uses the standard MIB-II Management Information Base module. Consequently, values for MIB objects can be retrieved from any SNMP-based network management software. In addition to the standard MIB-II, the switch also supports its own proprietary enterprise MIB as an extended Management Information Base. These MIBs may also be retrieved by specifying the MIB's Object-Identity (OID) at the network manager. MIB values can be either read-only or read-write.

Read-only MIB variables can be either constants that are programmed into the switch, or variables that change while the switch is in operation. Examples of read-only constants are the number of port and type of ports. Examples of read-only variables are

the statistics counters such as the number of errors that have occurred, or how many kilobytes of data have been received and forwarded through a port.

Read-write MIBs are variables usually related to user-customized configurations. Examples of these are the switch's IP address, spanning tree algorithm parameters, and port status.

If you use a third-party vendors' SNMP software to manage the switch, a diskette listing the switch's propriety enterprise MIBs can be obtained by request. If your software provides functions to browse or modify MIBs, you can also get the MIB values and change them (if the MIBs' attributes permit the write operation). This process however can be quite involved, since you must know the MIB OIDs and retrieve them one by one.

SNMP

The Simple Network Management Protocol (SNMP) is an OSI layer 7 (application layer) protocol for remotely monitoring and configuring network devices. SNMP enables network management stations to read and modify the settings of gateways, routers, switches, and other network devices. SNMP can be used to perform many of the same functions as a directly connected console, or can be used within an integrated network management software package such as HP OpenView.

SNMP performs the following functions:

- Sending and receiving SNMP packets through the IP protocol.
- Collecting information about the status and current configuration of network devices.
- Modifying the configuration of network devices.

The Summit24e2 has a software program called an 'agent' that processes SNMP requests, but the user program that makes the requests and collects the responses runs on a management station (a designated computer on the network). The SNMP agent and the user program both use the UDP/IP protocol to exchange packets.

Authentication

The authentication protocol ensures that the remote user SNMP application program discard packets from unauthorized users. Authentication is accomplished using

'community strings', which function like passwords. The remote user SNMP application must use the community string. SNMP community strings of up to 20 characters may be entered under the Remote Management Setup menu (see Figure 5-4) of the console program.

Packet Forwarding

The switch learns the network configuration and uses this information to forward packets. This reduces the traffic congestion on the network, because packets, instead of being transmitted to all segments, are transmitted to the destination only. Example: if Port 1 receives a packet destined for a station on Port 2, the switch transmits that packet through Port 2 only, and transmits nothing through the other ports.

MAC Address Aging Time

The Aging Time affects the learning process of the switch. Dynamic forwarding table entries, which are made up of the source and destination MAC addresses and their associated port numbers, are deleted from the table if they are not accessed within the aging time.

The aging time can be from 10 to 1,000,000 seconds with a default value of 300 seconds. A very long aging time can result in dynamic forwarding table entries that are out-of-date or no longer exist. This may cause incorrect packet forwarding decisions by the switch.

If the Aging Time is too short however, many entries may be aged out too soon. This will result in a high percentage of received packets whose source addresses cannot be found in the forwarding table, in which case the switch will broadcast the packet to all ports, negating the benefits of having a switch.

Static forwarding entries are not affected by the aging time.

Packet Filtering

The switch uses a filtering database to segment the network and control communication between segments. It can also filter packets off the network for intrusion control. Static filtering entries can be made by MAC address or IP address filtering.

Each port on the switch is a unique collision domain and the switch filters (discards) packets whose destination lies on the same port as where it originated. This keeps local packets from disrupting communications on other parts of the network.

For intrusion control, whenever a switch encounters a packet originating from or destined to a MAC address or an IP address entered into the filter table, the switch will discard the packet.

Some filtering is done automatically by the switch:

- Dynamic filtering - automatic learning and aging of MAC addresses and their location on the network. Filtering occurs to keep local traffic confined to its segment.
- Filtering done by the Spanning Tree Protocol, which can filter packets based on topology, making sure that signal loops don't occur.
- Filtering done for VLAN integrity. Packets from a member of a VLAN (VLAN 2, for example) destined for a device on another VLAN (VLAN 3) will be filtered.

Some filtering requires the manual entry of information into a filtering table:

- MAC address filtering - the manual entry of specific MAC addresses to be filtered from the network. Packets sent from one manually entered MAC address can be filtered from the network. The entry may be specified as either a source, a destination, or both.

Spanning Tree Protocol

The IEEE 802.1D Spanning Tree Protocol (STP) allows for the blocking of links between switches that form loops within the network. When multiple links between switches are detected, a primary link is established. Duplicated links are blocked from use and become standby links. The protocol allows for the duplicate links to be used in the event of a failure of the primary link. Once the Spanning Tree Protocol is configured and enabled, primary links are established and duplicated links are blocked automatically. The reactivation of the blocked links (at the time of a primary link failure) is also accomplished automatically - without operator intervention.

This automatic network reconfiguration provides maximum uptime to network users. However, the concepts of the Spanning Tree Algorithm and protocol are a complex subject and must be fully researched and understood. It is possible to cause serious degradation of the performance of the network if the Spanning Tree is incorrectly

configured. Please carefully read understand this section before making any changes from the default values.

The Summit24e2 allows two levels of spanning trees to be configured. The first level constructs a spanning tree among all links between network switches. This first level is referred to as the switch or global level. The second level is based on port groups. Groups of ports are configured as being members of a spanning tree and the algorithm and protocol are applied to the group of ports. This is referred to as the port or VLAN level.

Spanning Tree on the switch performs the following functions:

- Creates a single spanning tree from any combination of switching or bridging elements.
- Creates multiple spanning trees - from any combination of ports contained within a single switch, in user-specified groups (usually VLANs).
- Automatically reconfigures the spanning tree to compensate for the failure, addition, or removal of any element in the tree.
- Reconfigures the spanning tree without operator intervention.

STP Operation Levels

STP operates on two levels: the switch level and the port or VLAN level. The switch level forms a spanning tree consisting of links between one or more switches. The port level constructs a spanning tree consisting of groups of one or more ports. The STP operates in much the same way for both levels.

On the switch level, STP calculates the bridge identifier for each switch, then sets the root bridge and the designated bridges.

On the port level, STP sets the root port and designated ports.

Switch Level STP

User configurable switch STP parameters:

Bridge Identifier A combination of the user-set priority and the switch's MAC address. The Bridge Identifier consists of two parts: a 16-bit priority and a 48-bit Ethernet MAC address. The default value = 32768 + the MAC Address of the switch.

Priority	A relative priority for each switch — lower numbers give a higher priority and a greater chance of a given switch being elected as the root bridge. The default value = 32768.
Hello Time	The length of time between broadcasts of the hello message by the switch. The default value = 2 seconds.
Maximum Age Timer	Measures the age of a received BPDU for a port and ensures that the BPDU is discarded when its age exceeds the value of the maximum age timer. The default value = 20 seconds.
Forward Delay Timer	The amount of time spent by a port in the learning and listening states waiting for a BPDU that may return the port to the blocking state. The default value = 15 seconds.

Port Level STP

The VLAN or port STP parameters listed here may be configured by the user:

Port Priority	A relative priority for each port — lower numbers give a higher priority and a greater chance of a given port being elected as the root port. The default value = 32768.
Port Cost	A value used by STP to evaluate paths — STP calculates path costs and selects the path with the minimum cost as the active path. The default value = 19 for 100Mbps Fast Ethernet ports, and 4 for 1000Mbps Gigabit Ethernet ports.

It is highly recommended that STP port groups mirror the VLAN port groups. Any differences between the membership of STP port groups and the membership of VLAN port groups can cause severe problems.

Bridge Protocol Data Units

The Switch uses the following information for STP to stabilize network topology:

- The unique switch identifier
- The path cost to the root associated with each switch port
- The port identifier

This STP information is shared among switches on the network using Bridge Protocol Data Units (BPDUs). Each BPDU contains the following information:

- The unique identifier of the switch that the transmitting switch currently believes is the root switch
- The path cost to the root from the transmitting port
- The port identifier of the transmitting port

The switch sends BPDUs to communicate and construct the spanning-tree topology. All switches connected to the LAN receive the BPDU. BPDUs are not directly forwarded by the switch, but the receiving switch uses the information in the frame to calculate a BPDU, and, if the topology changes, initiates a BPDU transmission.

The communication between switches via BPDUs results in the following:

- One switch is elected as the root switch
- The shortest distance to the root switch is calculated for each switch
- A designated switch is selected. This is the switch closest to the root switch through which packets will be forwarded to the root.
- A port for each switch is selected. This is the port providing the best path from the switch to the root switch.
- Ports included in the STP are selected.

Creating a Stable STP Topology

If all switches have STP enabled with default settings, the switch with the lowest MAC address in the network will become the root switch. By increasing the priority (lowering the priority number) of the best switch, STP can be forced to select the best switch as the root switch.

When STP is enabled using the default parameters, the path between source and destination stations in a switched network might not be ideal. For instance, connecting higher-speed links to a port that has a higher number than the current root port can cause a root-port change. The goal is to make the fastest link the root port.

STP Port States

The BPDUs take some time to pass through a network. This propagation delay can result in topology changes where a port that transitioned directly from a Blocking state

to a Forwarding state could create temporary data loops. Ports must wait for new network topology information to propagate throughout the network before starting to forward packets. They must also wait for the packet lifetime to expire for BPDU packets that were forwarded based on the old topology. The forward delay timer is used to allow the network topology to stabilize after a topology change. In addition, STP specifies a series of states a port must transition through to further ensure that a stable network topology is created after a topology change.

Each port on a switch using STP exists is in one of the following five states:

Blocking	The port is blocked from forwarding or receiving packets.
Listening	The port is waiting to receive BPDU packets that may tell the port to go back to the blocking state.
Learning	The port is adding addresses to its forwarding database, but not yet forwarding packets.
Forwarding	The port is forwarding packets.
Disabled	The port only responds to network management messages and must return to the blocking state first.
Transition States	A port transitions from one state to another as follows:

- From initialization (switch boot) to blocking
- From blocking to listening or to disabled
- From listening to learning or to disabled
- From learning to forwarding or to disabled
- From forwarding to disabled
- From disabled to blocking

Figure 5.4 below illustrates the STP port transition states.

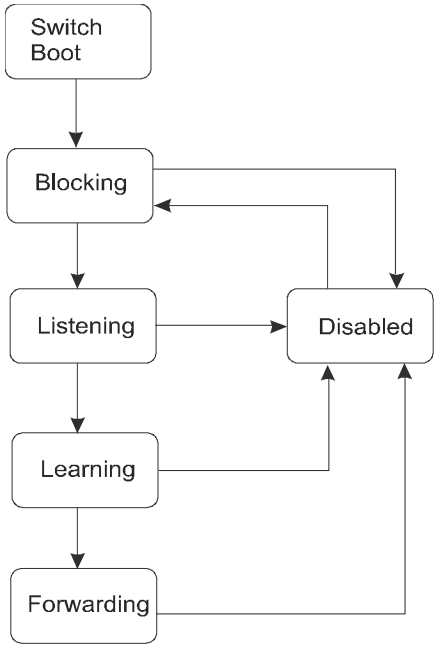

Figure 5-5: Port State Transition

When you enable STP, every port on every switch in the network goes through the blocking state and then transitions through the states of listening and learning at power up. If properly configured, each port stabilizes to the forwarding or blocking state.

No packets (except BPDUs) are forwarded from, or received by, STP enabled ports until the forwarding state is enabled for that port.

Default Spanning-Tree Configuration

The default Spanning Tree parameters are as follows:

Enable state	STP is enabled for all ports.
Port priority	128
Port cost	19
Bridge Priority	32,768

User-Changeable STP Parameters

The factory default setting should cover the majority of installations. However, it is advisable to keep the default settings as set at the factory; unless, it is absolutely necessary. The user changeable parameters in the switch are as follows:

Hello Time The Hello Time can be from 1 to 10 seconds. This is the interval between two transmissions of BPDU packets sent by the Root Bridge to tell all other switches that it is indeed the Root Bridge. If you set a Hello Time for your switch, and it is not the Root Bridge, the set Hello Time will be used if and when your switch becomes the Root Bridge. The Hello Time cannot be longer than the Max. Age. Otherwise, a configuration error will occur.

Max Age The Maximum Age Timer can be from 6 to 40 seconds. At the end of the Max. Age, if a BPDU has still not been received from the Root Bridge, the switch will start sending its own BPDU to all other switches for permission to become the Root Bridge. If it turns out the switch has the lowest Bridge Identifier, it will become the Root Bridge.

Forward Delay The Forward Delay can be from 4 to 30 seconds. This is the time any port on the switch spends in the listening state while moving from the blocking state to the forwarding state.

Priority A Priority for the switch can be set from 0 to 65535. 0 is equal to the highest Priority.

Observe the following formulas when setting the above parameters:

- Max. Age = 2 x (Forward Delay - 1 second)
- Max. Age = 2 x (Hello Time + 1 second)
- Port Priority A Port Priority can be from 0 to 255. The lower the number, the greater the probability the port will be chosen as the Root Port.
- Port Cost A Port Cost can be set from 1 to 65535. The lower the number, the greater the probability the port will be chosen to forward packets.

Illustration of STP

A simple illustration of three Bridges (or three switches) connected in a loop is depicted in Figure 5.5. In this example, you can anticipate some major network problems if the STP assistance is not applied. If Bridge A broadcasts a packet to Bridge B, Bridge B will broadcast it to Bridge C, and Bridge C will broadcast it to back to Bridge A ... and so on. The broadcast packet will be passed indefinitely in a loop, potentially causing a network failure.

STP can be applied as shown in Figure 5.6. In this example, STP breaks the loop by blocking the connection between Bridge B and C. The decision to block a particular connection is based on the STP calculation of the most current Bridge and Port settings. Now, if Bridge A broadcasts a packet to Bridge C, then Bridge C will drop the packet at port 2 and the broadcast will end there.

Setting-up STP using values other than the defaults, can be complex. Therefore, you are advised to keep the default factory settings and STP will automatically assign root bridges / ports and block loop connections. Influencing STP to choose a particular switch as the root bridge using the Priority setting, or influencing STP to choose a particular port to block using the Port Priority and Port Cost settings is, however, relatively straight forward.

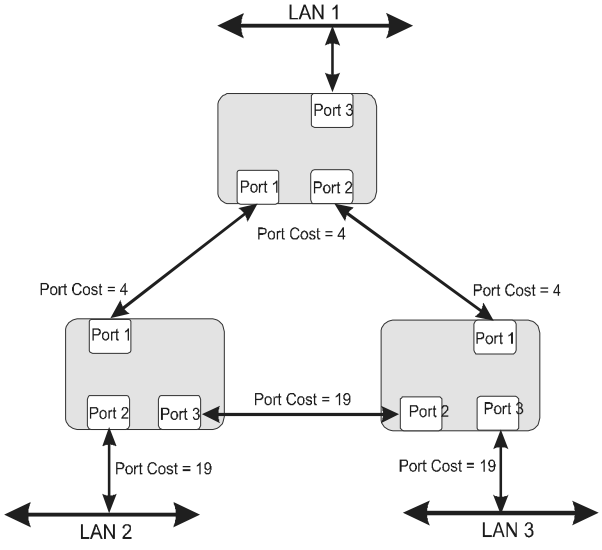

Figure 5-6: Sample Network without STP

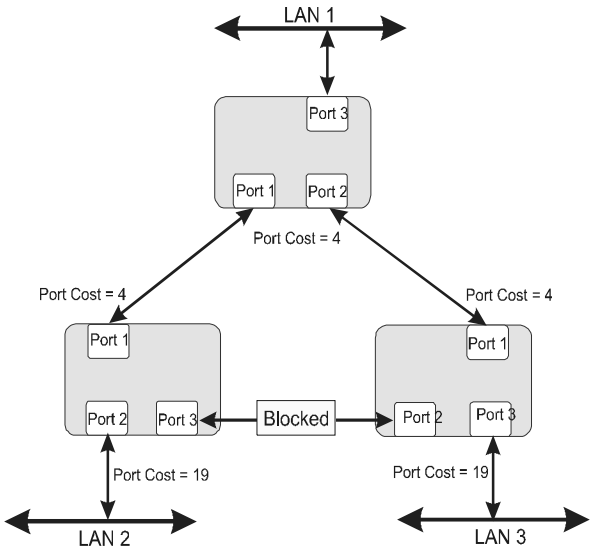

Figure 5-7: Sample Network using STP

The switch with the lowest Bridge ID (switch C) was elected the root bridge, and the ports were selected to give a high port cost between switches B and C. The two Gigabit ports (default port cost = 4) on switch A are connected to one Gigabit port on both switch B and C. The redundant link between switch B and C is deliberately chosen as a 100 Mbps Fast Ethernet link (default port cost = 19). Gigabit ports could be used, but the port cost should be increased from the default to ensure that the link between switch B and switch C is the blocked link.

Note also that the example network topology is intended to provide redundancy to protect the network against a link or port failure - not a switch failure or removal. For example, a failure of switch A would isolate LAN 1 from connecting to LAN 2 or LAN 3.

Port Aggregation

Port aggregation is used to combine a multiple ports to make a single high-bandwidth data pipeline. The participating ports are called members of a link aggregation group (sometime called a port trunk group), with one port designated as the master port of the group. Since all members of the link aggregation group must be configured to

operate in the same manner, the configuration of the master port is applied to all members of the link aggregation group. Thus, when configuring the ports in a link aggregation group, you only need to configure the master port.

The switch supports up to 6 link aggregation groups, which may include from 2 to 8 contiguous copper ports each. An additional Gigabit link aggregation group can be added for the two Gigabit ports.

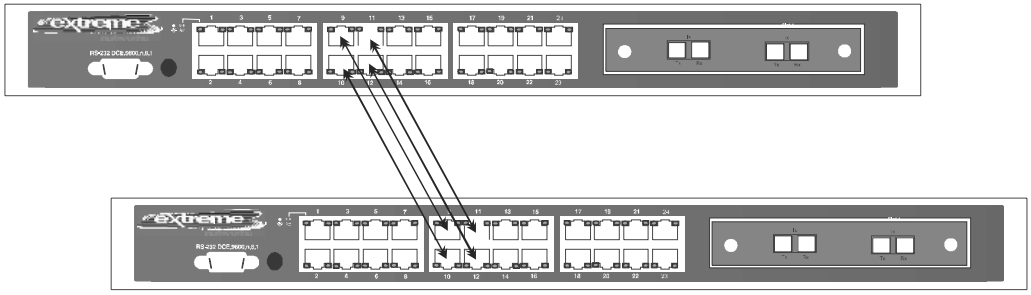

Figure 5-8: Link Agggregation

Data transmitted to a specific host (destination address) will always be transmitted over the same port in a trunk group. This allows packets in a data stream to arrive in the same order they were sent. A trunk connection can be made with any other switch that maintains host-to-host data streams over a single trunk port. Switches that use a load-balancing scheme that sends the packets of a host-to-host data stream over multiple trunk ports cannot have a trunk connection with the Switch.

VLANs

A Virtual Local Area Network (VLAN) is a network topology configured according to a logical scheme rather than the physical layout. VLANs can be used to combine any collection of LAN segments into an autonomous user group that appears as a single LAN. VLANs also logically segment the network into different broadcast domains so that packets are forwarded only between ports within the VLAN. Typically, a VLAN corresponds to a particular subnet, although not necessarily.

VLANs can enhance performance by conserving bandwidth, and improve security by limiting traffic to specific domains.

A VLAN is a collection of end nodes grouped by logic instead of physical location. End nodes that frequently communicate with each other are assigned to the same VLAN, regardless of where they are physically on the network. Logically, a VLAN can be equated to a broadcast domain, because broadcast packets are forwarded to only members of the VLAN on which the broadcast was initiated.

- No matter what basis is used to uniquely identify end nodes and assign these nodes VLAN membership, packets cannot cross VLANs without a network device performing a routing function between the VLANs.
- The switch supports only IEEE 802.1Q VLANs. The port untagging function can be used to remove the 802.1Q tag from packet headers to maintain compatibility with devices that are tag-unaware.
- By default the switch assigns all ports to a single 802.1Q VLAN named DEFAULT_VLAN. The DEFAULT_VLAN has a VID = 1.

Sharing Resources Across VLANs

Network resources such as printers and servers however, can be shared across 802.1Q VLANs. This is achieved by setting up overlapping (asymmetric) VLANs as shown in the diagram below.

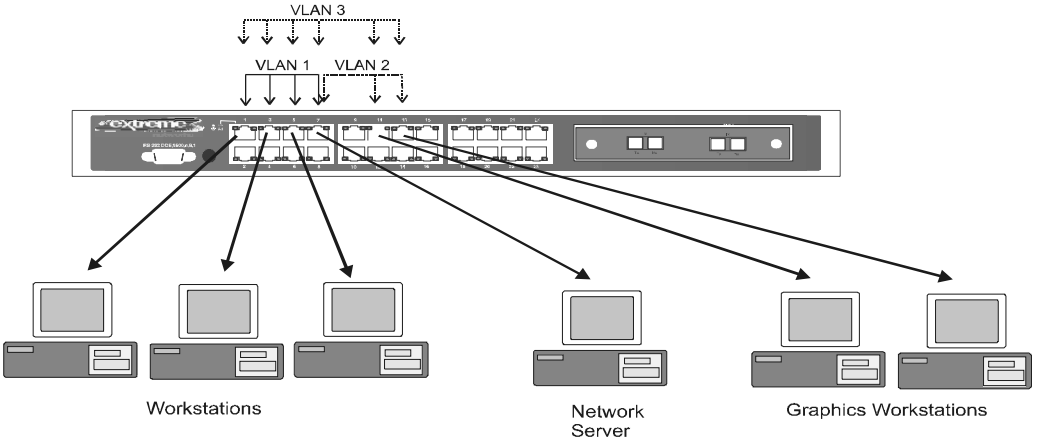

Figure 5-9: Sharing Resources Across VLANs

In the above example, there are three different 802.1Q VLANs and each port can transmit packets on one of them according to their Port VLAN ID (PVID). However, a port can receive packets on all VLANs (VID) that it belongs to.

Port-based VLANs

Port-based VLANs are a simplified version of the 802.1Q VLANs described in the previous section. In port-based VLANs, all the 802.1Q settings are pre-configured allowing you to quickly and easily setup and maintain port-based VLANs on your network.

In port-based VLANs, broadcast, multicast and unknown packets will be limited to within the VLAN. Thus, port-based VLANs effectively segment your network into broadcast domains. Furthermore, ports can only belong to a single VLAN.

Because port-based VLANs are uncomplicated and fairly rigid in their implementation, they are best used for network administrators who wish to quickly and easily setup VLANs in order to limit the effect of broadcast packets on their network.

For the most secure implementation, make sure that end stations are directly connected to the switch. Attaching a hub, switch or other repeater to the port causes all stations attached to the repeater to become members of the Port-based VLAN.

To setup port-based VLANs, simply select one of 64 VLAN ID numbers, name the VLAN and specify which ports will be members. All other ports will automatically be forbidden membership, even dynamically, as a port can belong to only one VLAN.

IEEE 802.1Q VLANs

To help you understand 802.1Q VLANs as implemented by the switch, it is necessary to understand the following:

Tagging	The act of putting 802.1Q VLAN information (a tag) into the header of a packet.
Untagging	The act of stripping 802.1Q VLAN information out of the packet header.
Ingress Port	A port on a switch where packets are flowing into the switch and VLAN decisions must be made.

Egress Port A port on a switch where packets are flowing out of the switch, either to another switch or to an end station, and tagging decisions must be made.

IEEE 802.1Q (tagged) VLANs are implemented on the Switch. 802.1Q VLANs require tagging, which enables them to span the entire network (assuming all switches on the network are IEEE 802.1Q-compliant).

VLANs allow a network to be segmented in order to reduce the size of broadcast domains. All packets entering a VLAN will only be forwarded to the stations (over IEEE 802.1Q enabled switches) that are members of that VLAN, and this includes broadcast, multicast and unicast packets from unknown sources.

VLANs can also provide a level of security to your network. IEEE 802.1Q VLANs will only deliver packets between stations that are members of the VLAN.

Any port can be configured as either tagging or untagging. The untagging feature of IEEE 802.1Q VLANs allow VLANs to work with legacy switches that don't recognize VLAN tags in packet headers. The tagging feature allows VLANs to span multiple 802.1Q-compliant switches through a single physical connection and allows Spanning Tree to be enabled on all ports and work normally.

The IEEE 802.1Q standard restricts the forwarding of untagged packets to the VLAN the receiving port is a member of.

The main characteristics of IEEE 802.1Q are as follows:

- Assigns packets to VLANs by filtering.
- Assumes the presence of a single global spanning tree.
- Uses an explicit tagging scheme with one-level tagging.

802.1Q Packet Forwarding Decisions

Packet forwarding decisions are made based upon the following three types of rules:

- **Ingress rules** - rules relevant to the classification of received frames belonging to a VLAN.
- **Forwarding rules between ports** - decides filter or forward the packet
- **Egress rules** - determines if the packet must be sent tagged or untagged.

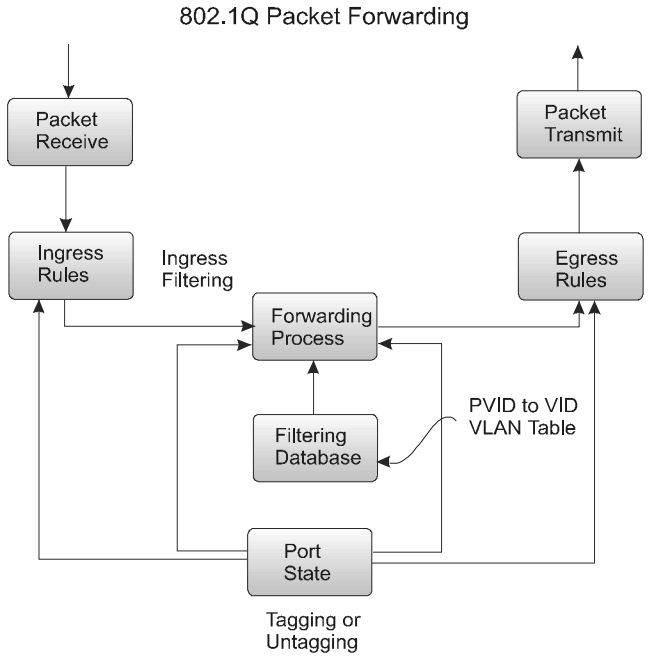

Figure 5-10: 802.1Q Packet Forwarding

802.1Q VLAN Tags

The figure below shows the 802.1Q VLAN tag. There are four additional octets inserted after the source MAC address. Their presence is indicated by a value of 0x8100 in the EtherType field. When a packet's EtherType field is equal to 0x8100, the packet carries the IEEE 802.1Q/802.1p tag. The tag is contained in the following two octets and consists of 3 bits or user priority, 1 bit of Canonical Format Identifier (CFI - used for encapsulating Token Ring packets so they can be carried across Ethernet backbones) and 12 bits of VLAN ID (VID). The 3 bits of user priority are used by 802.1p. The VID is the VLAN identifier and is used by the 802.1Q standard. Because the VID is 12 bits long, 4094 unique VLANs can be identified.

The tag is inserted into the packet header making the entire packet longer by 4 octets. All of the information contained in the packet originally is retained.

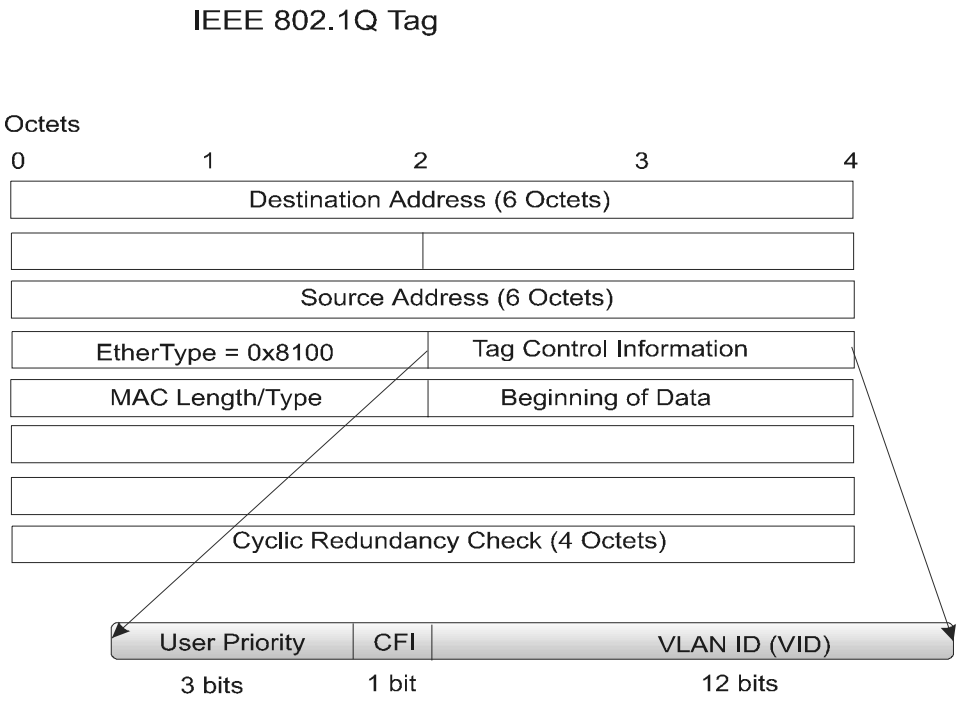

Figure 5-11: 802.1Q VLAN Tag

The EtherType and VLAN ID are inserted after the MAC source address, but before the originial EtherType/Length or Logical Link Control. Because the packet is now a bit longer than it was originally, the Cyclic Redundancy Check (CRC) must be recalculated.

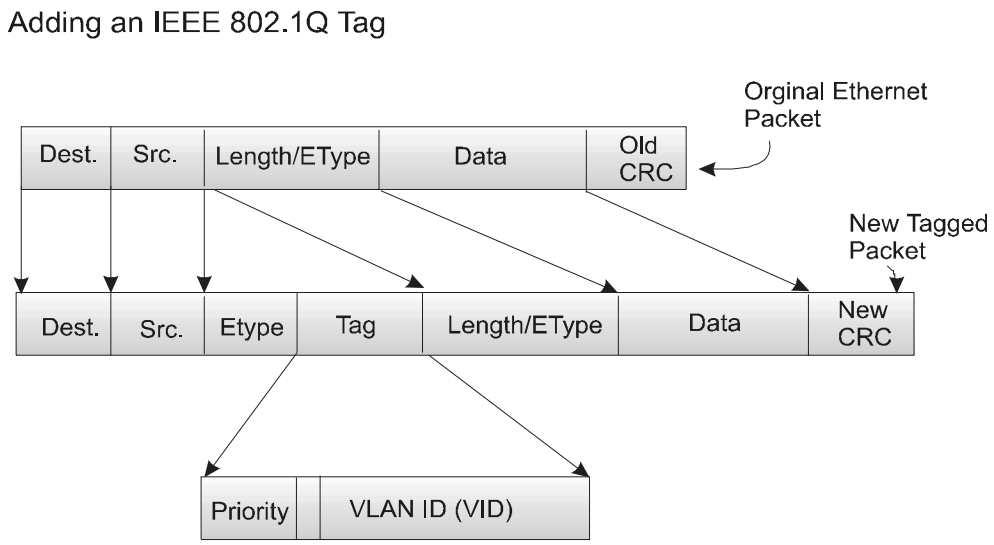

Figure 5-12: Adding 802.1Q Tag to a Packet Header

Port VLAN ID

Packets that are tagged (are carrying the 802.1Q VID information) can be transmitted from one 802.1Q compliant network device to another with the VLAN information intact. This allows 802.1Q VLANs to span network devices (and indeed, the entire network - if all network devices are 802.1Q compliant).

Unfortunately, not all network devices are 802.1Q compliant. These devices are referred to as tag-unaware. 802.1Q devices are referred to as tag-aware.

Prior to the adoption 802.1Q VLANs, port-based and MAC-based VLANs were in common use. These VLANs relied upon a Port VLAN ID (PVID) to forward packets. A packet received on a given port would be assigned that port's PVID and then be forwarded to the port that corresponded to the packet's destination address (found in the switch's forwarding table). If the PVID of the port that received the packet is different from the PVID of the port that is to transmit the packet, the switch will drop the packet.

Within the switch, different PVIDs mean different VLANs. (remember that two VLANs cannot communicate without an external router). So, VLAN identification based upon the PVIDs cannot create VLANs that extend outside a given switch (or switch stack).

Every physical port on a switch has a PVID. 802.1Q ports are also assigned a PVID, for use within the switch. If no VLANs are defined on the switch, all ports are then assigned to a default VLAN with a PVID equal to 1. Untagged packets are assigned the PVID of the port on which they were received. Forwarding decisions are based upon this PVID, in so far as VLANs are concerned. Tagged packets are forwarded according to the VID contained within the tag. Tagged packets are also assigned a PVID, but the PVID is not used to make packet forwarding decisions, the VID is.

Tag-aware switches must keep a table to relate PVIDs within the switch to VIDs on the network. The switch will compare the VID of a packet to be transmitted to the VID of the port that is to transmit the packet. If the two VIDs are different, the switch will drop the packet. Because of the existence of the PVID for untagged packets and the VID for tagged packets, tag-aware and tag-unaware network devices can coexist on the same network.

A switch port can have only one PVID, but can have as many VIDs as the switch has memory in its VLAN table to store them.

Because some devices on a network may be tag-unaware, a decision must be made at each port on a tag-aware device before packets are transmitted - should the packet to be transmitted have a tag or not? If the transmitting port is connected to a tag-unaware device, the packet should be untagged. If the transmitting port is connected to a tag-aware device, the packet should be tagged.

Tagging and Untagging Packets

Every port on an 802.1Q compliant switch can be configured as tagging or untagging.

Ports with tagging enabled will put the VID number, priority and other VLAN information into the header of all packets that flow into and out of it. If a packet has previously been tagged, the port will not alter the packet, thus keeping the VLAN information intact. The VLAN information in the tag can then be used by other 802.1Q compliant devices on the network to make packet forwarding decisions.

Ports with untagging enabled will strip the 802.1Q tag from all packets that flow into and out of those ports. If the packet doesn't have an 802.1Q VLAN tag, the port will not alter the packet. Thus, all packets received by and forwarded by an untagging port will have no 802.1Q VLAN information. (Remember that the PVID is only used internally within the switch). Untagging is used to send packets from an 802.1Q-compliant network device to a non-compliant network device.

Ingress Filtering

A port on a switch where packets are flowing into the switch and VLAN decisions must be made is referred to as an ingress port. If ingress filtering is enabled for a port, the switch will examine the VLAN information in the packet header (if present) and decide whether or not to forward the packet.

If the packet is tagged with VLAN information, the ingress port will first determine if the ingress port itself is a member of the tagged VLAN. If it is not, the packet will be dropped. If the ingress port is a member of the 802.1Q VLAN, the switch then determines if the destination port is a member of the 802.1Q VLAN. If it is not, the packet is dropped. If the destination port is a member of the 802.1Q VLAN, the packet is forwarded and the destination port transmits it to its attached network segment.

If the packet is not tagged with VLAN information, the ingress port will tag the packet with its own PVID as a VID (if the port is a tagging port). The switch then determines if the destination port is a member of the same VLAN (has the same VID) as the ingress port. If it does not, the packet is dropped. If it has the same VID, the packet is forwarded and the destination port transmits it on its attached network segment.

This process is referred to as ingress filtering and is used to conserve bandwidth within the switch by dropping packets that are not on the same VLAN as the ingress port at the point of reception. This eliminates the subsequent processing of packets that will just be dropped by the destination port.

Configuring VLANs

The switch initially configures one VLAN, VID = 1, called the DEFAULT_VLAN. The factory default setting assigns all ports on the switch to the DEFAULT_VLAN. As new VLANs are configured, there respective member ports are removed from the DEFAULT_VLAN.

Packets cannot be transmitted accross VLANs. If a member of one VLAN wants to connect to another VLAN, the link must be through an external router.

If no VLANs are configured on the switch all packets will be forwarded to any destination port. Packets with unknown source addresses will be flooded to all ports. Broadcast and multicast packets will also be flooded to all ports.

Broadcast Storms

Broadcast storms consist of broadcast packets that flood and/or are looped on a network causing noticeable performance degradation and in extreme cases, network failure. Broadcast storms can be caused by malfunctioning NICs, bad cable connections and applications or protocols that generate broadcast traffic, among others.

Broadcast storms have long been a concern for network administrators with routers traditionally being used to prevent their occurrence, and if that failed, limit their scope. However, with the advent of VLANs, switches are now able to limit broadcast domains better and cheaper than routers. Also, many switches, including the Summit24e2, have broadcast sensors and filters built into each port to further control broadcast storms.

Segmenting Broadcast Domains

VLANs can be used to segment broadcast domains. They do this by forwarding packets only to ports that are members of the same VLAN. Other parts of the network are effectively shielded. Thus, the smaller the broadcast domain, the smaller effect a broadcast storm will have. Because VLANs are implemented at each switch port, they can be quite effective in limiting the scope of broadcast storms.

Eliminating Broadcast Storms

SNMP agents can be programmed to monitor the number of broadcast packets on switch ports and act on the data. When the number of broadcast packets on a given port rise past an assigned threshold, an action can be triggered. When enabled, the usual action is to block the port from receiving broadcast packets. This will discard all broadcast frames arriving at the port from the attached segment. Not only does this isolate the broadcast domain, but it actually starts removing broadcast packets from the affected segment. When the number of broadcast packets falls to an acceptable level (below the trigger level), the SNMP agent can remove the blocking condition, returning the port to its normal operational state.

In the Summit24e2, the default trigger threshold is set to 128,000 broadcast packets per second (128 Kpps) for both 100 Mbps Fast Ethernet ports and the 1000 Mbps Gigabit Ethernet ports. The thresholds can be set separately for the two types of ports and can easily be modified by using a normal SNMP management program or through the console interface.

Multicasting

Multicasting enables a single network source to send packets to multiple destinations with persistent connections. The main advantage to multicasting is to decrease network load for communications that would otherwise use broadcasting.

Multicast Groups

There are three types of IP v4 addresses: unicast, broadcast, and multicast. Unicast addresses are used to transmit messages from a single network device to another, single network device. Broadcast packets are sent to all devices on the subnetwork. Multicast defines a group of network devices or computers that will receive the multicast packets. The members of this group are not necessarily on the same subnetwork. Specially designated multicast addresses are used to send multicast packets to the group members.

Multicast Addressing

Class D IP addresses are assigned to a group of network devices that comprise a multicast group. The four most significant four bits of a Class D address are set to "1110". The following 28 bits is referred to as the 'multicast group ID'. Some of the range of Class D addresses are registered with the Internet Assigned Numbers Authority (IANA) for special purposes. For example, the block of multicast addresses ranging from 224.0.0.1 to 224.0.0.225 is reserved for use by routing protocols and some other low-level topology discovery and maintenance protocols.

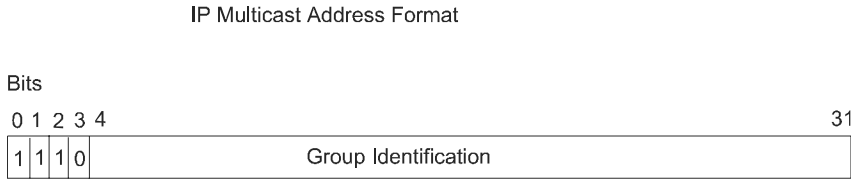

Figure 5-13: Class D Multicast Address

Internet Group Management Protocol (IGMP)

End users that want to receive multicast packets must be able to inform nearby routers that they want to become a multicast group member of the group these packets are being sent to. The Internet Group Management Protocol (IGMP) is used by multicast

routers to maintain multicast group membership. IGMP is also used to coordinate between multiple multicast routers that may be present on a network by electing one of the multicast routers as the 'querier'. This router then keep track of the membership of multicast groups that have active members on the network. IGMP is used to determine whether the router should forward multicast packets it receives to the subnetworks it is attached to or not. A multicast router that has received a multicast packet will check to determine if there is at least one member of a multicast group that has requested to receive multicast packets from this source. If there is one member, the packet is forwarded. If there are no members, the packet is dropped.

6 Using the Console Interface

The Summit24e2 24-port Ethernet switch supports a console management interface that allows you to set up and control your switch, either with an ordinary terminal (or terminal emulator) or over the network using the TCP/IP Telnet protocol. You can use this facility to perform many basic network management functions. In addition, the console program will allow you to set up the switch for management using an SNMP-based network management system. This chapter describes how to use the console interface to access the switch, change its settings, and monitor its operation. Included in this chapter are the following:

- Setting Up a Console
- Connecting to the switch Using Telnet
- Console Use Conventions
- First Time Connecting to the switch
- Logging On as a Registered User
- Setting Up the switch
- switch Utilities
- Network Monitoring

Console Screen Heirachary

The console screens and menus are organized as shown below (by title). The menus shown preceeded by a bullet (•) can be reached from the console's Main Menu. Menus

that can only be accessed from a previous menu (sub-menus) are shown indented with a dash (—) preceding the title of the sub-menu and are listed below the menu from which they can be accessed.

Basic Setup:

- switch Information
- Basic Network Setup
- Serial Port Settings
- Configure Ports
- Setup User Accounts
- Network Management Setup
  - SNMP Configuration
  - Management Station IP Setup
- Utilities
  - Download Firmware from TFTP Server
  - Download Configruation from TFTP Server
  - Upload Settings to TFTP Server
  - Upload History Log to TFTP Server
  - Ping Test
- Network Monitoring
  - Port Utilization
  - Port Error Packets
  - Port Packet Analysis
  - Browse MAC Address
  - ARP Table
  - Routing Table
  - Browse Router Port
  - IGMP Snooping
  - switch History
- Save Changes
- Reboot

  - Reboot
  - Save Configuration & Reboot
  - Reboot & Load Factory Default Configuration
  - Reboot & Load Factory Default Configuration Except IP Address
- Logout

Advanced Setup:

- Spanning Tree
  - Port Settings
- Forwarding
  - Unicast MAC Address Setting
  - Static Multicast Settings
  - Static/Default Routes
  - Static ARP
- Priority
  - Priority Queue
  - Setup MAC Address Priority
- Mirroring
  - Target Port Selection
  - Port Mirroring Settings
- IGMP Configuration
  - switch IGMP Snooping
  - IGMP Snooping Settings
  - Static Router Port Settings
- VLANs
  - Edit VLANs
- Link Aggregation
- Exception Handling

Setting Up a Console

First-time configuration must be carried out through a "console," that is, either (a) a VT100-type serial data terminal, or (b) a computer running communications software set to emulate a VT100. The console must be connected to the Diagnostics port. This is an RS-232 port with a 9-socket D-shell connector and DTE-type wiring. Make the connection as follows:

1 Obtain suitable cabling for the connection.

You can use an ordinary RS-232 cable. One end of the cable (or cable/adapter combination) must have a 9-pin D-shell connector suitable for the Diagnostics port; the other end must have a connector suitable for the management stations serial communications port.

2 Power down the devices, attach the cable (or cable/adapter combination) to the correct ports, and restore power.

3 Set the console to use 9600 baud the following communication parameters for your terminal:

- VT-100/ANSI compatible
- No parity checking (sometimes referred to as "no parity")
- 8 data bits (sometimes called a "word length" of 8 bits)
- 1 stop bit (sometimes referred to as a 1-bit stop interval)
- VT-100/ANSI compatible
- Arrow keys enabled

Connecting to the Switch Using Telnet

Once you have set an IP address for your switch, you can use a Telnet program (in a VT-100 compatible terminal mode) to access and control the switch. Most of the screens are identical, whether accessed from the console port or from a Telnet interface. You can also use a Web-based browser to manage the switch. See the next chapter, "Web-Based Network Management," for further information.

Console Usage Conventions

The console interface makes use of the following conventions:

1 Items in <angle brackets> can be toggled on or off using the space bar.

2 Items in [square brackets] can be changed by typing in a new value. You can use the Backspace and Delete keys to erase characters behind and in front of the cursor.

3 The up and down arrow keys, the left and right arrow keys, the Tab key and the Backspace key can be used to move between selected items. It is recommended that you use the Tab key and Backspace key for moving around the console.

4 Items in UPPERCASE are commands. Moving the selection to a command and pressing **Enter** will execute that command, e.g. APPLY, etc.

Note the command APPLY only applies for the current session. Use Save Changes from the Main Menu for permanent changes. An asterisk "" indicates that a change has been made but won't take effect until the switch has been rebooted.*

First Time Connecting To The switch

The switch supports user-based security that can allow you to prevent unauthorized users from accessing the switch or changing its settings. This section tells how to log onto the switch.

The passwords used to access the switch are case-sensitive; therefore, "S" is not the same as "s."

When you first connect to the switch, you will be presented with the first login screen (shown below). If this screen does not appear, Press **Ctrl+R** (hold down the Ctrl key, press and release the R key, and release Ctrl) to call it up. **Ctrl+R** can also be used at any time to refresh the screen.

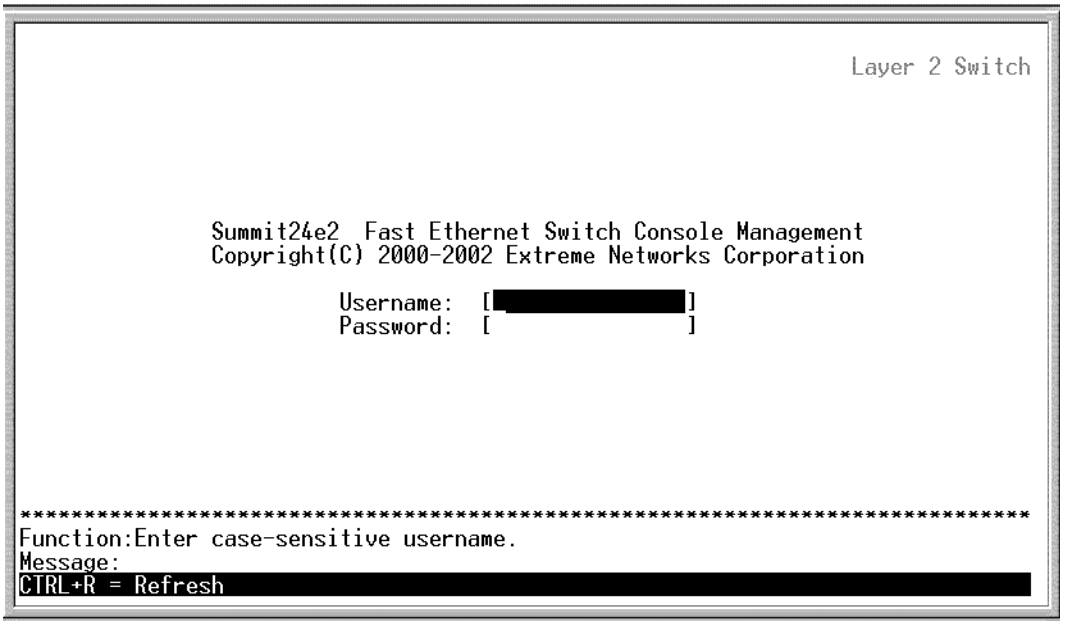

Figure 6-1: Initial Screen, First Time Connecting to the switch

The initial username and password are both admin. Specify admin in both the username and password fields when initially logging in.

Press **Enter** or **Return** in the Username and Password fields. You will be given access to the Main Menu **as** shown in Figure 6-1 below:

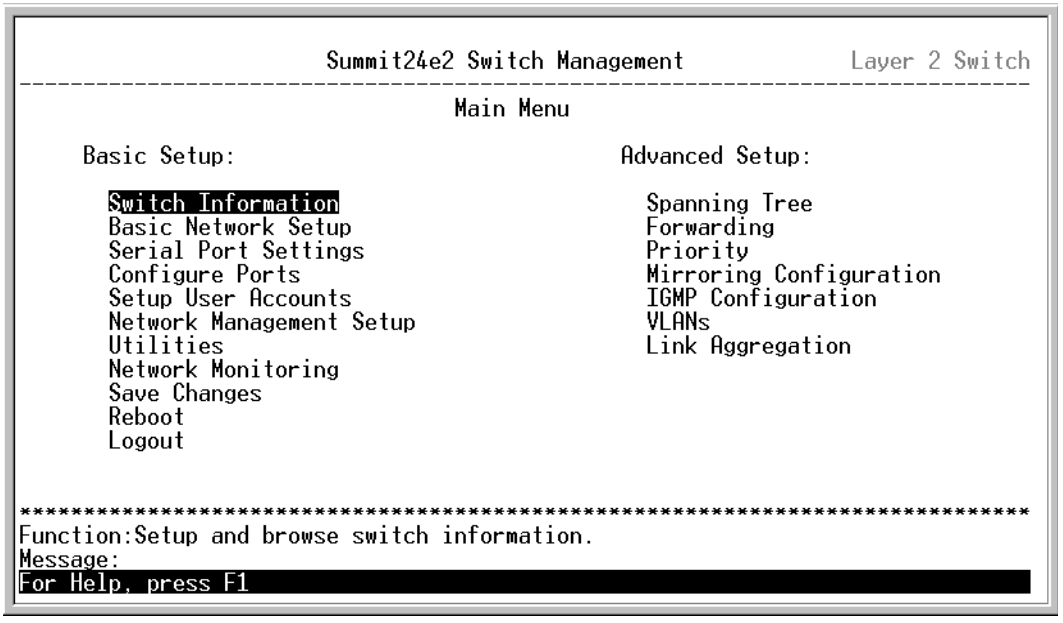

Figure 6-2: Main Menu

The first user automatically gets Administrator (Root) privileges (See Table 6-1). It is recommended to create at least one Administrator-level user for the switch.

User Accounts Management

From the screen above, move the cursor to the Setup User Accounts menu and press **Enter**. The Users Accounts menu appears.

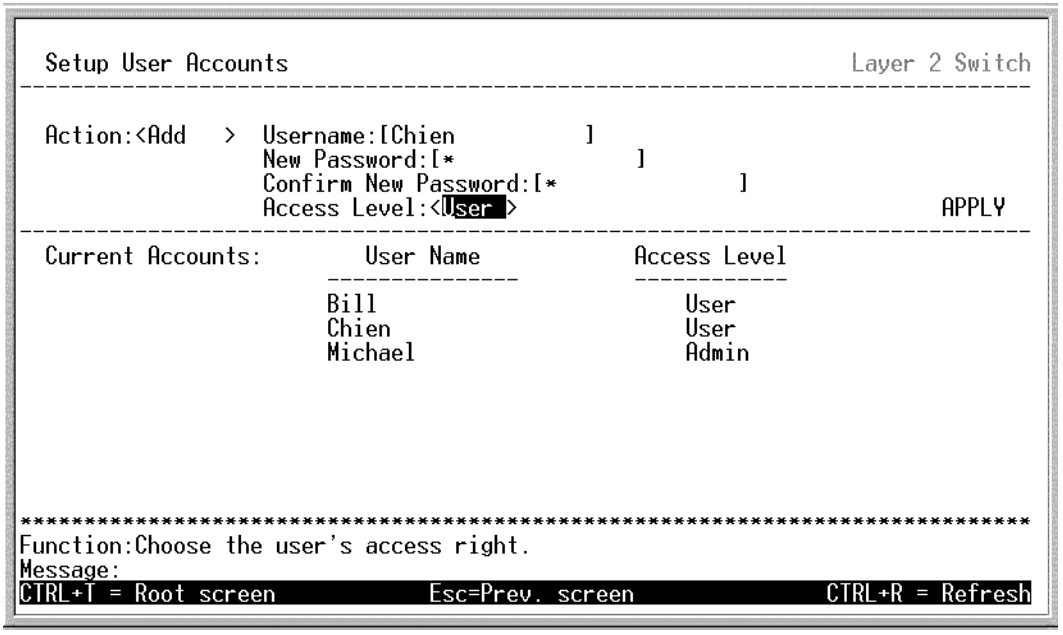

Figure 6-3: User Accounts Menu

1. Toggle the Action:<Add> field using the space bar to choose Add, Update, or Delete.
2. Type in the Username for the user account you wish to change and enter the Old Password for that user account.
3. You can now modify the password or the privilege level for this user account.
4. If the password is to be changed, type in the New Password you have chosen, and press **Enter**. Type in the same new password in the following field to verify that you have not mistyped it.
5. If the privilege level is to be changed, toggle the Access Level:<Root> field until the appropriate level is displayed - Root, User+ or User.
6. Highlight APPLY and press enter to make the change effective.
7. You must enter the configuration changes into the non-volatile ram (NV-RAM) using Save Changes from the Main Menu if you want the configuration to be used after a switch reboot.
8. Press **Esc** to return to the previous screen or **Ctrl+T** to go to the root screen

Root, User+ and Normal User Privileges

There are three levels of user privileges: Root , User+ and User. Some menu selections available to users with Root privileges may not be available to those with User+ or User privileges.

The following table summarizes Root, User+ and User privileges. The menus shown are the menus for the two types of users:

Table 6-1: Root, User+ and User Privileges

Management	Root Privilege	User+ Privileges	User Privilege
Configuration	Yes	Yes, view only	Yes, view only
Network Monitoring	Yes	Yes, view only	Yes, view only
Community Strings and Trap Stations	Yes	Yes, view only	Yes, view only
Update Firmware and Configuration Files	Yes	No	No
System Utilities	Yes	Yes, Ping only	Yes, Ping only
Factory Reset	Yes	No	No
Restart System	Yes	Yes	No
User Accounts Management			
Add/Update/Delete User Accounts	Yes	No	No
View User Accounts	Yes	Yes	Yes

After establishing a User Account with Administrator-level privileges, press **Esc** twice. Then choose the Save Changes menu (see below). Pressing any key will return to the Main Menu. You are now ready to operate the switch.

Saving Changes

The Summit24e2 has two levels of memory, normal RAM and non-volatile or NV-RAM. Settings need to be changed in all screens by choosing APPLY and pressing **Enter**. When this is done, the settings will be immediately applied to the switching software in RAM, and will immediately take effect. Some settings, though, require you to restart the

switch before they will take effect. Restarting the switch will erase all settings in RAM and reload them from the NV-RAM. Thus, it is necessary to save all settings to the NV-RAM before restarting the switch.

To retain any modifications made in the current session by saving them into the NV-RAM, you must choose Save Changes from the Main Menu. The following screen will appear to indicate your new settings have been saved:

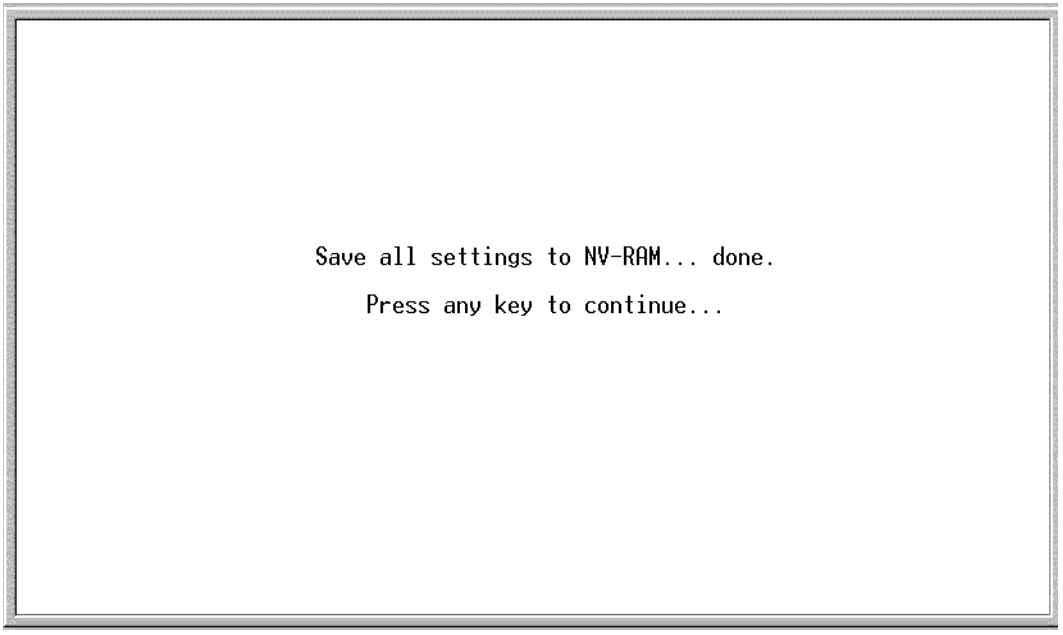

Figure 6-4: Save Changes Confirmation Screen

After the settings have been saved to NV-RAM, they will become the default settings for the switch, and they will be used by the switch every time it is powered on, reset, or rebooted. The only exception to this is a factory reset, which will clear all settings and restore them to their initial values listed in the appendix, which were present when the switch was purchased.

Logging On As a Registered User

To log in once you have created a registered user:

1 Type in your username and press **Enter**.

2 Type in your password and press **Enter**.

3 The Main Menu screen will be displayed based on your Administrator or Normal User access level or privilege.

Setup User Accounts

To add a new user:

1 Choose Setup User Accounts from the Main Menu. The following menu appears:

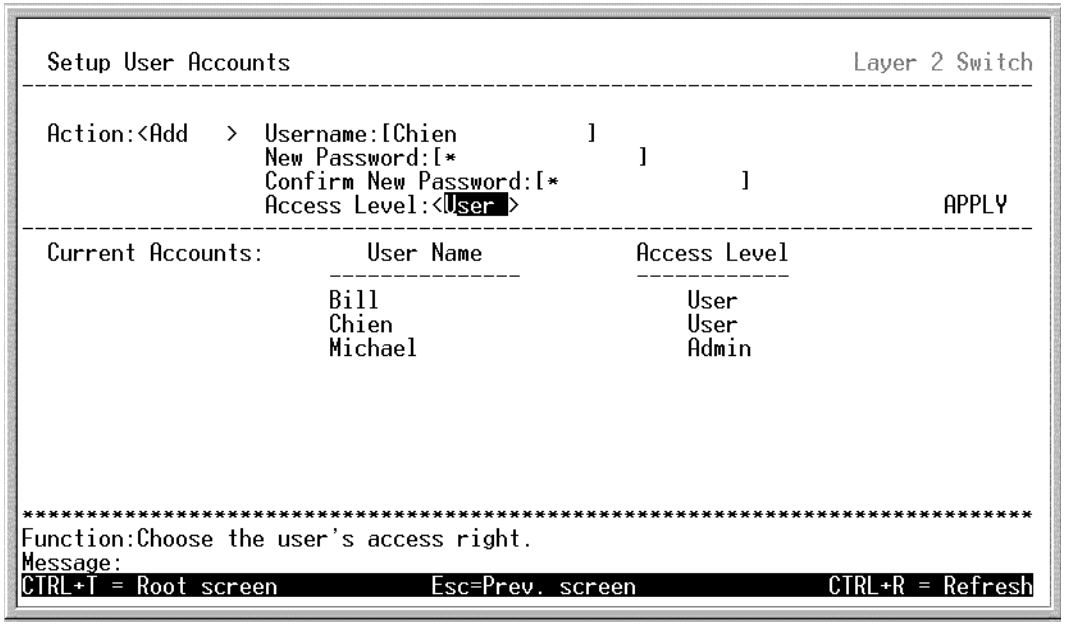

Figure 6-5: Setup User Accounts Menu

2 Toggle the Action:<Add> field using the space bar to choose Add, Update, or Delete.

3 Type in the Username for the user account you wish to add.

4 Enter the New Password for that user account.

5 Enter the same password again in Confirm New Password to confirm the choice of password.

6 Choose the privilege level by toggling the Access Level:<Root> field until the appropriate level is displayed - Root, User+ or User.

7 Highlight APPLY and press enter to make the change effective.

8 You must enter the configuration changes into the non-volatile RAM (NV-RAM) using Save Changes from the Main Menu if you want the configuration to be used after a switch reboot.

9 Press **Esc** to return to the previous screen or **Ctrl+T** to go to the root screenType in your Username and press **Enter**.

Modifying user names and password for user accounts can only be done by a Root-level user.

Update/Delete User Accounts

Access to the console, whether using the console port or via TELNET, is controlled using a user name and password. Up to eight of these user names can be defined. The console interface will not let you delete the current logged-in user, however, in order to prevent accidentally deleting all of the users with Root privileges.

Only users with the Root privilege can delete users.

To view a user account:

Choose Setup User Accounts from the Main menu. The following screen appears:

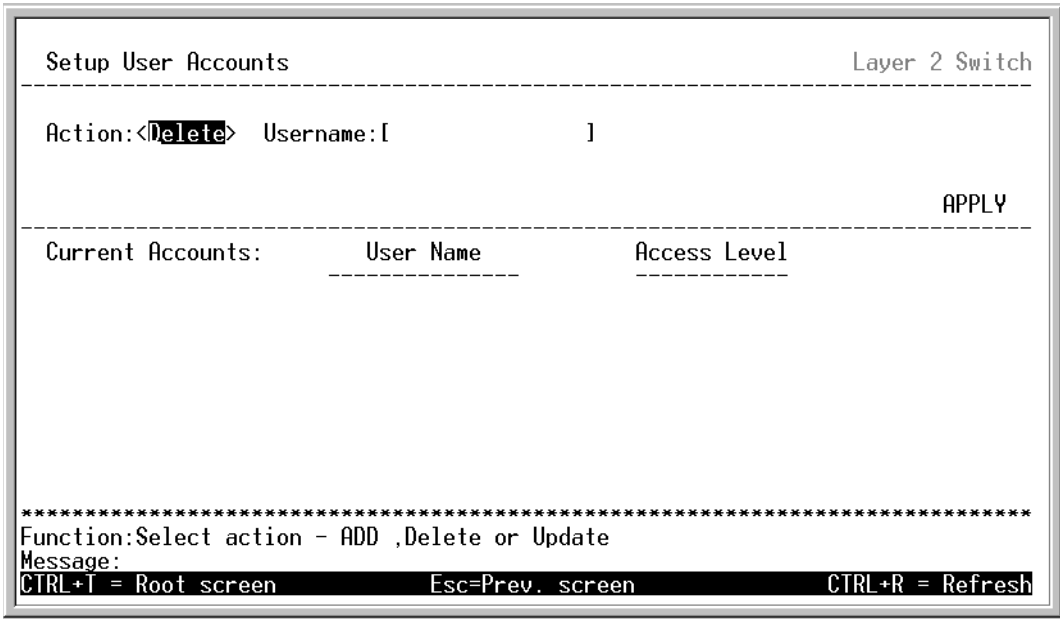

Figure 6-6: View/Delete User Accounts Menu

To delete a user:

1. Toggle the Action:<Delete> field using the space bar to choose Add, Update, or Delete.
2. Type in the Username for the user account you wish to delete.
3. Select APPLY and press **Enter** to let the user deletion take effect.

Setting Up the switch

switch management functions are grouped into two major groups in the console, Basic Setup and Advanced functions. The remaining sections of this chapter deal with how you can use the console to setup these functions to implement an efficient network management strategy.

Basic Setup

This section will help prepare the switch user by describing the following menus and their sub-menus:

- Switch Information
- Basic Network Setup
- Serial Port Settings
- Configure Ports
- Setup User Account
- Network Management Setup
  - SNMP Configuration
  - Management Station IP Setup
- Utilities
  - Download Firmware from TFTP Server
  - Download Configuration from TFTP Server
  - Upload Settings to TFTP Server
  - Upload History Log to TFTP Server
  - Ping Test
- Network Monitoring
  - Port Utilization
  - Port Error Packets
  - Port Packet Analysis
  - Browse MAC Address
  - Browse ARP Table
  - Browse Routing Table
  - Browse Router Port
  - IGMP Snooping
  - switch History
- Save Changes
- System Reboot

- Logout

Switch Information

Choose switch Information to access the first item on the Summit24e2 Main Menu. The following menu appears:

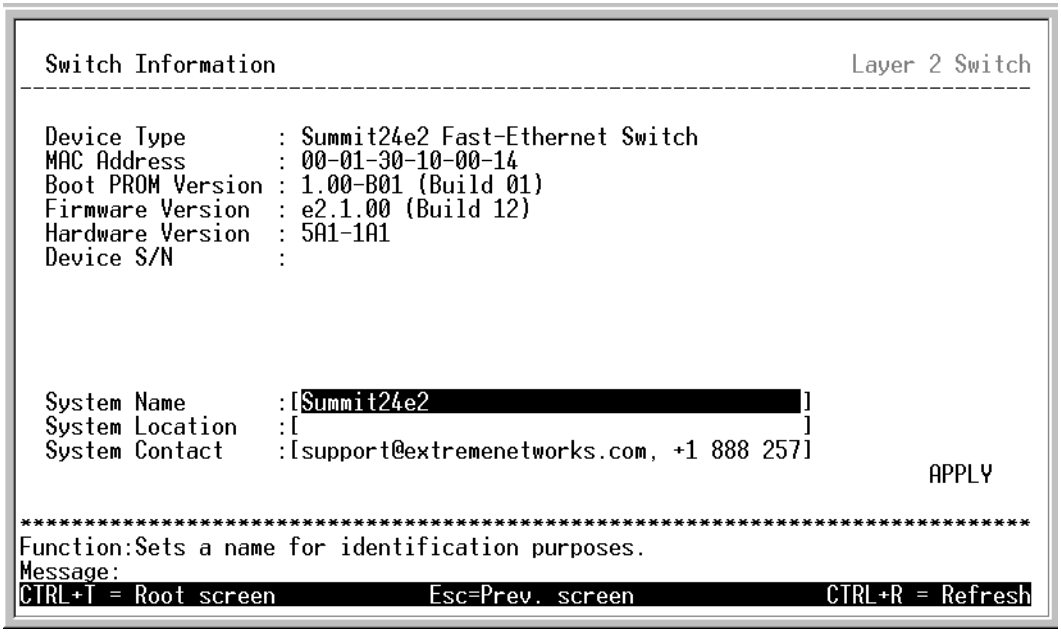

Figure 6-7: Switch Information Menu

The switch Information menu displays the type of switch (Summit24e2), which (if any) external modules are installed, and the switch's MAC address (assigned by the factory and unchangeable). In addition, the Boot PROM and Firmware Version numbers are shown. This information is helpful to keep track of PROM and Firmware updates and to obtain the switch's MAC address for entry into another network device's address table - if necessary.

You can also enter the name of the system, its location, and the name and telephone number of the System Administrator. It is recommended that the person responsible for the maintenance of the network system that this switch is installed on be listed here.

Basic Network Setup

Some settings such as the switch IP address and subnet mask must be entered to allow the switch to be managed from an SNMP-based Network Management System or to be able to access the switch using the TELNET protocol.

The Basic Network Setup menu lets you specify how the switch will be assigned an IP address to allow the switch to be identified on the network. In addition, you may specify a subnet mask and default gateway.

Highlight Basic Network Setup to access the first item on the Configuration menu. The following screen appears:

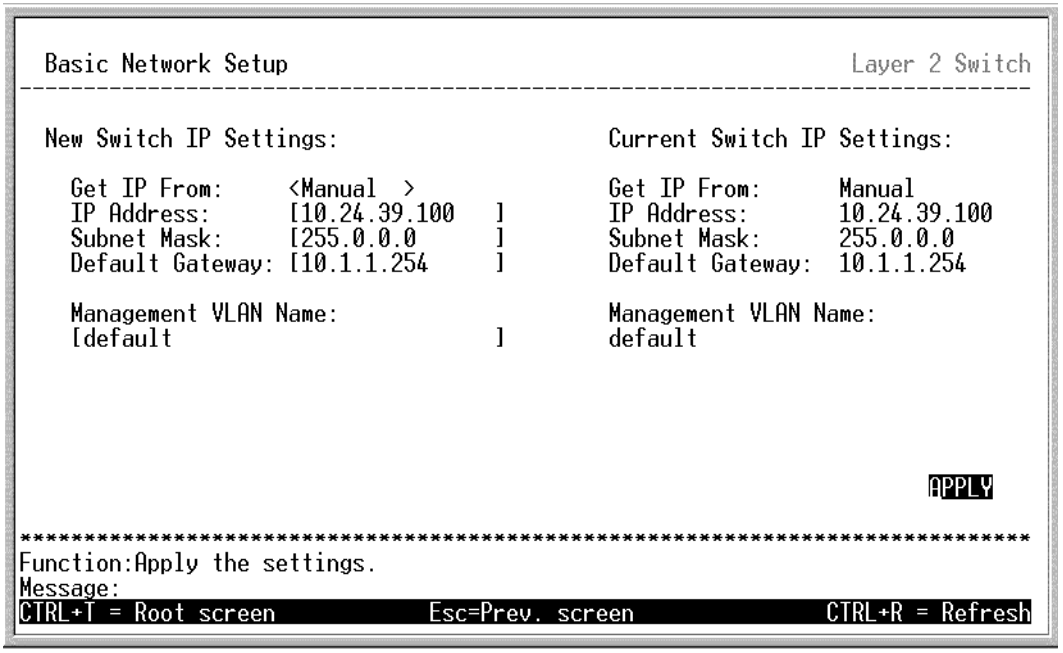

Figure 6-8: Basic Network Setup Menu

The fields listed under the New switch Settings heading are those that are currently being used by the switch. Those fields listed under the Restart Settings heading are those which will be used after the switch has been reset. Fields that can be set are:

Get IP Address From Determines whether the switch should get its IP address settings from the user (Manual), a BOOTP server, or a DHCP server. If Manual is chosen, the switch will use the IP address, Subnet Mask and Default Gateway settings defined in this screen after saving the changes and rebooting. If BOOTP is chosen, the switch will send out a BOOTP broadcast request when it is powered up. The BOOTP protocol allows IP addresses, network masks, and default gateways to be assigned by a central BOOTP server. If this option is set, the switch will get its IP settings from the BOOTP server upon being rebooted. If DHCP is chosen, a Dynamic Host Configuration Protocol request will be sent when the switch is rebooted.

IP Address Determines the IP address used by the switch for receiving SNMP and TELNET communications. These fields should be of the form xxx.xxx.xxx.xxx, where each xxx is a number (represented in decimal) between 0 and 255. This address should be a unique address on a network assigned to you by the central Internet authorities.

Subnet Mask Bitmask that determines the extent of the subnet that the switch is on. Should be of the form xxx.xxx.xxx.xxx, where each xxx is a number (represented in decimal) between 0 and 255. If no subnetting is being done, the value should be 255.0.0.0 for a Class A network, 255.255.0.0 for a Class B network, and 255.255.255.0 for a Class C network.

Default Gateway IP address that determines where frames with a destination outside the current IP subnet should be sent. This is usually the address of a router or a host acting as an IP gateway. If your network is not part of an inter-network, or you do not want the switch to be accessible outside your local network, you can leave this field unchanged.

Management VLAN Name Allows a management VLAN Name to be entered. This allows switch management from a host within that VLAN to use either TELNET or an SNMP-based network manager. The default VLAN Name is default which includes the entire switch until VLANs are configured.

Configuring the Serial Port

You can use the Serial Port Settings screen to enter settings for the switch's RS-232C serial port for console management.

To configure the serial port, highlight Serial Port Settings from the Main Menu and press **Enter**. The following screen appears:

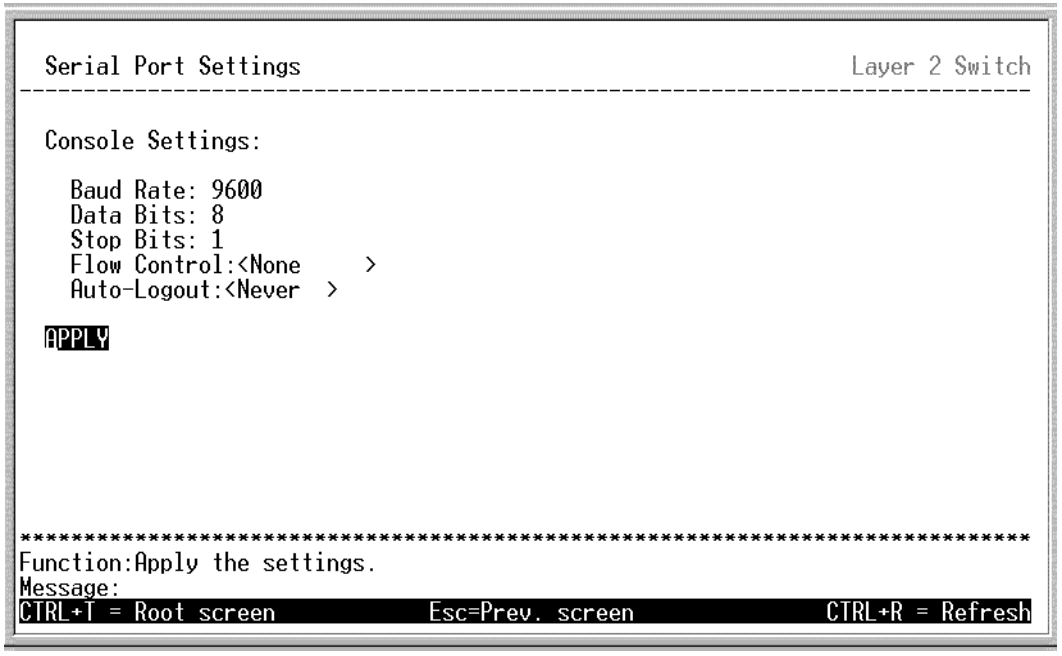

Figure 6-9: Serial Port Settings Screen

The following fields can then be set:

Baud Rate Displays the serial bit rate that will be used to communicate with the management host. The default setting is **9600**, and cannot be changed.

Data Bits Displays the number of bits in the data stream that will represent data. The default is **8** data bits and cannot be changed.

Flow Control: <None> Sets the flow control scheme for the switch's RS-232C port. Can be set to **None**, or **XOn/XOff**. This setting must match the setting on the host computer's RS-232C port, and is set to **None** by default.

Auto-Logout Sets the time the interface can be idle before the switch automatically logs-out the user. The options are 2 mins, 5 mins, 10 mins, 15 mins, or Never.

Configure Ports

The Configure Ports screen allows you to change port state, the speed/duplex combination and flow control of all ports. If you are configuring Gigabit Ethernet ports, use the same screen. The options for configuring Gigabit Ethernet ports differ from the 10/100 Mbps ports, therefore instructions for configuring the Gigabit ports are presented in a separate section below. To configure the 10/100 Mbps ports, highlight Port Configuration on the Main Menu and the following screen will appear:

```
 Configure Ports                                                  Layer 2 Switch
--------------------------------------------------------------------------------
 View Ports:<1 to 12  >  Configure Port from [1 ] to [1 ] State:<Enabled >
 Speed/Duplex:<Auto     > Flow Control: Off                                APPLY
--------------------------------------------------------------------------------
 Port     State       Settings         Connection
 -----   --------  ---------------     ---------------
 1       Enabled     Auto/Enabled       100M/Full/802.3x
 2       Enabled     Auto/Enabled           Link Down
 3       Enabled     Auto/Enabled           Link Down
 4       Enabled     Auto/Enabled           Link Down
 5       Enabled     Auto/Enabled           Link Down
 6       Enabled     Auto/Enabled           Link Down
 7       Enabled     Auto/Enabled           Link Down
 8       Enabled     Auto/Enabled           Link Down
 9       Enabled     Auto/Enabled           Link Down
 10      Enabled     Auto/Enabled           Link Down
 11      Enabled     Auto/Enabled           Link Down
 12      Enabled     Auto/Enabled           Link Down
********************************************************************************
Function:Select the scope of ports for display and configuration.
Message:
CTRL+T = Root screen          Esc=Prev. screen              CTRL+R = Refresh
```

Figure 6-10: Configure Ports Screen

Configure the ports by selecting options in the following fields:

View Ports — Toggle the **View Ports:< >** field, using the space bar, to view the configuration of ports 1-12, ports 13-24 or Slot-1. To configure the Slot Module ports, see the instructions below.

Configure Ports from [ ] to [ ] — To configure a specific port, or a sequential group of ports within the range defined in the View Ports field, use the **Configure Port from [ ] to [ ]** field. To configure a single port, type the port number in both the upper and lower limit of the range.

State — Toggle the **State:< >** field to either Enable or Disable the specified port or sequential port group.

Speed/Duplex — Toggle the **Speed/Duplex:< >** field to select the speed and duplex state of the port. Auto - auto-negotiation, negotiates settings for operation with 10 and 100 Mbps devices, in full- or half-duplex modes. The Auto setting allows the port to automatically determine the fastest settings of the device to which the port is connected, and then to use those settings. For the 24 copper ports choose Auto, 100M / Full, 100M / Half, 10M / Full, 10M / Half. There is no automatic adjustment of port settings with any option other than Auto.

Flow Control — For 10 / 100 Mbps ports configured as half-duplex, backpressure is always enabled.

For 10 / 100 Mbps ports configured as full-duplex, no flow control is supported.

Configure the Gigabit Ethernet Ports

You can configure the Gigabit Ethernet ports using the Configure Ports Screen. The nature of Gigabit Ethernet requires that additional factors be considered during configuration In addition, the type of port will effect how these settings function.

Configure the Gigabit Ethernet ports using the same fields in the Configure Ports screen with the following considerations:

View Ports — Toggle the **View Ports:< >** field to select Slot-1.

Configure Ports from [ ] to [ ] Use the **Configure Port from [ ] to [ ]** field to configure one or both of the Slot 1 ports. Enter 1 for Slot Module Port 1 (S1P1), and 2 for Slot Module Port 2 (S1P2).

State Toggle the **State:< >** field to either Enable or Disable the specified port or ports.

Speed/Duplex For most circumstances, the default setting Auto will be appropriate. However, if the switch is used on a network that employs certain older Gigabit Ethernet devices, it may be necessary to force 1000 Mbps full-duplex operation. Consult the specifications of these older devices for more information. In addition, the type of port effects how the Auto setting functions. 1000BASE-TX ports are capable of true Auto-negotiation in that they can operate at 10, 100 or 1000 Mbps and in full- or half-duplex mode. GBIC ports will only operate at 1000 Mbps in full-duplex mode. Toggle the **Speed/Duplex:< >** field to either select the either Auto or 1000MFULL.

Flow Control For 1000 Mbps ports, flow control can be disabled. If the Flow Control:<Off> field is toggled to “Rx”, then the Summit24e2 switch will not transmit 802.3x PAUSE frames when it becomes congested, but will honor PAUSE frames it receives from another switch’s connected port.

Network Management Setup

Highlight **Network Management Setup** from the **Main Menu** and press **Enter.**

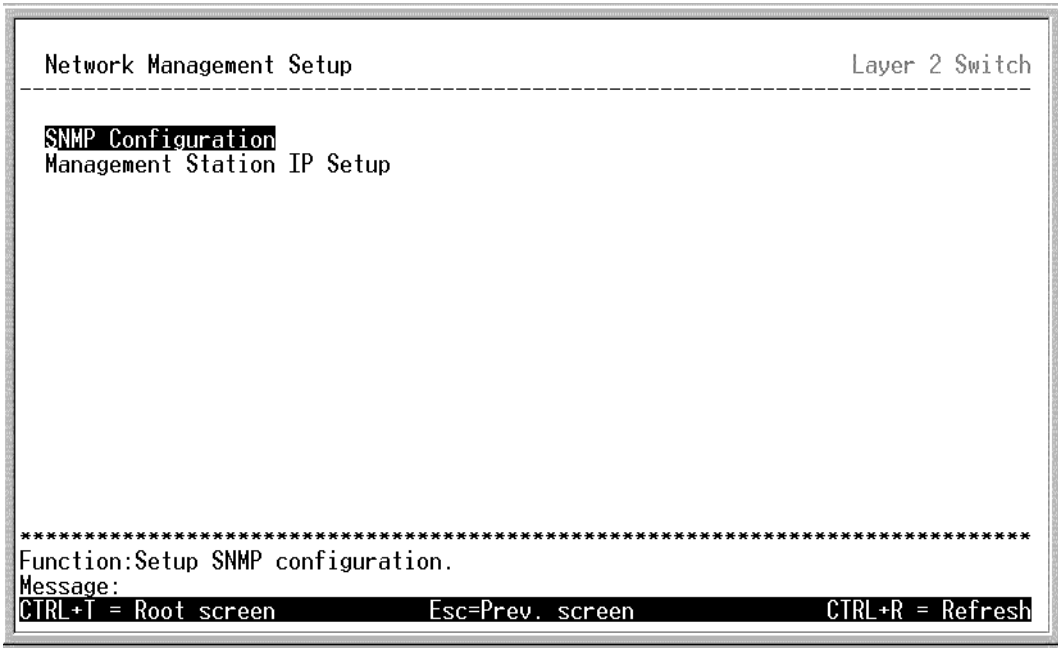

Figure 6-11: Network Management Setup

SNMP Configuration

Highlight **SNMP Configuration** from the **Network Management Setup Menu** and press **Enter.**

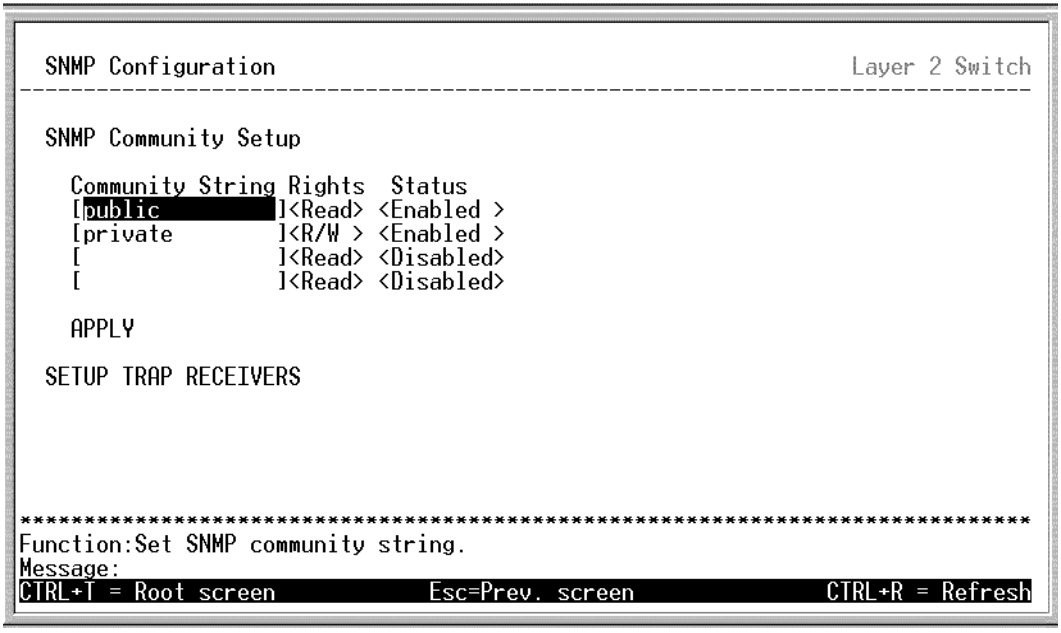

Figure 6-12: SNMP Configuration

Setting Trap Receivers

The Setup Trap Receivers feature allows the switch to send traps (messages about errors, etc.) to management stations on the network. Highlight Setup Trap Receivers and press **Enter**. The trap recipients can be setup from the following screen:

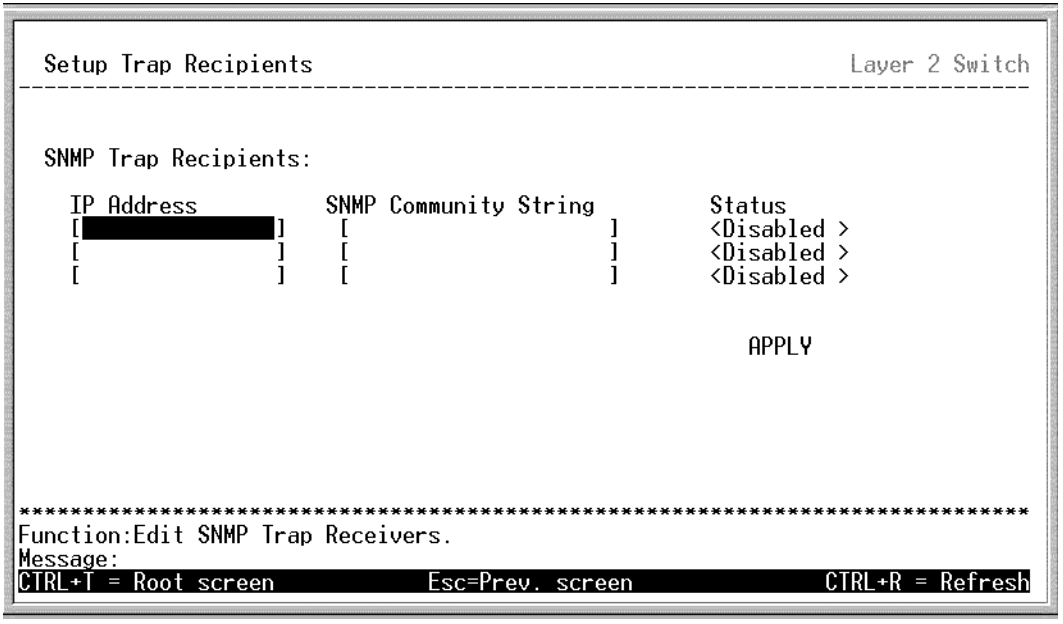

Figure 6-13: Setup SNMP Trap Recievers Menu

Fields that can be set in the Setup Trap Recipients menu include:

IP Address	The IP address of a management station (usually a computer) that is configured to receive the SNMP traps from the switch.
SNMP Community String	Similar to a password in that stations that do not know the correct string cannot receive or request SNMP information from the switch.
Status	Toggle between Enabled and Disabled to enable or disable the receipt of SNMP traps by the listed management stations.

Up to four SNMP trap recipients can be entered.

Management Station IP Address Setup

The switch allows you to specify up to four IP addresses that are allowed to manage the switch via in-band SNMP or TELNET based management software. These IP addresses must be members of the Management VLAN.

If no IP addresses are specified, then there is nothing to prevent any IP address from accessing the switch, provided the user knows the Username and Password.

Highlight Management Station IP Setup from the SNMP Configuration Menu and press **Enter.**

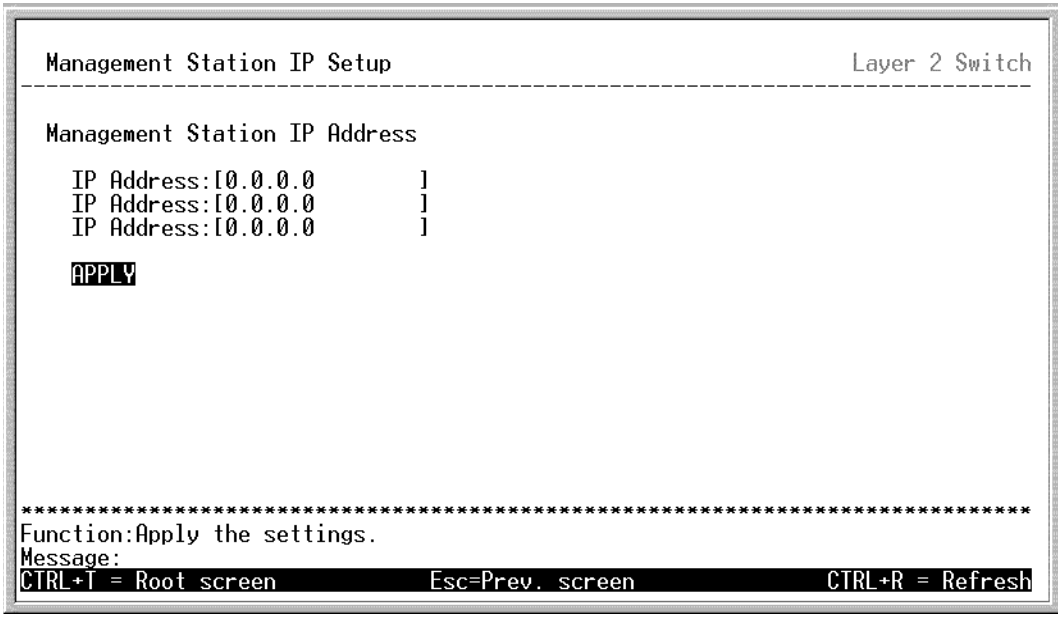

Figure 6-14: Management Station IP Setup

Switch Utilities

When it is necessary to upgrade firmware or use a saved switch configuration file located on a TFTP (Trivial File Transfer Protocol) server, use the switch Utilities menu to perform these operations. Additionally, you can save switch configuration files and switch history logs to a TFTP server. The switch Utilities menu is also where you can perform a Ping test to confirm connectivity and communication with other devices on the network.

Highlight Switch Utilities on the Main Menu to and press **Enter** to access the menu:

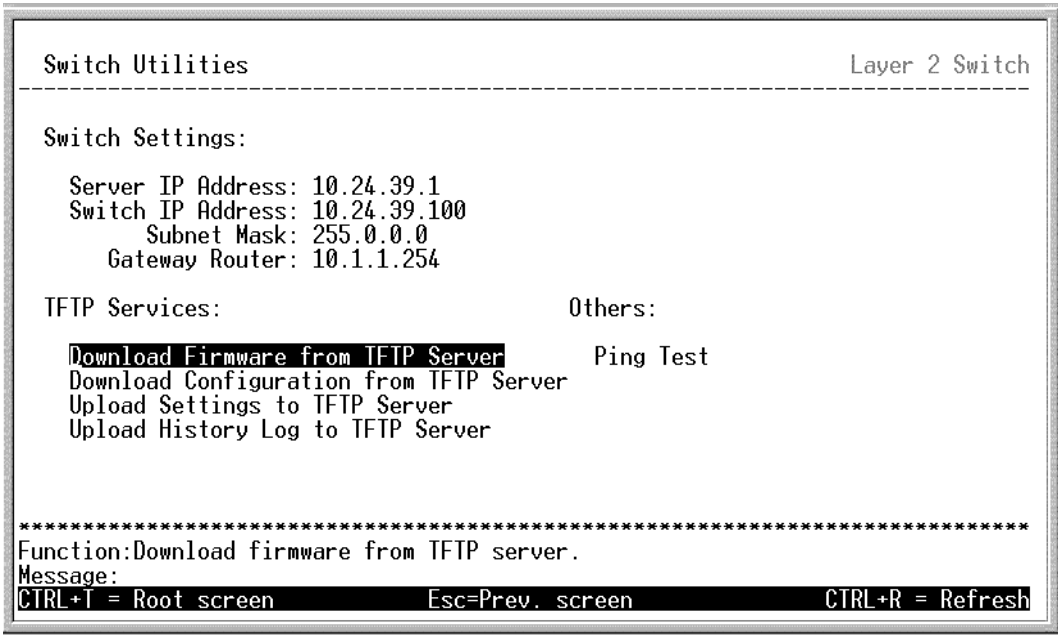

Figure 6-15: Switch Utilities Menu

Upgrade Firmware from a TFTP Server

To download and install a new firmware file via a TFTP server:

Highlight Upgrade Firmware from TFTP Server on the switch Utilities Menu and press **Enter** to see the Upgrade Firmware screen:

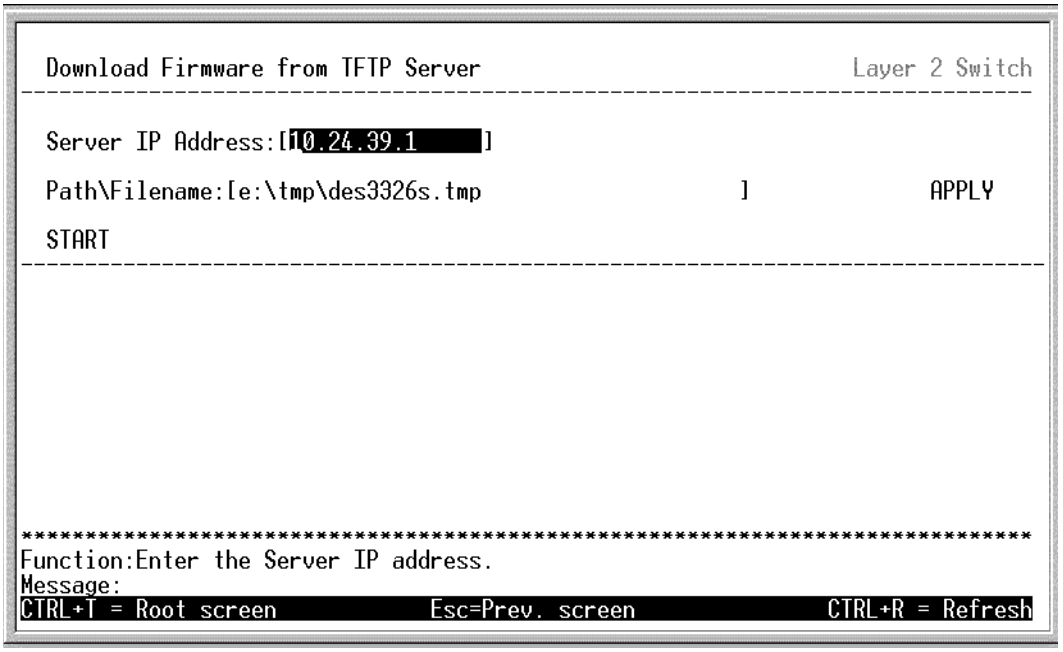

Figure 6-16: Download Firmware from TFTP Server

Trivial File Transfer Protocol (TFTP) services allow the switch firmware to be upgraded by transferring a new firmware file from a TFTP server to the switch. A configuration file can also be loaded into the switch from a TFTP server, switch settings can be saved to the TFTP server, and a history log can be uploaded from the switch to the TFTP server.

Enter the IP address of the TFTP server in the Server IP address:[ ] field.

The TFTP server must be on the same IP subnet as the switch.

Enter the path and the filename to the firmware file on the TFTP server in the Path\Filename: [ ] field. In the above example, the firmware file is in the root directory of the C drive of the TFTP server.

The TFTP server must be running TFTP server software to perform the file transfer. TFTP server software is a part of many network management software packages, or can be obtained as a separate program.

4 Highlight APPLY and press **Enter** to record the IP address of the TFTP server.

5 Highlight START and press **Enter** to initiate the file transfer.

Download a Configuration File From a TFTP Server

Downloading a configuration file from a TFTP server is a procedure similar to upgrading firmware via a TFTP server. To download a switch configuration file from a TFTP server:

1 Highlight Use Configuration File on TFTP Server on the switch Utilities Menu and press **Enter**.

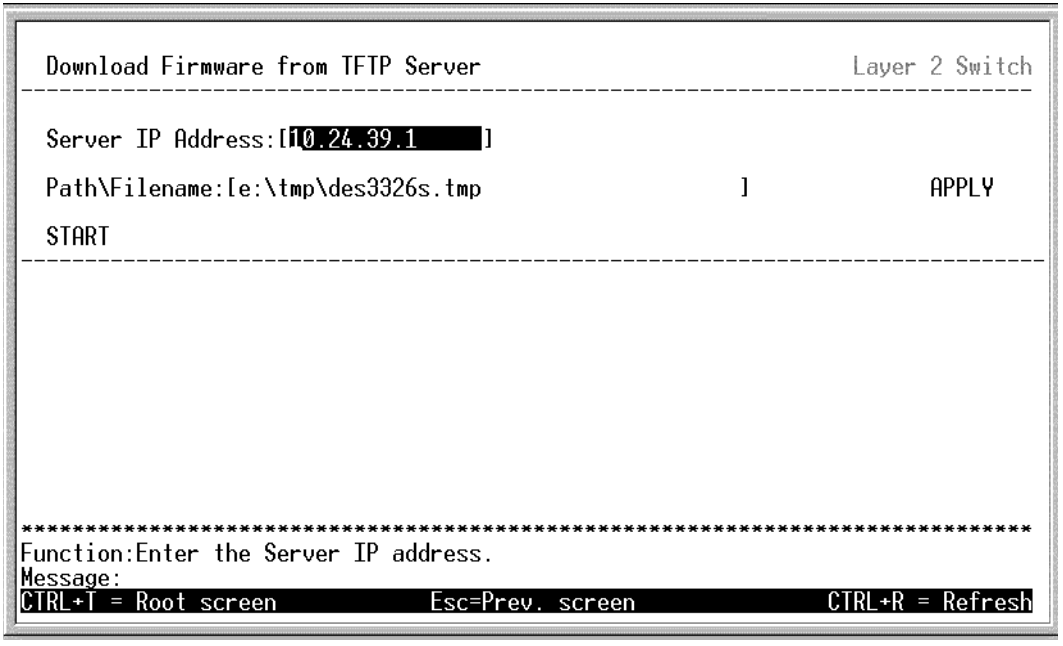

Figure 6-17: Download Configuration from TFTP Server

2 Enter the IP address of the TFTP server in the Server IP address:[ ] field

3 Specify the location of the switch configuration file on the TFTP server in the Path\Filename:[ ] field.

4 Highlight APPLY and press **Enter** to record the IP address of the TFTP server.

5 Highlight START and press enter to initiate the file transfer.

Save Configuration File to a TFTP Server

If you wish to save the current switch configuration settings, you have two options. You can save the settings to a defined section as described in Configure Section Settings, or you can upload the settings to a file on a TFTP server.

Note: uploading the current configuration settings will upload the settings that are now defined on the switch even if they have not been saved to NV-RAM.

To upload the current switch settings:

Highlight Save Settings to TFTP Server on the switch Utilities Menu and press **Enter**.

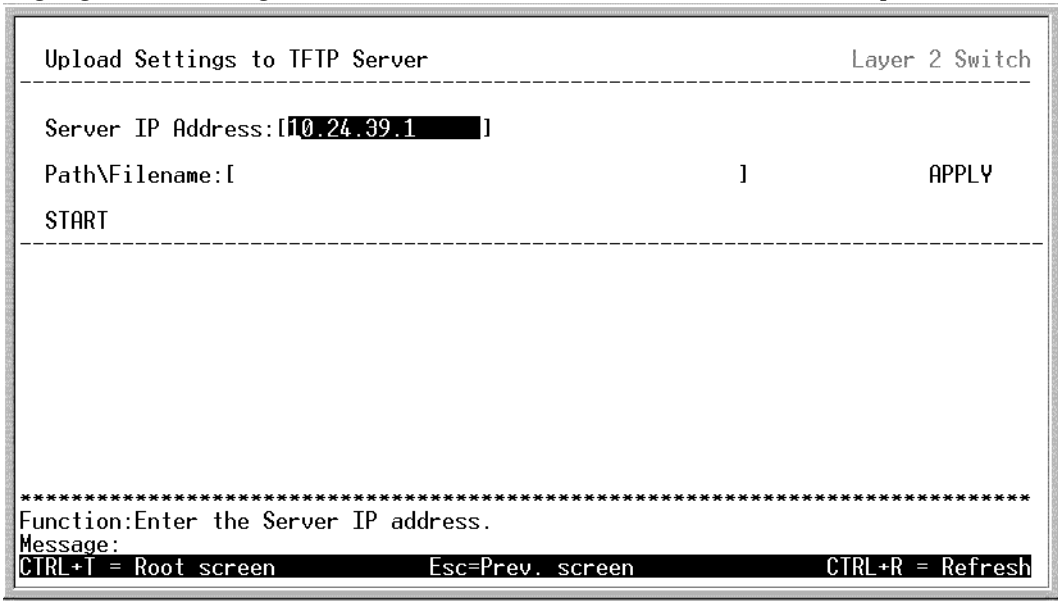

Figure 6-18: Upload Settings to TFTP Server Menu

6 Enter the IP address of the TFTP server in the Server IP address:[ ] field.

7 Specify the location for the switch configuration file on the TFTP server in the Path\Filename:[ ] field.

8 Highlight APPLY and press **Enter** to record the IP address of the TFTP server.

9 Highlight START and press enter to initiate the file transfer.

Save switch History to TFTP Server

If you wish to save the a switch history log to a TFTP server:

1 Highlight Save History Log to TFTP Server on the switch Utilities Menu and press **Enter**.

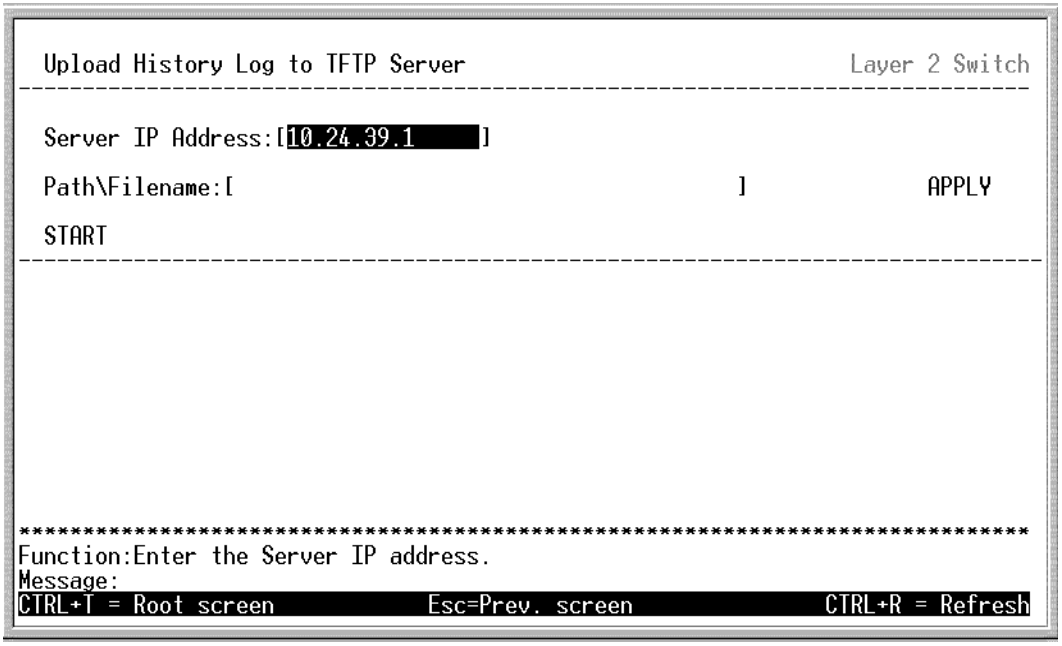

Figure 6-19: Upload History Log to TFTP Server

2 Enter the IP address of the TFTP server in the Server IP Address:[ ] field.

3 Specify the location for the switch configuration file on the TFTP server in the Path\Filename:[ ] field.

4 Highlight APPLY and press **Enter** to record the IP address of the TFTP server.

5 Highlight START and press enter to initiate the file transfer.

Ping Test

The switch can be used to test conduct Ping tests. To perform a Ping test:

1 Highlight Ping Test on the switch Utilities Menu and press **Enter.**

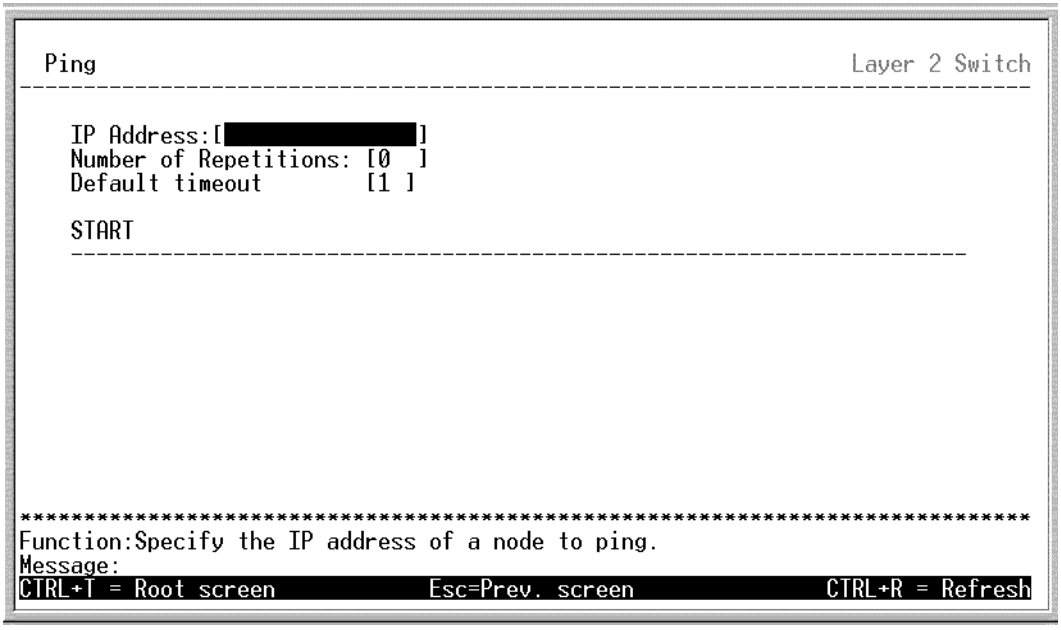

Figure 6-20: Ping Test

2 In the Ping test screen, enter the IP address of the network device to be pinged in the IP Address: [ ] field.

3 Enter the number of test packets to be sent (3 is usually enough).

4 Highlight START and press enter to initiate the ping program.

Network Monitoring

The Network Monitoring menu offers eight items including:

- Link Utilization Averages
- Port Error Packets Statistics
- Port Packet Analysis
- Browse MAC Address
- ARP Table
- Routing Table
- Browse Router Port
- IGMP Snooping
- Switch History

Choose Network Monitoring from the Main Menu. The following to see the following screen:

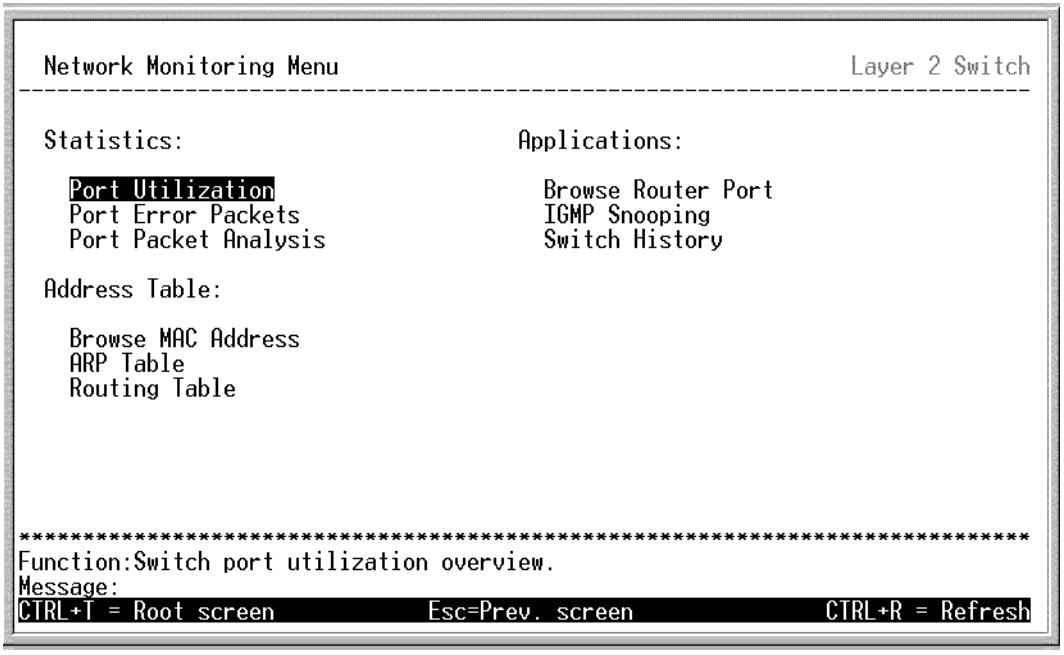

Figure 6-21: Network Monitoring Menu

Link Utilization Averages

To view the link utilization averages, highlight Link Utilization Averages on the Network Monitoring menu and press **Enter**.

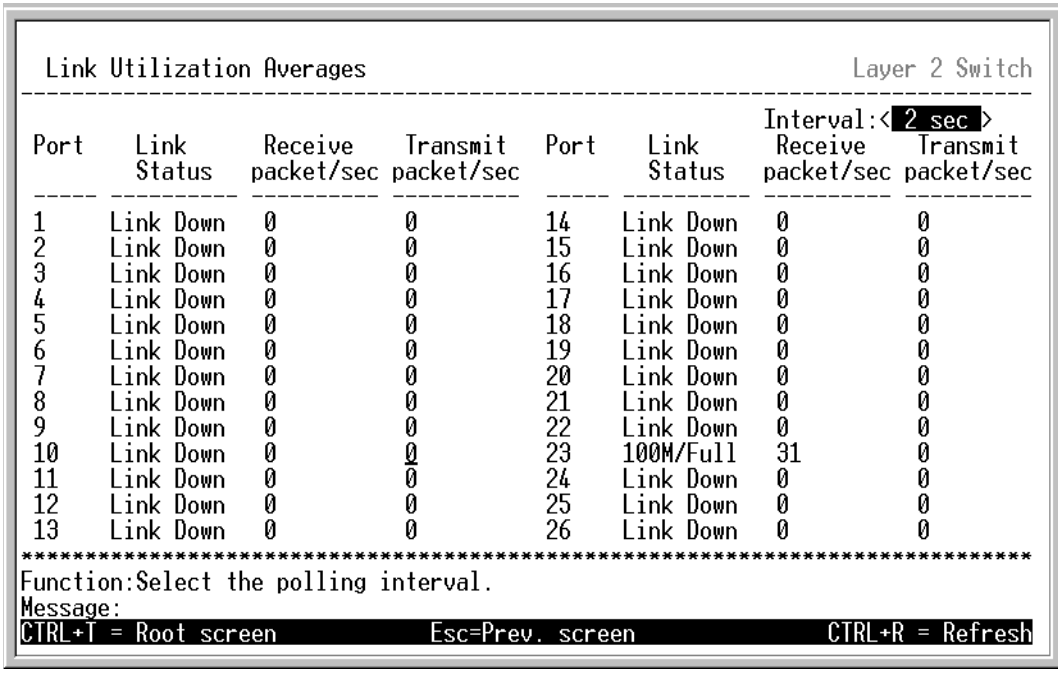

Figure 6-22: Link Utilization Averages

The Link Utilization Averages screen displays the number of packets transmitted and received per second and calculates the percentage of the total available bandwidth being used on the port (displayed under %Util.).

The Interval at which the displayed statistics are refreshed may be changed by highlighting Interval: < > and toggling to select <2 sec>, <5 sec>, <15 sec>, <30 sec>, <1 min> and <Suspend>.

Port Error Statistics

To view the error statistics for a port, highlight Port Error Packets on the Network Monitoring menu and press **Enter**.

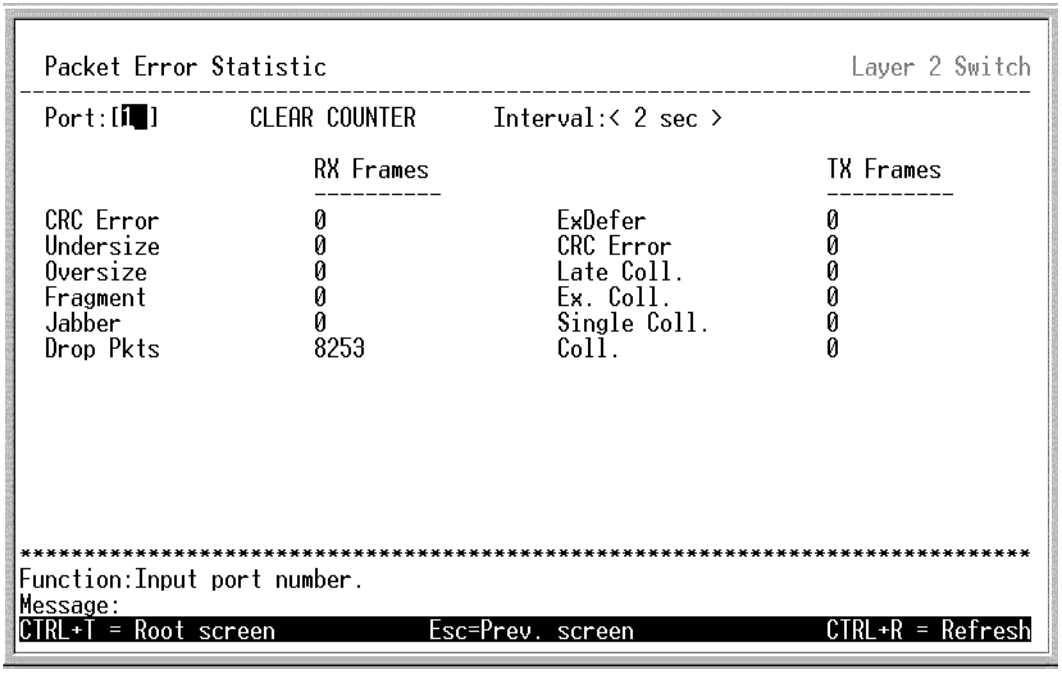

Figure 6-23: Port Error Statistics

Use the Module:<Base Unit> field to toggle either Base Unit or Slot-1 to select which group of ports will be displayed. The Gigabit Ethernet ports are in Slot-1.

Enter the port number of the port to be viewed in the Port: [ ] field. The Interval:< > field can be toggled from 2 seconds to 1 minute, or suspend. This sets the interval at which the error statistics are updated.

The following are definitions of the terms used in the Port Error Statistics menu:

Interval:<2 sec> The interval (in seconds) that the table is updated. The default is 2 seconds.

RX Frames The number of packets recieved.

CRC Error	For 10 Mbps ports, the counter records CRC errors (FCS or alignment errors), for 100 Mbps ports, the counter records the sum of CRC errors and code errors (frames received with rxerror signal).
Undersize	The total number of frames received that were less than 64 octets long (excluding framing bits, but including FCS octets) and were otherwise well formed.
Oversize	The total number of frames received that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
Fragments	The total number of frames recieved that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS or an alignment error.
Jabber	The total number of frames received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS or an alignment error.
Drop Pkts	The total number of frames for which the first transmission attempt on a particular interface was delayed because the medium was busy.
CRC Error	For 10 Mbps ports, the counter records CRC errors (FCS or alignment errors). For 100 Mbps ports, the counter records the sum of CRC errors and code errors (frames received with rxerror signal).
Late Coll.	Late Collisions — the number of times that a collision is detected later than 512 bit-times into the transmission of a packet.
Ex. Coll.	Excessive Collisions — the number of frames for which transmission failed due to excessive collisions.
Single Coll.	Single Collision Frames — the number of successfully transmitted frames for which transmission is inhibited by more than one collision.
Coll.	An estimate of the total number of collisions on this network segment.

Port Packet Analysis

The Port Packet Analysis table displays the size of packets received or transmitted by a given switch port. In addition, statistics on the number and rate of unicast, multicast, and broadcast packets received by the switch are displayed.

To view the Port Packet Analysis Screen, highlight Port Packet Analysis on the Network Monitoring Menu.

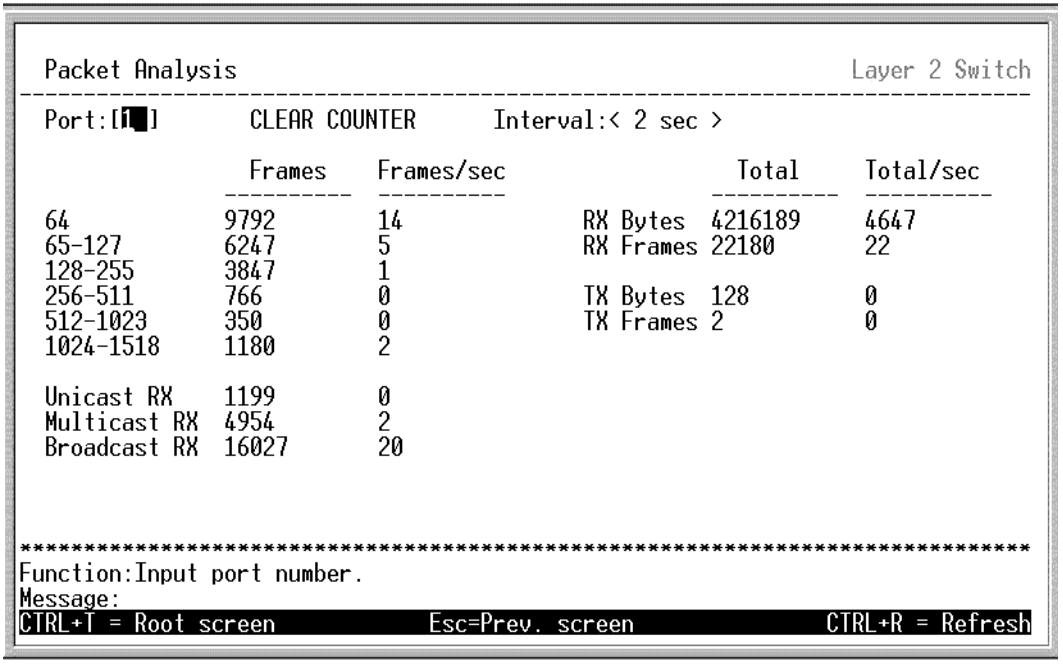

Figure 6-24: Port Packet Analysis

In addition to the size of packets received or transmitted by the selected port, statistics on the number of unicast, multicast, and broadcast packets are displayed. Toggle to select the Base Unit or Slot-1 in the Module field. Choose the Port number for which you wish to view statistics and Interval to refresh the screen.

The following are definitions of the terms used in the Port Packet Analysis menu:

Interval:<2 sec> The interval (in seconds) that the table is updated. The default is 2 seconds.

Frames	Displays the number of frames received or transmitted by the switch with the size, in octets, given by the column on the right.
Frames/sec	The number of frames transmitted or received, per second, by the switch.
Unicast RX	The number of unicast packets received by the switch in total number (under Frames) and by rate (under Frames / sec).
Multicast RX	The number of multicast packets received by the switch in total number and by rate.
Broadcast RX	The number of broadcast frames received by the switch in total number and by rate.
RX Bytes	The number of bytes (octets) received by the switch in total number and by rate.
RX Frames	The number of frames received by the switch in total number and by rate.
TX Bytes	The number of bytes (octets) transmitted by the switch in total number and by rate.
TX Frames	The number of frames transmitted by the switch in total number and by rate.

MAC Address Forwarding Table

To view the MAC address forwarding table, highlight Browse MAC Address on the Network Monitoring menu and press **Enter.**

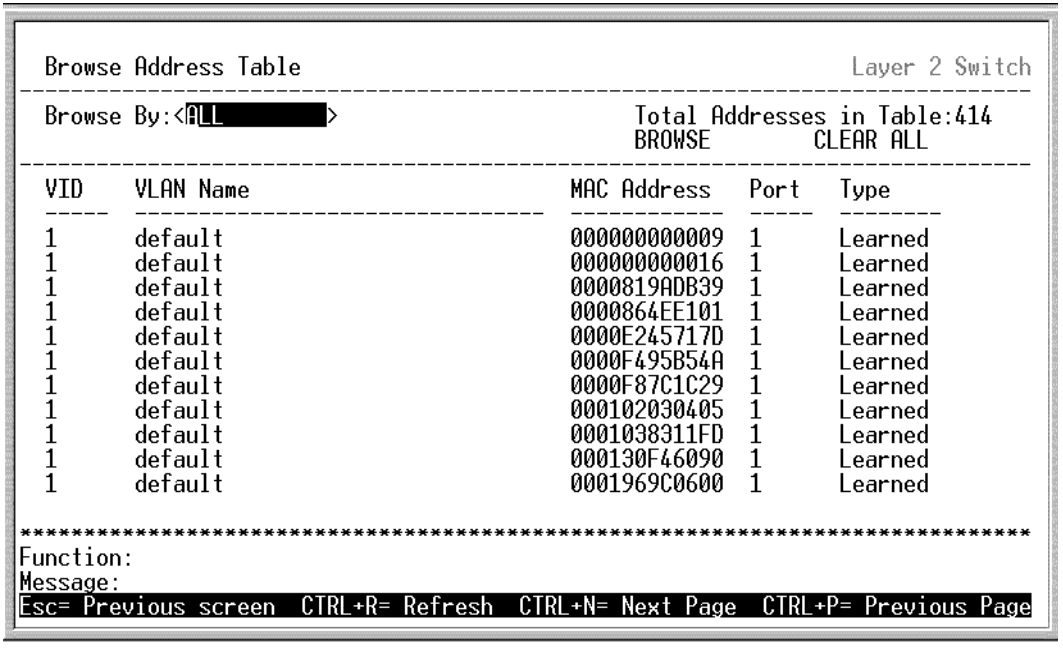

Figure 6-25: Browse MAC Address

The Browse By:<ALL> field can be toggled between ALL, MAC Address, Port, and VLAN. This sets a filter to determine which MAC addresses from the forwarding table are displayed. ALL specifies no filter.

To search for a particular MAC address:

Toggle the Browse By:<ALL> field to MAC Address. A MAC Address:[000000000000] field will appear. Enter the MAC address in the field and press **Enter**.

Browse the ARP Table

To view the switch's Address Resolution Protocol (ARP) table, highlight ARP Table on the Network Monitoring menu and press **Enter.**

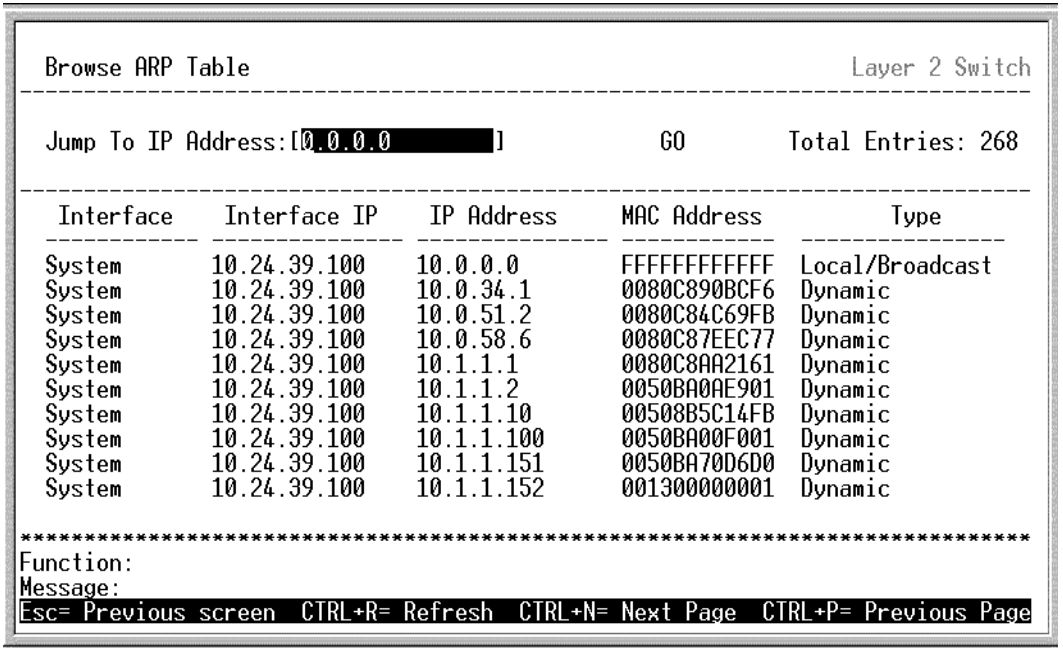

Figure 6-26: Browse ARP Table

To display a particular IP address, enter the IP address in the Jump to IP Address:[0.0.0.0] field, highlight GO and press **Enter.**

Browse the Routing Table

To view the contents of the switch's routing table, highlight Routing Table on the Network Monitoring menu and press **Enter.**

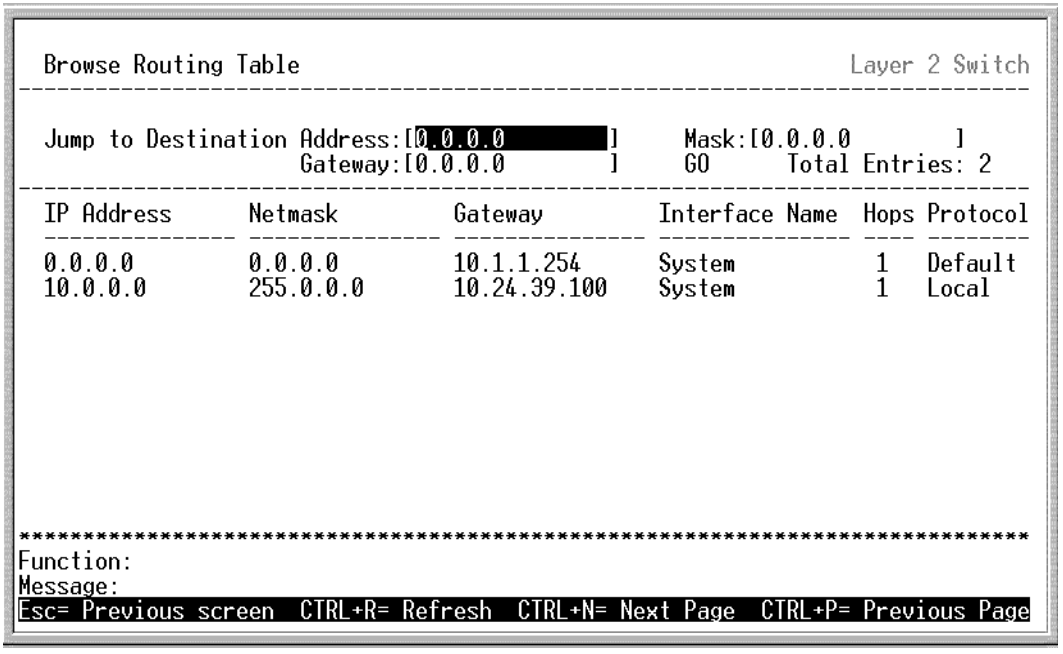

Figure 6-27: Browse Routing Table

To display a particular destination IP address, enter the IP address in the Jump to Destination Address:[0.0.0.0], the corresponding subnet mask in the Mask:[0.0.0.0] field, and the gateway address in the Gateway:[0.0.0.0] field, highlight GO and press **Enter.**

Browse Router Port

This displays which of the switch's ports are currently configured as router ports. A router port configured by a user (using the management console) is designated as a static router port, displayed using an "S". A router port that has been dynamically configured by the switch is displayed using a "D".

To view the router port table, highlight Browse Router Port from the Network Monitoring menu and press **Enter.**

A router port is simply a port that has a multicast router connected to it. Generally, this router would have a connection to a WAN or to the Internet.

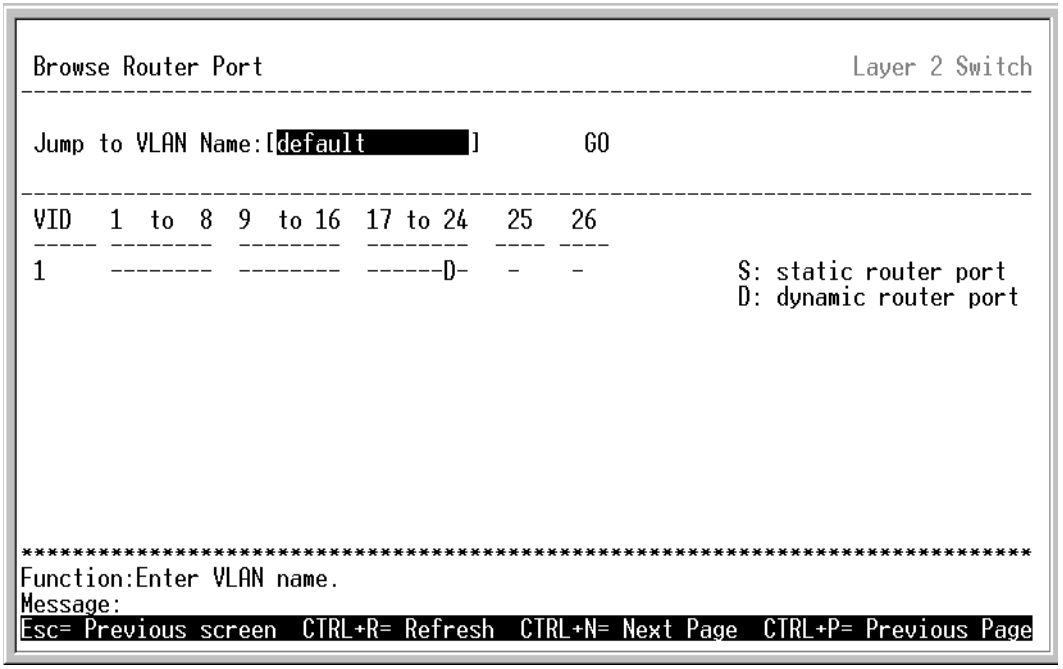

Figure 6-28: Browse Router Port

The Jump to VLAN Name:[default] field allows the entry of a VLAN name for any VLAN configured on the switch. Enter the name of the VLAN, highlight GO and press **Enter.**

All of the ports that have been designated as router ports, for the specified VLAN, will be indicated with an “S” — for staticly configured router ports, or a “D” — for router ports that are dynamically configured by the switch.

IGMP Snooping Status

This allows the switch’s IGMP snooping table to be viewed. IGMP snooping allows the switch to read the multicast group address from IGMP packets that pass through the switch. The ports where the IGMP packets were snooped are displayed using an “M”. the number of IGMP reports that were snooped are also displayed in the Reports: field.

To view the IGMP snooping table, highlight IGMP Snooping Status from the Network Monitoring menu and press **Enter.**

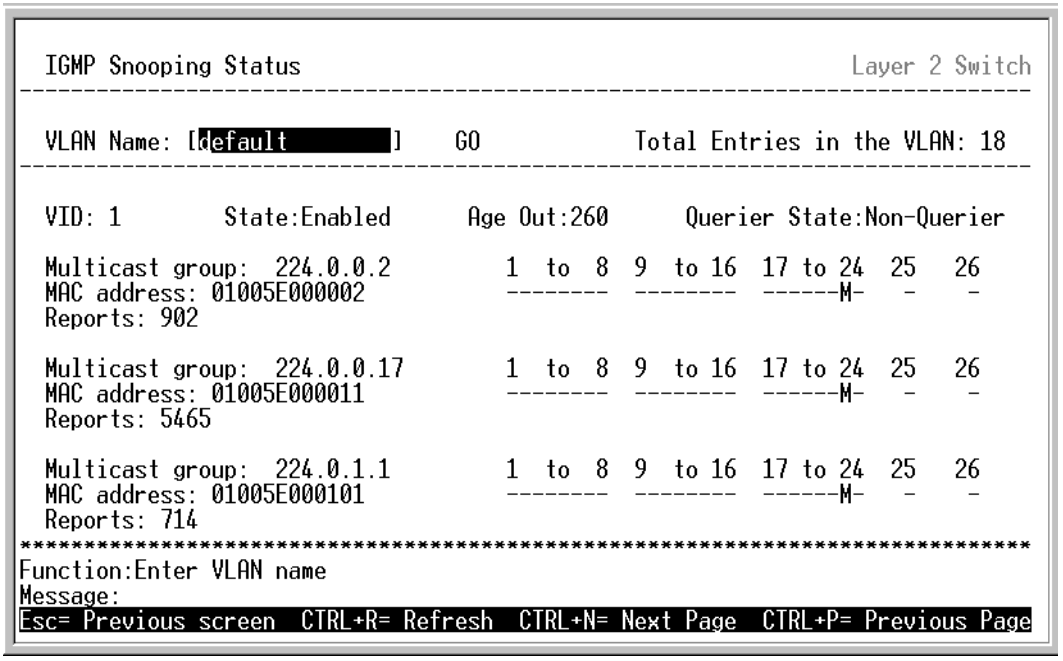

Figure 6-29: IGMP Snooping Status

The IGMP settings for the switch are also displayed.

Switch History Log

This allows the switch history log to be viewed. The switch records all traps, in sequence, that identify events on the switch since the last cold boot of the switch. A cold boot (power down and restart) erases the switch's history log.

To view the switch history log, highlight switch History from the Network Monitoring menu and press **Enter**.

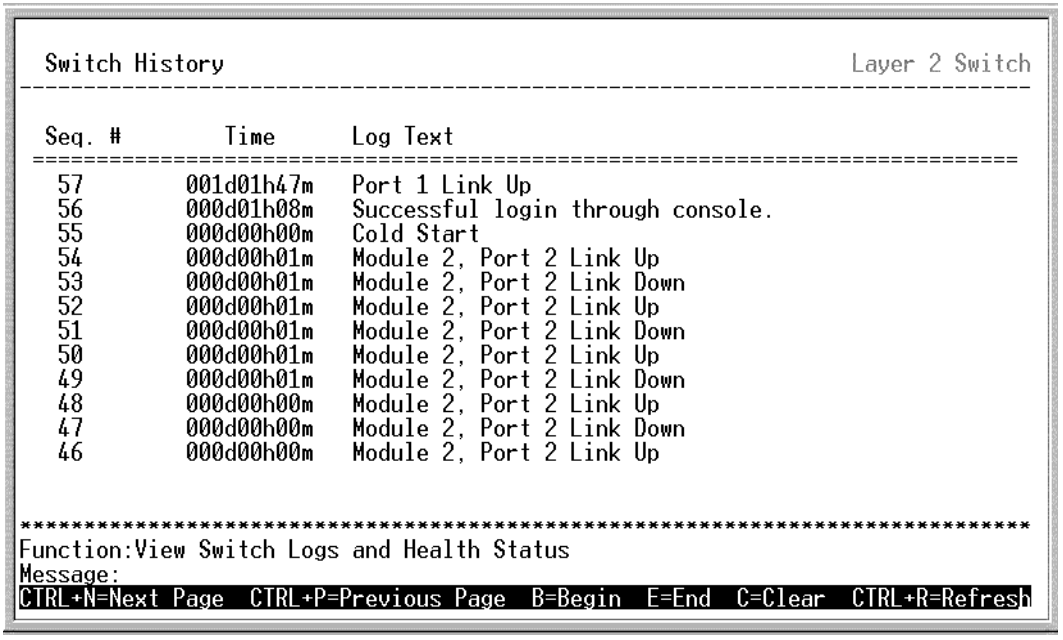

Figure 6-30: Switch History

Reboot

The Summit24e2 switch offers several options to reboot the switch. The switch can be rebooted without entering the current configuration into NV-RAM, in which case all settings not saved with the Save Changes command will be lost. Or the switch can enter the current configuration into NV-RAM and then reboot.

In addition, there are two options to return the switch's configuration to the default values entered at the factory. The full factory default values (including the default IP address and subnet mask) can be loaded into NV-RAM, or the factory default values except the user-defined IP address can be loaded.

The factory defaults are listed in Appendix B of this manual.

The System Reboot menu is shown in Figure 6-31 below.

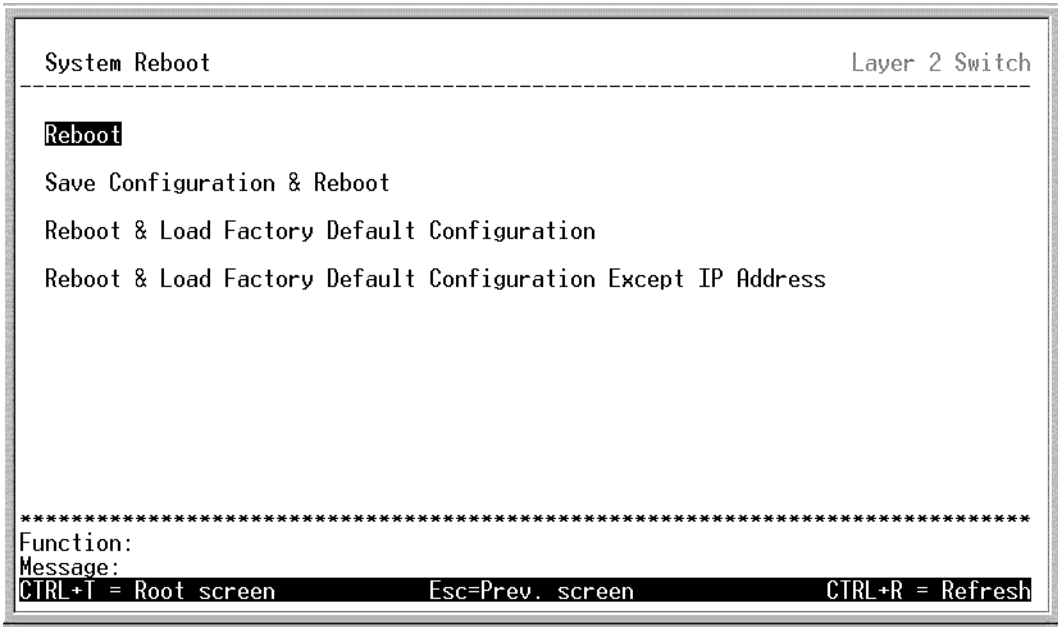

Figure 6-31: System Reboot

Advanced Setup

The Advanced Setup features include:

- Spanning Tree
  - Port Settings
- Forwarding
  - Unicast MAC Address Setting
  - Static Multicast Settings
  - Static/Default Routes
  - Static ARP
- Priority
  - Priority Queue

  - Setup MAC Address Priority
- Mirroring
  - Target Port Selection
  - Port Mirroring Settings
- IGMP Configuration
  - switch IGMP Snooping
  - IGMP Snooping Settings
  - Static Router Port Settings
- VLANs
  - Edit VLANs
- Link Aggregation
- Exception Handling

Configuring the Spanning Tree Protocol

The Spanning Tree protocol is used to prevent loops in a network in which alternative connections exist between switches. The Protocol Parameters allow you to change the behind the scene parameters of the Spanning Tree Protocol at the bridge level. The parameters for this section have been fully explained in “STA Operation Levels” on page 5-9 and “User-Changeable STA Parameters” on page 5-11. It is recommended that you read these sections, as well as “Spanning Tree Algorithm” on page 5-8 before changing any of the parameters.

The factory default settings should cover the majority of installations. However, it is advisable to keep the default settings as set at the factory; unless, it is absolutely necessary.

STP Parameter Settings

To globally configure the Protocol Parameters:

Choose Spanning Tree from the Main Menu appearing under Advanced Setup. The following Configure Spanning Tree Protocol menu will be displayed:

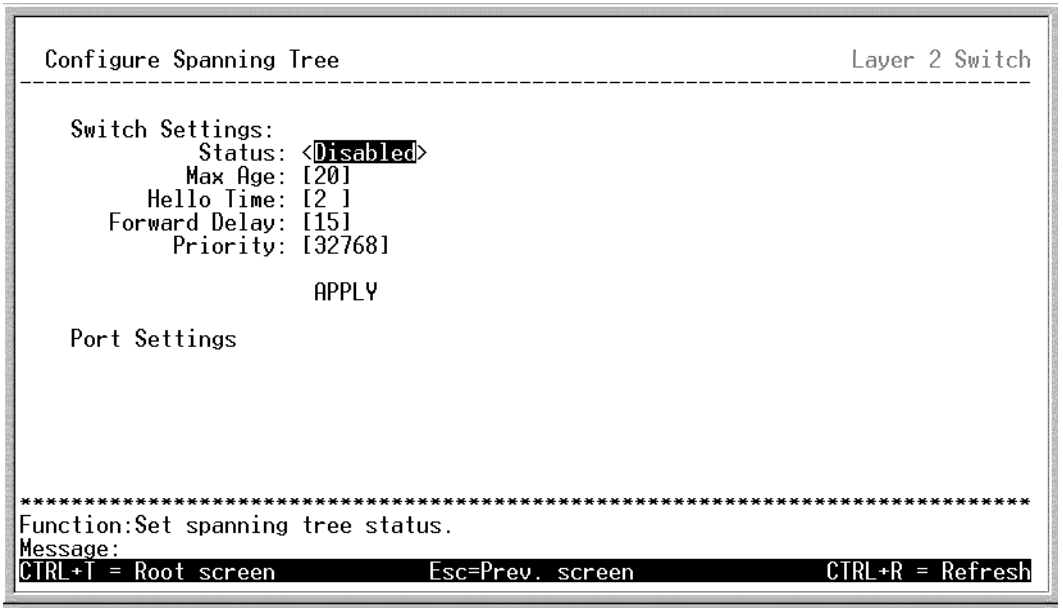

Figure 6-32: Configure Spanning Tree Menu

The user-changeable parameters in the switch are as follows:

Status Allows the spanning tree algorithm to be **Disabled** or **Enabled**, globally for the switch.

Max Age The Max. Age can be set from 6 to 40 seconds. At the end of the Max. Age, if a BPDU has still not been received from the Root Bridge, your switch will start sending its own BPDU to all other switches for permission to become the Root Bridge. If it turns out that your switch has the lowest Bridge Identifier, it will become the Root Bridge.

Hello Time Hello Time:[ ] The Hello Time can be set from 1 to 10 seconds. This is the interval between two transmissions of BPDU packets sent by the Root Bridge to tell all other switches that it is indeed the Root Bridge. If you set a Hello Time for your switch, and it is not the Root Bridge, the set Hello Time will be used if and when your switch becomes the Root Bridge.

Forward Delay Forward Delay:[ ] The Forward Delay can be from 4 to 30 seconds. This is the time any port on the switch spends in the listening state while moving from the blocking state to the forwarding state.

Priority Priority:[ ] A Priority for the switch can be set from 0 to 65535. 0 is equal to the highest Priority. This number is used in the voting process between switches on the network to determine which switch will be the root switch. A low number indicates a high priority, and a high probability that this switch will be elected as the root switch.

Observe the following formulas when setting the Spanning Tree parameters:

Max. Age = 2 x (Forward Delay - 1 second)

Max. Age = 2 x (Hello Time + 1 second)

The Hello Time cannot be longer than the Max. Age. Otherwise, a configuration error will occur.

Port Spanning Tree Settings

To configure STP port settings, toggle the View Ports:< > field to the range of ports to be configured. The Fast Ethernet ports displayed for configuration in groups of 12 and the two Gigabit Ethernet ports are displayed together. Select the port(s) to be configured in the Configure Port from [ ] to [ ] field.

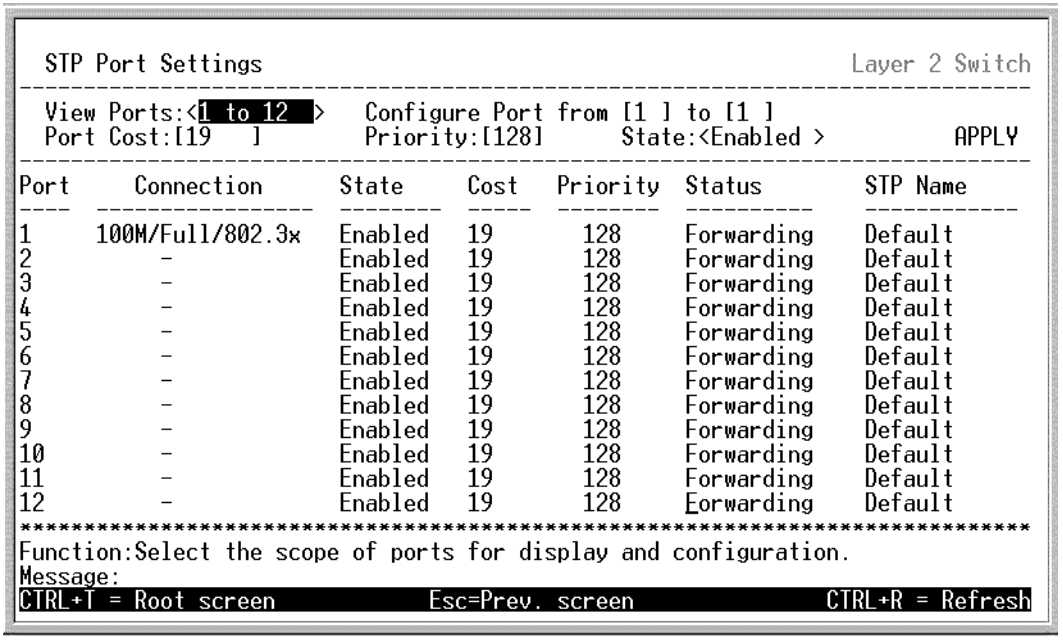

Figure 6-33: Spanning Tree Port Settings Menu

The Port Group STP parameters that can be configured are:

Port Cost — A Port Cost can be set from 1 to 65535. The lower the number, the greater the probability the port will be chosen to forward packets.

Priority — A Port Priority can be from 0 to 255. The lower the number, the greater the probability the port will be chosen as the Root Port.

State:<Disabled> — Allows the spanning tree protocol to be **Disabled** or **Enabled** for an individual port, or a group of ports, as selected above.

Forwarding

When an incoming packet is processed in the switch, the forwarding table is consulted in order to decide how to handle the packet. Specifically, the switch must decide to

either filter the packet off the network or to forward it through the port and VLAN (if enabled) on which its destination lies.

Dynamic MAC forwarding and MAC filtering are functions of the Learning Process, or the process of observing network traffic. The Learning Process is an automatic and continuous function of the switch, the only variable that can be configured for this process is the MAC address age out time. Static MAC forwarding uses static entries that may be added and removed from the database by the user. They are not automatically removed by any age-out mechanism.

The MAC Forwarding and IP Forwarding screens allow you to stop or start address learning for specified MAC addresses and change the way the switch treats MAC address table entries.

MAC Address Aging Time

Highlight Forwarding Menu from the Main Menu and press **Enter**.

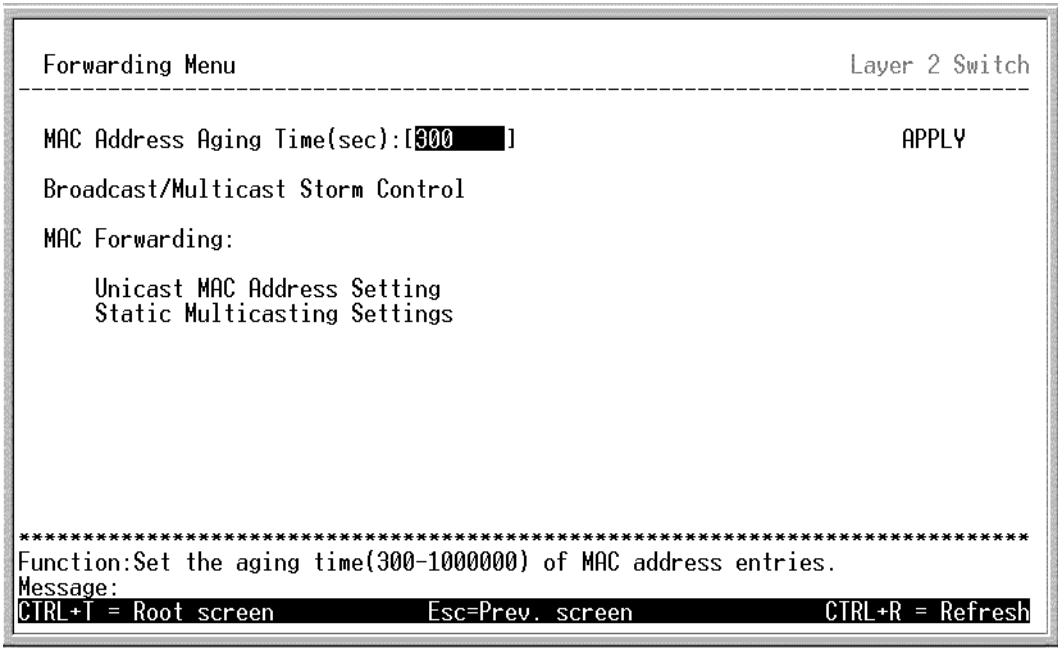

Figure 6-34: Forwarding Menu

The following fields can then be set:

MAC Address Aging Time(sec):[300] This field specifies the length of time a learned MAC Address will remain in the forwarding table without being accessed (that is, how long a learned MAC Address is allowed to remain idle). The Aging Time can be set to any value between 10 and 1,000,000 seconds. The default is **300** seconds

Note: a very long Aging Time can result with the out-of-date Dynamic Entries that may cause incorrect packet filtering/forwarding decisions. A very short aging time may cause entries to be aged out to soon, resulting in a high percentage of received packets whose source addresses cannot be found in the address table, in which case the switch will broadcast the packet to all ports, negating many of the benefits of having a switch.

Broadcast/Multicast Storm Control

Highlight Broadcast/Multicast Storm Control from the Forwarding menu and press **Enter.**

```
 Broadcast/Multicast Storm Control                              Layer 2 Switch
--------------------------------------------------------------------------------

           Upper Threshold(Kpps)   Broadcast Storm Mode   Multicast Storm Mode
           =====================   ====================   ====================
Group 1:           [128]                <Disabled>             <Disabled>
Group 2:           [128]                <Disabled>             <Disabled>
Group 3:           [128]                <Disabled>             <Disabled>
Group 4:           [128]                <Disabled>             <Disabled>
Group 5:           [128]                <Disabled>             <Disabled>
                                                                      APPLY

********************************************************************************
Function:Set the upper threshold of this storm control for ports 1-8(0-255).
Message:
CTRL+T = Root screen          Esc=Prev. screen              CTRL+R = Refresh
```

Figure 6-35: Broadcast/Multicast Storm Control Menu

Use the entry fields described below for the parameters that control how the switch will react to broadcast and multicast storms.

The switch ports are grouped in the following groups:

- Group 1 - ports 1 through 8
- Group 2 - ports 9 through 16
- Group 3 - ports 17 through 24
- Group 4 - Gigabit port 25
- Group 5 - Gigabit port 26

The following fields can be set:

Upper Threshold (Kbp):[128]	This is the number of Broadcast / Multicast in Kbs received by the switch that will trigger the switch's reaction to a Broadcast / Multicast storm.
Broadcast Storm Mode:<Disabled>	Toggle to select Enabled or Disabled using the space bar to globally enable or disable the switch's reaction to Broadcast storms, triggered at the threshold set above.
Multicast Storm Mode:<Disabled>	This field can be toggled between Enabled and Disabled using the space bar. This enables or disables, globally, the switch's reaction to Multicast storms, triggered at the threshold set above.

Highlight APPLY and press **Enter** to apply the Multicast Forwarding settings. Use the Save Changes menu to save the settings to NV-RAM.

Unicast MAC Address Forwarding

Highlight Unicast MAC Address Setting from the Forwarding Menu and press **Enter.**

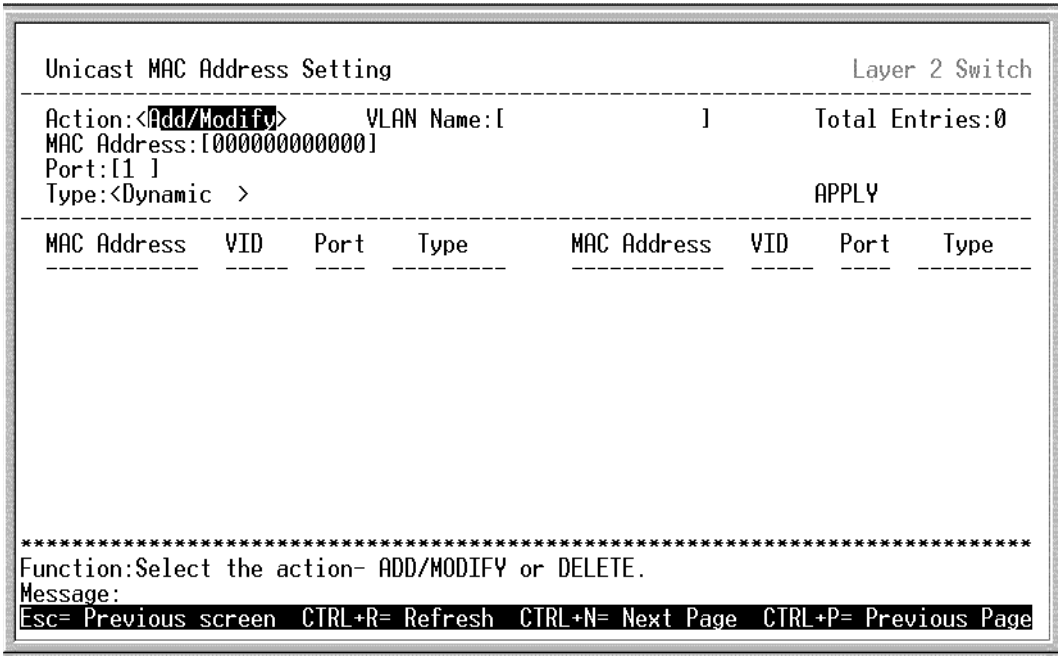

Figure 6-36: Unicast MAC Address Setting Menu

The following fields can be set:

Action: <Add/Modify>	Toggle to choose **Add/Modify** or **Delete** to add, change or delete an entry.
VLAN Name: [ ]	The name of the VLAN on which the MAC address resides.
MAC Address: [000000000000]	The MAC address to be added to the static forwarding table. Below the entry fields in the menu, a display of a total of ten destination addresses per page can be seen. The switch can hold up to 256 entries.
Port:[ ]	The port number corresponding to the MAC destination address. The switch will always forward traffic to the specified device through this port.

Type:<Static> Can be toggled between **Static,** and **BlackHole** using the space bar. Static enters the MAC address into the switch's unicast forwarding table. BlackHole causes all packets with the specified MAC address as their destingation to be dropped.

Highlight APPLY and press **Enter** to apply the Multicast Forwarding settings. Use the Save Changes menu to save the settings to NV-RAM.

Multicast MAC Address Forwarding

Highlight Static Multicasting Forwarding from the Forwarding Menu and press **Enter.**

```
 Setup Static Multicast Forwarding                                Layer 2 Switch
--------------------------------------------------------------------------------
 Action: <Add/Modify>     VLAN Name: [default                              ]
 Multicast MAC Address[000000000000]
 Port    1  to  8  9  to 16  17 to 24  25   26
 (E/-)  [--------][--------][--------] [-]  [-]
                                              Total Entries:0        APPLY
--------------------------------------------------------------------------------
 MAC Address   VID   1  to  8  9  to 16  17 to 24   25   26
 ------------  ----- --------  --------  --------  ---- ----

********************************************************************************
Function:Select the action- ADD/MODIFY or DELETE.
Message:
Esc= Previous screen  CTRL+R= Refresh  CTRL+N= Next Page  CTRL+P= Previous Page
```

Figure 6-37: Static Multicast Settings Menu

The following fields can be set:

Action: <Add/Modify> Toggle to choose **Add/Modify** or **Delete** to add, change or delete an entry.

VLAN Name: [ ]	The name of the VLAN on which the MAC address resides.
Multicast MAC Address: [000000000000]	The multicast MAC address to be added to the static forwarding table. Below the entry fields in the menu, a display of a total of ten destination addresses per page can be seen. The switch can hold up to 256 entries.
Port:(E/-):[ ][ ][ ]	A port number should be chosen for each corresponding destination address. The switch will always forward traffic to the specified device through this port. Each port can be an Egress, Forbidden, or a Non-member of the multicast group, on a per-VLAN basis. To set a port's multicast group membership status, highlight the first field of **Port:(E/-): [ ][ ][ ]**. Each port's multicast group membership can be set individually by highlighting the port's entry using the arrow keys, and then toggling between E, F, or - using the space bar. Use the following definitions to guide you: **E** — Egress membership specifies the port as being a static member of the multicast group. Egress Member Ports are ports that will be transmitting traffic for the multicast group. **-** — Non-member status specifies the port as not being a member of the multicast group, but the port can become a member of the multicast group dynamically.
Type:<Dynamic>	Can be toggled between **Dynamic**, **Static,** and **BlackHole** using the space bar.

Highlight APPLY and press **Enter** to apply the Multicast Forwarding settings. Use the Save Changes menu to save the settings to NV-RAM.

Priority

Highlight Priority from the Main Menu and press **Enter.**

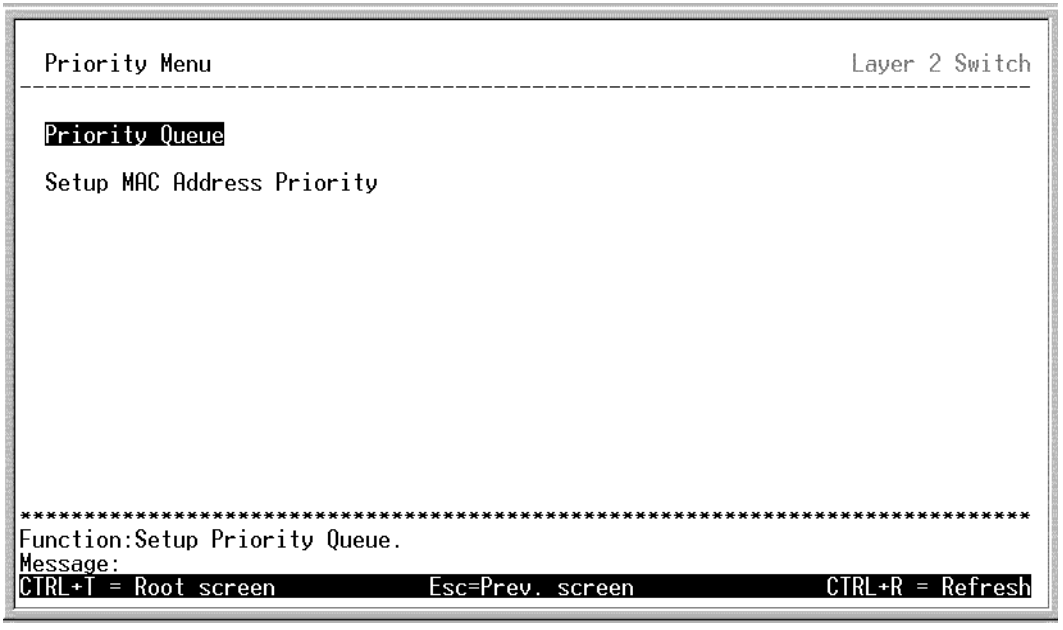

Figure 6-38: Priority Menu

Priority Queue Configuration

The Summit24e2 switch has 4 priority queues. These priority queues are numbered from Q0 — the lowest priority queue — to Q3 — the highest priority queue. The eight priority queues specified in IEEE 802.1p (Q0 to Q7) are mapped to the switch's priority queues as follows:

- Q1 and Q0 are assigned to the switch's Q0 queue.
- Q3 and Q2 are assigned to the switch's Q1 queue.
- Q5 and Q4 are assigned to the switch's Q2 queue.
- Q7 and Q6 are assigned to the switch's Q3 queue.

The switch's four priority queues are emptied in a round-robin fashion — begining with the highest priority queue, and proceeding to the lowest priority queue before returning to the highest priority queue.

To configure the switch's priority queing, highlight Priority Queue from the Priority Menu and press **Enter.**

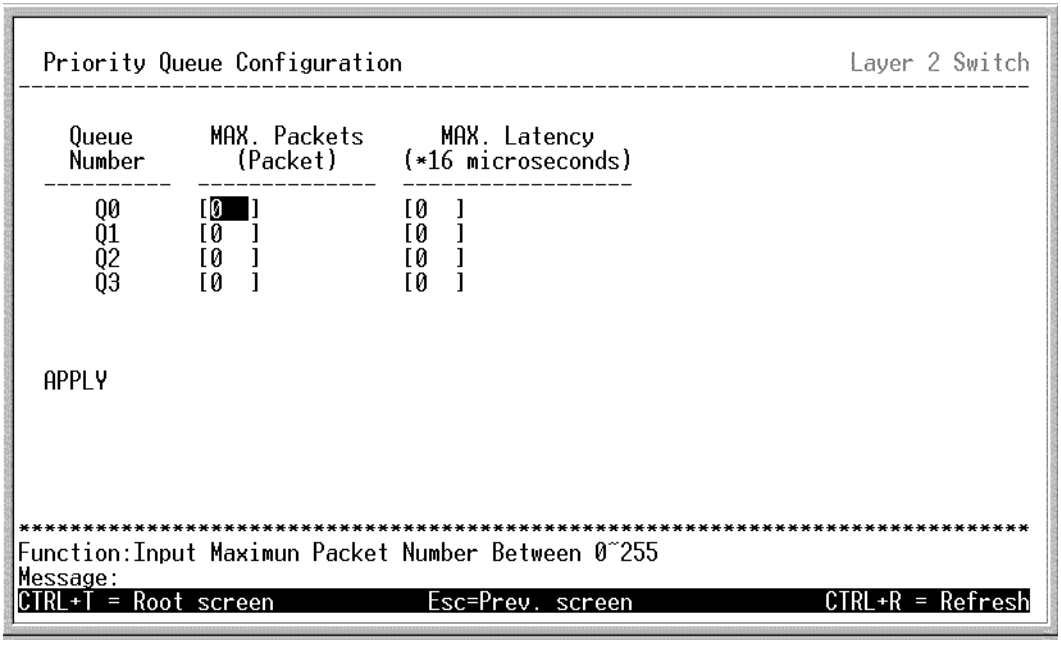

Figure 6-39: Priority Queue Configuration Menu

The priority queues are numbered in the order of the lowest priority queue — Q0 — to the highest priority queue — Q3.

The MAX. Packets field specifies the number of packets that a que will transmit before surrendering the transmit buffer to the next lower priority queue in a round-robin fashion.

The MAX. Latency field specifies the maximum amount of time — in multiples of 16 microseconds — that a queue will have to wait before being given access to the transmit buffer. The MAX. Latency is a priority que timer. When it expires, it overides the round-robin and gives the priority que that it was set for access to the transmit buffer.

There is a small amount of additional latency introduced because the priority que that is transmitting at the time the MAX. Latency time expires will finish transmitting its current packet before giving up the transmit buffer.

The following fields can be set:

Queue Number	Allows the selection of any one of the four priority queues available on the switch. Q0 is the lowest priority que. Q3 is the highest priority que.
MAX. Packets (Packet)	Allows the specification of between 0 and 255 packets that will be transmitted by the specified que before the next lower priority que is allowed to transmit packets.
MAX. Latency (+16 microseconds)	Allows the specification of the maximum latency of between 0 and 255 multiples of 16 microseconds. This establishes a priority que timer that overides the round-robin and allows the priority que (for which the timer was set) the next access to the transmit buffer.

If all of the MAX. Packets and MAX Latency fields — for all four priority queues — are set to zero, then the priority queues are emptied in a strict priority round-robin. Q3 (the highest priority que) will transmit all of the packets it has accumulated before Q2 is allowed to transmit. Q2 will then transmit all of the packets it has accumulated before Q1 is allowed to transmit, and so on. Finally, when Q0 has transmitted all of the packets it has accumulated, Q3 will again be allowed to transmit packets.

A strict priority round-robin priority scheme is recommended for most installations that do not have specific priority or latency requirments.

The MAX. Packets fields can be configured to provide control over the maximum number of packets that can be transmitted before a given priority que will be allowed to transmit its accumulated packets.

The resulting priority scheme is best illustrated using an example:

If the MAX. Packets fields for Q0, Q1, Q2, and Q3 are set to 2, 4, 8, and 12 respectively, then the priority queues are emptied as follows:

(assuming all queues have many packets to transmit)

1 The first 12 packets from Q3 will be sent.

2 8 packets from Q2 will be sent.

3 4 packets from Q1 will be sent.

4 2 packets from Q0 will be sent.

5 12 more packets from Q3 will be sent, ... and so on.

In this example, Q2 will transmit 8 packets and then have a maximum of 18 packets transmitted by the other queues before it can send 8 more packets.

The MAX. Packets fields should be either all zeros (strict priority) or non-zero for all priority queues.

If the MAX. Latency field for Q2 (in the above example) is set to 1, then a 16 microsecond timer is set for Q2. When Q2 finishes sending its alloted 8 packets, the timer starts. If 16 microseconds elapse before Q2 is again allowed to send 8 packets, the currently transmitting que is allowed to completly finish transmitting its current packet, and then the transmit buffer is given to Q2. Q2 can then transmit its next 8 packets. When Q2 is finished transmitting its alloted 8 packets, Q1 is then allowed to transmit its alloted 4 packets, and so on.

The MAX. Latency should be set to a non-zero value for only one priority que. All other priority queues should have this field set to zero.

If a MAX. Latency timer is established for more than one priority que, it is possible that the timer will expire and overide the round-robin repeatadly such that only the two (or more) queues for which a MAX. Latency timer has been defined are allowed to transmit packets — leaving the other queues unable to send their full allotment of packets or even unable to send any packets at all.

MAC Address Priority

The Summit24e2 switch allows packets sent to (or received from) specific MAC addresses to be assigned to one of the four priority queues configured above.

Highlight MAC Address Priority from the Priority Menu and press **Enter.**

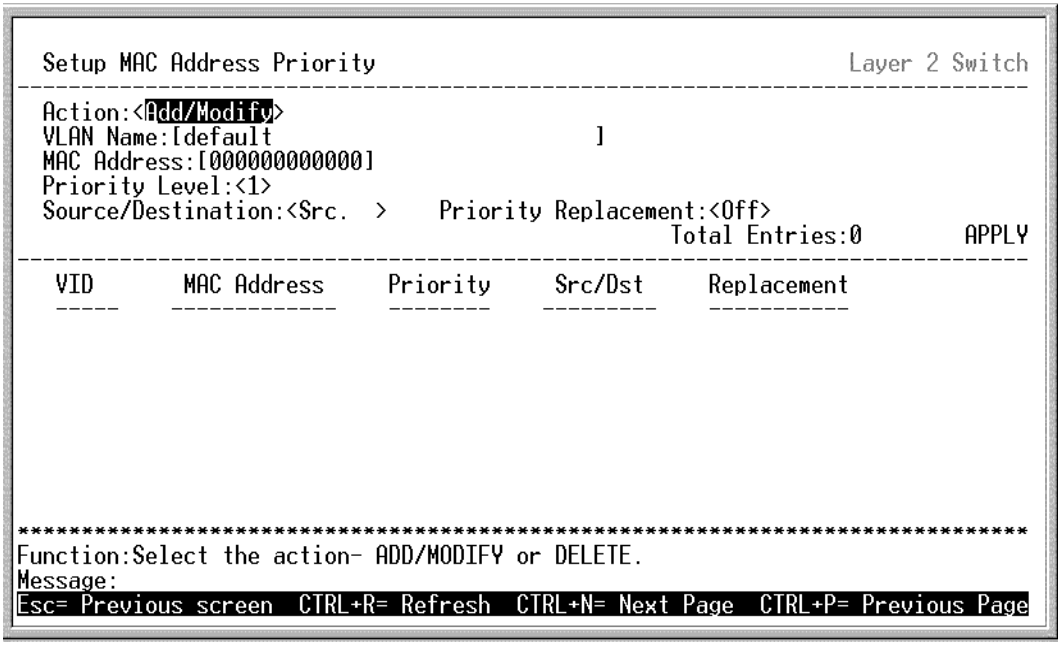

Figure 6-40: Setup MAC Address Priority Menu

Action: <Add/Modify> Toggle to choose Add / Modify or Delete to add, modify or delete an entry.

VLAN Name:[default] The name of the VLAN the MAC address entered below belongs to.

MAC Address: [000000000000] The MAC address for which priority is to be set.

Priority Level:<1> Allows the selection of which of the four priority queues packets with the above MAC address as their destination address will be put in.

Source/Destination: <Src.> Can be toggled between Scr. (Source), Dst**.** (Destination), and Either. This specifies the above MAC address as a source of packets to forward (packets coming into the switch from this MAC address), a destination (packets flowing out of the switch to this destination, or either (both a source and a destination.

Priority Replacement: <Off> Can be toggled between On and Off using the space bar. This specifies whether the switch should replace a packet's existing priority tag with the priority specified above. The priority specified above is used by the switch internally, but the packet's priority tag will remain unchanged (perhaps for use by a subsequent network device) if Priority Replacement is set to **Off**. The packet's priority tag will be changed to the priority set above, if Priority Replacement is set to **On**.

Port Mirroring

The switch allows you to copy frames transmitted and received on a port and redirect the copies to another port. You can attach a monitoring device to the mirrored port, such as a sniffer or an RMON probe to view details about the packets passing through the first port. This is useful for network monitoring and troubleshooting purposes.

Choose Mirroring Configuration on the Main Menu under Advanced Setup to access the following screen:

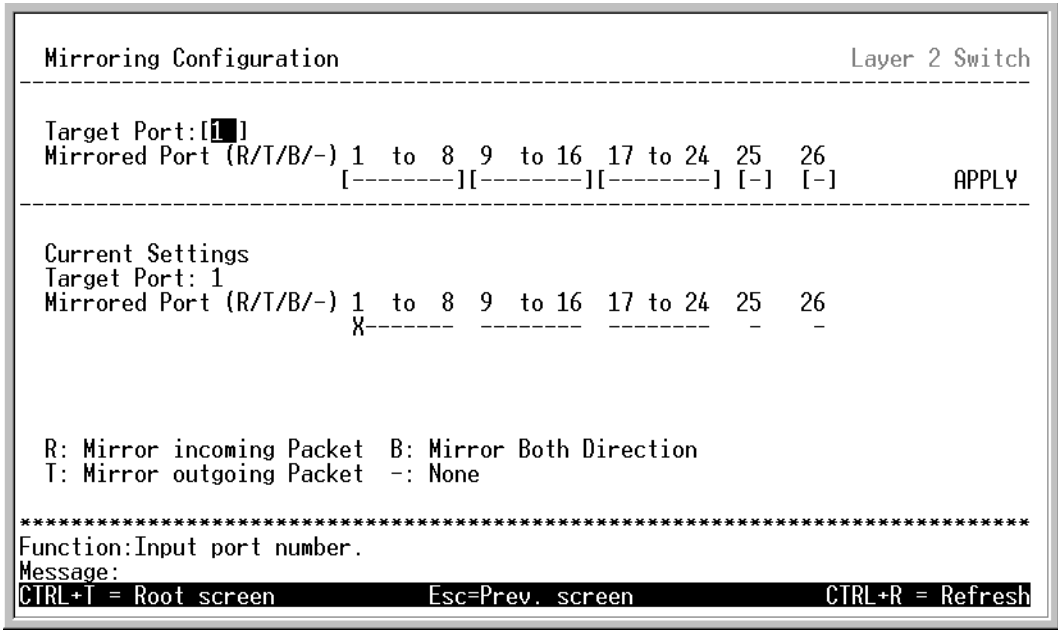

Figure 6-41: Port Mirroring Menu

The target port is the port chosen to receive duplicated packets used for analysis. This port is where a monitoring/troubleshooting device such as a sniffer or an RMON probe is connected for such analysis. Only one target port can be chosen and the target port is displayed with an "X" corresponding to the target port's number. To select the mirrored port(s) and the target port for mirroring:

1. Enter the target port number in the Target Port [ ] field.
2. Use the arrow keys to highlight the port(s) to be mirrored in the fields below the Mirrored Port (R/T/B/-) field.
3. Toggle the field corresponding the the port number of the port to be mirrored. The field can be toggled between R — mirror packets leaving the port, T — mirror packets coming into the port, B — mirror packets either leaving or coming into the port (both directions), or - — designating the port as one not being mirrored.
4. Highlight APPLY and press **Enter**.
5. The Screen will now list the port that has just been selected as the mirror target port. After selecting the target port, press **Esc** to return to the previous screen (Mirroring Menu).

Note: you should not mirror a fast port onto a slower port. For example, if you try to mirror the traffic from a 100 Mbps port onto a 10 Mbps port, this can cause throughput problems. The port you are copying frames to should always support an equal or higher speed than the port from which you are sending the copies to. Also, the target port cannot be a member of a trunk group.

IGMP Configuration

Highlight IGMP Configuration from the Main Menu and press **Enter.**

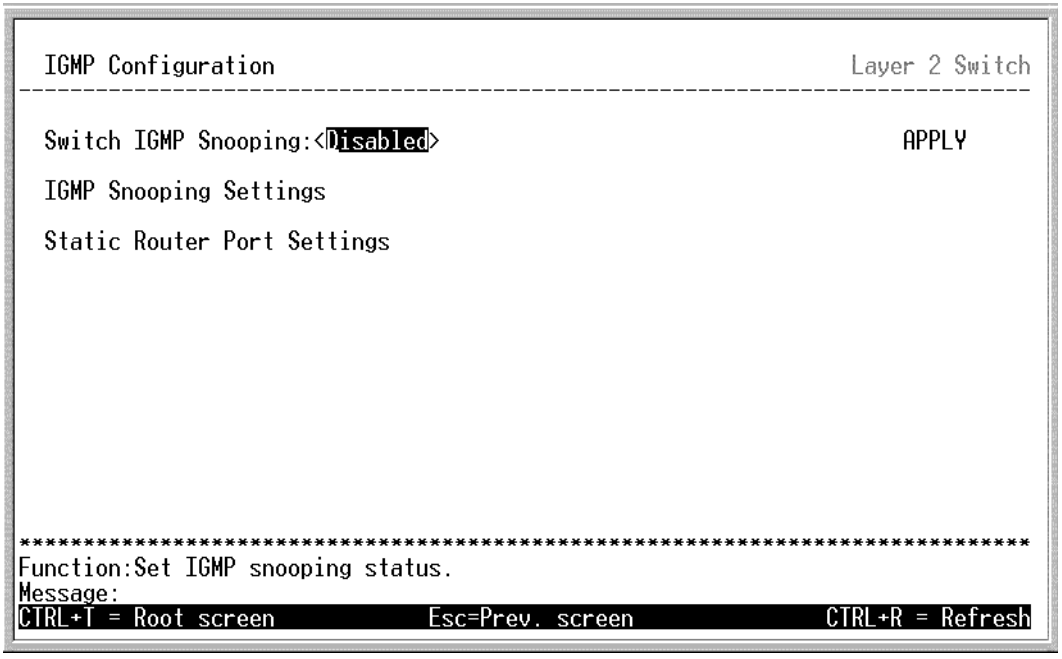

Figure 6-42: IGMP Configuration Menu

The following fields can be set:

Switch IGMP Snooping: <Disabled>	Can be toggled between Disabled and Enabled using the space bar. This allows IGMP to be disabled or Enabled, globally, on the switch.

IGMP Snooping Settings

Highlight IGMP Snooping Settings from the IGMP Configuration Menu and press **Enter.**

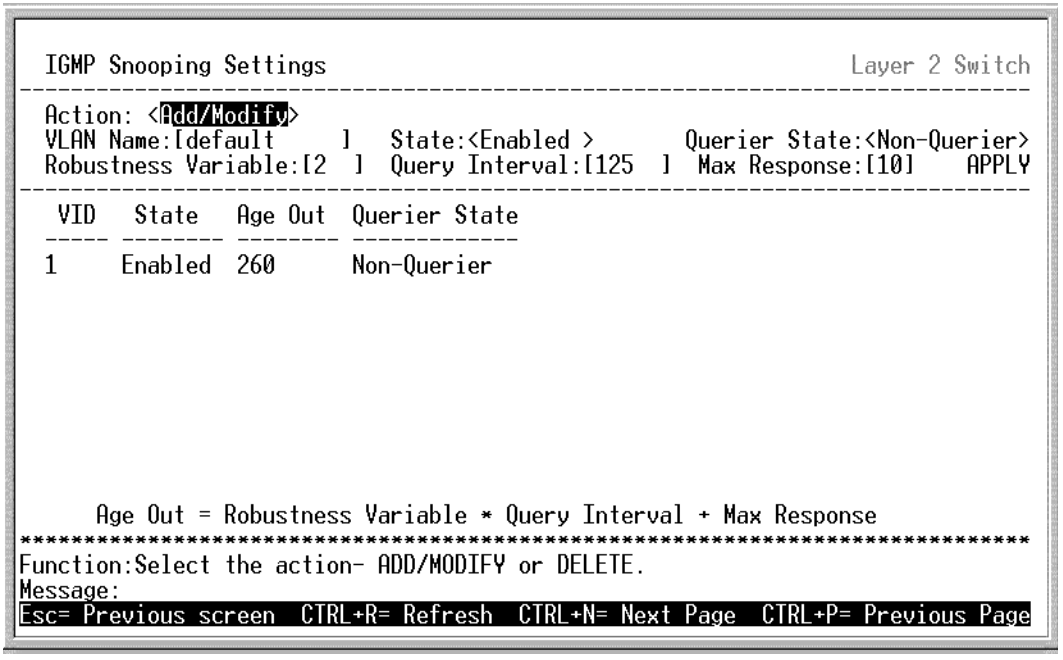

Figure 6-43: IGMP Snooping Settings Menu

Use the following fields to edit the settings for IGMP Snooping:

Action: <Add/Modify>	Toggle either Add/Modify or Delete to add, change or delete an entry.
VLAN Name: [default]	For port based VLANs only. Enter the name of the VLAN number for which you wish to enable, disable or modify IGMP snooping settings
State: <Enabled>	Toggle to enable or disable IGMP snooping for the chosen VLAN.
Querier State: <Non-Querier>	Can be toggled between Non-Querier and Querier. This allows the switch to be specified as a IGMP Querier (sends IGMP query packets) or a Non-Querier (does not send IGMP query packets).

Robustness Variable: [2]	The Robustness Variable field allows an entry between 1 and 255 that defines the maximum time (in seconds) between the receipt of IGMP queries. If this timer expires without the receipt of another IGMP query, the switch assumes the querier is no longer present.
Query Interval: [125]	The Query field allows an entry between 1 and 9,999 seconds and defines the time between transmitting IGMP queries.
Max Response: [10]	The Max-Response field allows an entry between 1 and 254 and defines the maximum time allowed before sending a response report to a query measured in units of 1/10 of a second. This is used to adjust the "leave latency", the time internal between the moment the last host leaves a group and when the routing protocol is notified there are no more members.

Highlight APPLY and press **Enter** to apply the IGMP Snooping settings. Use the Save Changes menu to save the settings to NV-RAM.

Static Router Port Settings

A router port allows UDP multicast and IGMP packets to be forwarded to a designated port on the switch regardless of VLAN configuration.

A static router port is a port that has a multicast router attached to it. Generally, this router would have a connection to a WAN or to the Internet. Establishing a router port will allow multicast packets coming from the router to be propagated through the network, as well as allowing multicast messages coming from the network to be propagated to the router.

The purpose of a router port is to enable UDP multicast packets, and IGMP multicast group membership messages to reach a multicast router attached to the switch. Routers do not implement IGMP snooping or transmit/forward IGMP report packets. Thus, forwarding all IP UDP multicast packets to a static router port on the switch guarantees that all multicast routers can reach all multicast group members.

To setup a static router port, highlight Static Router Port Settings from the IGMP Configuration Menu and press **Enter**.

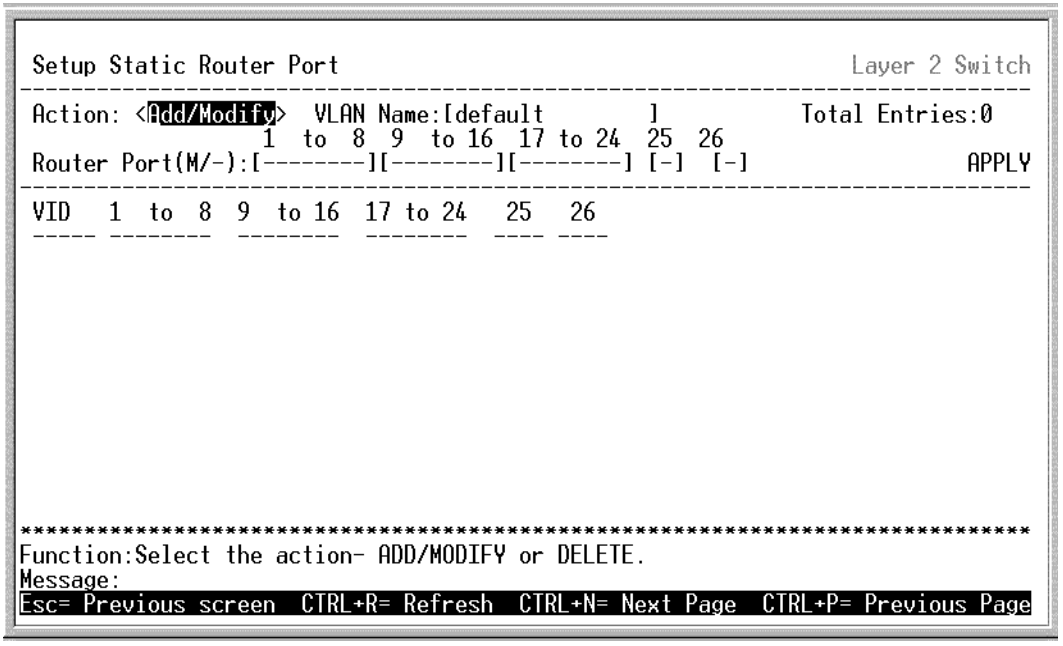

Figure 6-44: Static Router Port Settings Menu

All IGMP Report packets will be forwarded to the router port.

IGMP queries (from the router port) will be flooded to all ports.

A router port will be dynamically configured when IGMP query packets, RIPv2 multicast, DVMRP multicast, PIM-DM multicast packets are detected flowing into a port.

Action: <Add/Modify> Can be toggled between Add / Modify and Delete to allow the addition or the deletion of a static router port to the switch's static router port table.

VLAN Name: [default] Allows the entry of the name of the VLAN the static router port will be a member of. The default VLAN Name is default.

Router Port (M/-): [ ][ ][ ] Highlight the first field of Router Port (M/-):[ ][ ][ ]. Each port can be set individually as a router port by highlighting the port's entry using the arrow keys, and then toggling between M and – using the space bar. M defines the port as a member of the static router port table (that is, the port is a static router port), while the dash (-) defines the port as a non-member (the port is not a static router port). The dash (-) allows the port to be dynamically assigned as a router port by the switch.

Highlight APPLY and press enter to make the changes current. Use Save Changes from the Main Menu to enter the changes into NV-RAM.

To delete an entry, select Delete and enter the VLAN Name of the VLAN for which the router port table entry is to be deleted. Highlight APPLY and press enter. The entry for the VLAN will be deleted. Use Save Changes from the Main Menu to enter the changes into NV-RAM.

VLANs

The switch supports IEEE 802.1Q and Port-based VLANs. The VLAN configuration menus are accessed from VLANs on the Main Menu**.**

Configure VLANs

The switch reserves one VLAN, VID = 1, called 'default' for internal use. The factory default setting assigns all ports on the switch to the 'default'. As new VLANs are configured, their respective member ports are removed from the 'default' VLAN. If the 'default' VLAN is reconfigured, all ports are again assigned to it.

If you are unsure about your knowledge of VLANs, please review the VLANs section of Chapter 5 before configuring the switch for VLANs.

Note: the switch can only support either 802.1Q or Port-based VLANs at any given time; it cannot support more than one type of VLAN at the same time.

To create a new 802.1Q VLAN, or to delete or modify an existing 802.1Q VLAN:

Highlight **VLANs** from the **Main Menu** and press **Enter.**

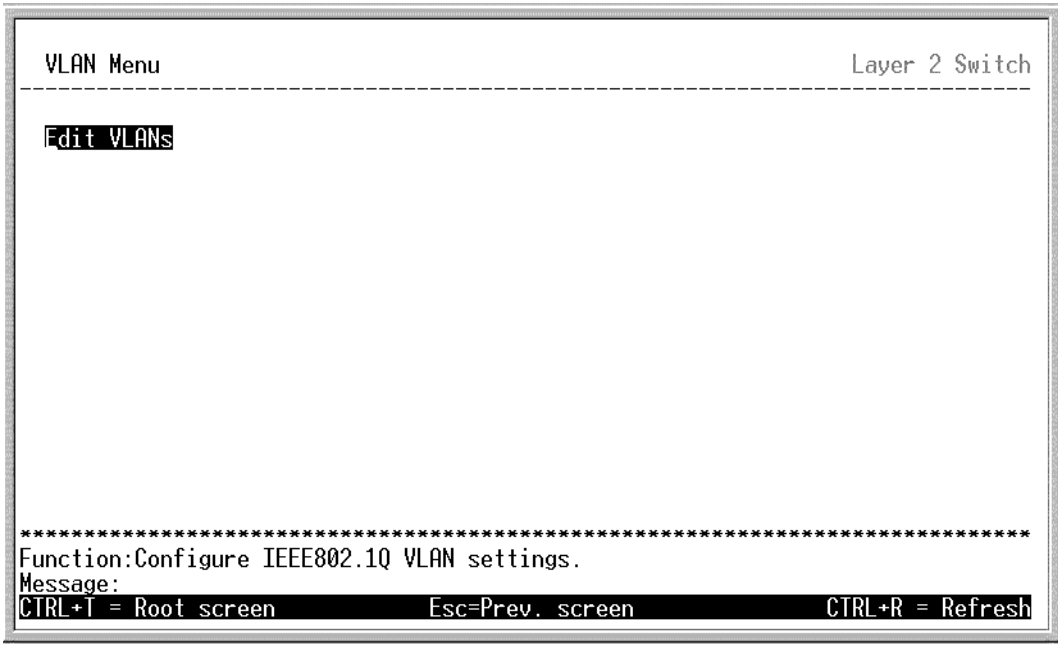

Figure 6-45: VLANs Menu

Highlight Edit VLANs from the VLAN Menu and press **Enter.**

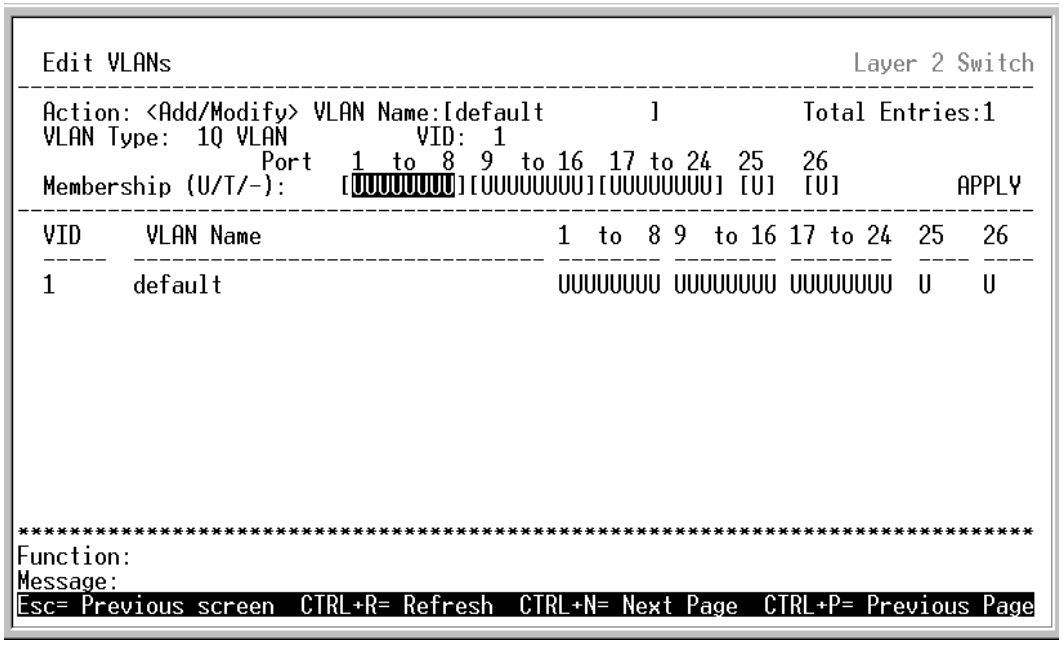

Figure 6-46: Edit VLANs Menu

The following fields can be set:

To create, change or delete an 802.1Q VLAN, use the following fields of the Edit VLANs Menu:

Action: <Add/Modify>	Toggle to select Add/Modify or Delete to add, change or delete a VLAN.
VLAN Name: [ ]	Type in a name for the VLAN to be added, modified or deleted.
VLAN Type:<Port Based VLAN>	Toggle to select <802.1Q VLAN> or <Port Based VLAN> according to the type of VLAN used. To configure port based VLANs see the next section.
Membership (M/): [ ][ ][ ]	Assign a VLAN ID number for the VLAN group.

VLAN membership can be set individually for each port. At the same time, a port can be Tagged or Untagged. To set a port's VLAN membership and its status as being a tagged or untagged port, highlight the first field of (U/T/-):[ ][ ][ ]. Use the arrow keys to move the cursor to the selected ports, then toggle between U, T, or - using the space bar. The status of each port is defined as follows:

U Specifies the port as being an Untagged static member of the VLAN. When an untagged packet is transmitted by the port, the packet header remains unchanged. When a tagged packet exits the port, the tag is stripped and the packet is changed to an untagged packet.

T Specifies the port as a Tagged member of the VLAN. When an untagged packet is transmitted by the port, the packet header is changed to include the 32-bit tag associated with the PVID (Port VLAN Identifier). When a tagged packet exits the port, the packet header is unchanged.

- Specifies the port as not being a member of the VLAN, but the port can become a member of the VLAN dynamically.

When port VLAN membership and its tagging function have been determined for the VLAN group, highlight APPLY and press **Enter** to perform the chosen action (i.e. add, modify or delete a VLAN). Save the changes into NV-RAM using the Save Changes menu.

Note: if the port is attached to a device that is not IEEE 802.1Q VLAN compliant (VLAN-tag unaware), then the port should be set to U - Untagged. If the port is attached to a device that is IEEE 802.1Q VLAN compliant, (VLAN-tag aware), then the port should be set to T - Tagged.

Configuring Port-Based VLANs

To configure port-based VLANs:

1. Highlight VLANs from the Main Menu and press **Enter**.
2. Highlight Edit VLANs and press **Enter**

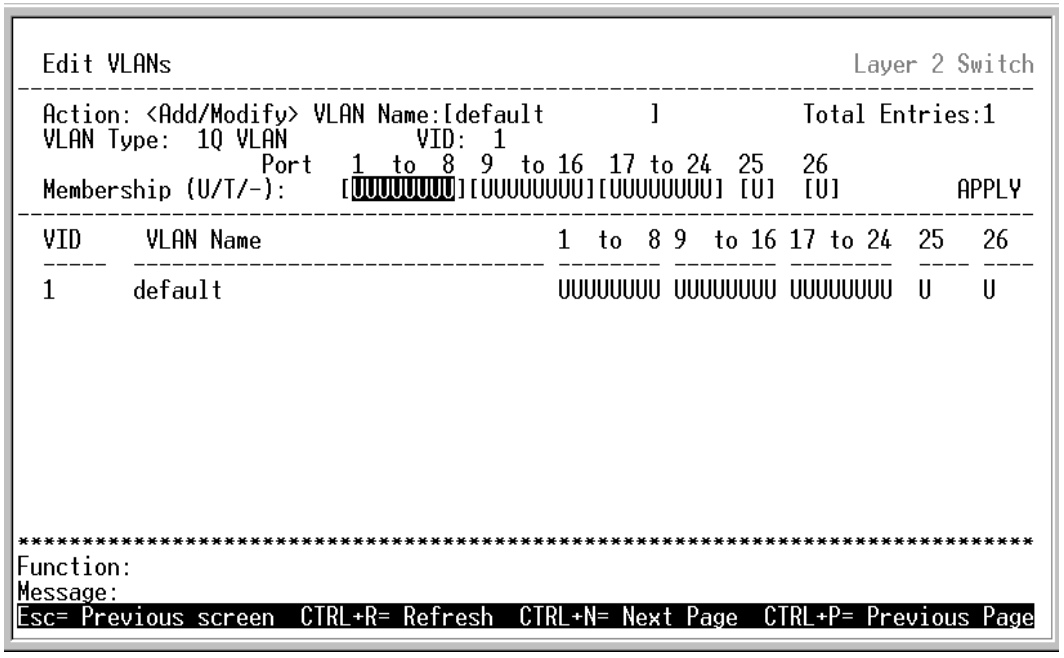

Figure 6-47: Edit VLANs Menu

To create, change or delete a Port-Based VLAN, use the following fields of the Edit VLANs Menu:

Action: <Add/Modify> Toggle to select <Add/Modify> or <Delete> to add, change or delete a VLAN.

VLAN Name: [ ] Type in a name for the VLAN to be added, modified or deleted.

VLAN Type: <Port Based VLAN> Toggle to select <Port Based VLAN>.

VLAN membership is set individually for each port. To set a port's VLAN membership, highlight the first field of (M/-):[ ][ ][ ]. Use the arrow keys to move the cursor to the selected ports, then toggle between M or - using the space bar. The status of each port is defined as follows:

M Assigns membership status to the port. Each port may be assigned to only one VLAN.

- Indicates non-member status of the port.

Highlight APPLY and press **Enter** to for the Port-based VLAN settings to go into effect. Use the Save Changes menu to save the settings to NV-RAM.

Link Aggregation

Link aggregation allows several ports to be grouped together and to act as a single link. This gives a bandwidth that is a multiple of a single link's bandwidth.

Link aggregation is most commonly used to link a bandwidth intensive network device or devices — such as a server or server farm — to the backbone of a network.

Note: the switch allows the creation of up to 6 link aggregation groups, each group consisting of up of up to 8 links (ports). The aggregated links must be contiguous (they must have sequential port numbers) and each group must fall within an 8 port boundary (groups may be within ports 1 to 8, ports 9 to 16, and ports 17 to 24), except the two Gigabit ports - which can only belong to a single link aggregation group. A link aggregation group may not cross an 8 port boundary, starting with port 1 (a group may not contain ports 8 and 9, for example) and all of the ports in the group must be members of the same VLAN. Further, the linked ports must all be of the same speed and should be configured as full-duplex.

The configuration of the lowest numbered port in the group becomes the configuration for all of the ports in the aggregation group. This port is called the base port of the group, and all configuration options - including the VLAN configuration - that can be applied to the base port are applied to the entire link aggregation group.

Load balancing is automatically applied to the links in the aggregation group, and a link failure within the group causes the network traffic to be directed to the remaining links in the group.

The Spanning Tree Protocol will treat a link aggregation group as a single link, on the switch level. On the port level, the STP will use the port parameters of the base port in the calculation of port cost and in determining the state of the link aggregation group. If two redundant link aggregation groups are configured on the switch, STP will block one entire group — in the same way STP will block a single port that has a redundant link.

To configure a link aggregation group follow these steps:

Highlight **Link Aggregation** on the **Main Menu** and press **Enter**. The following menu appears:

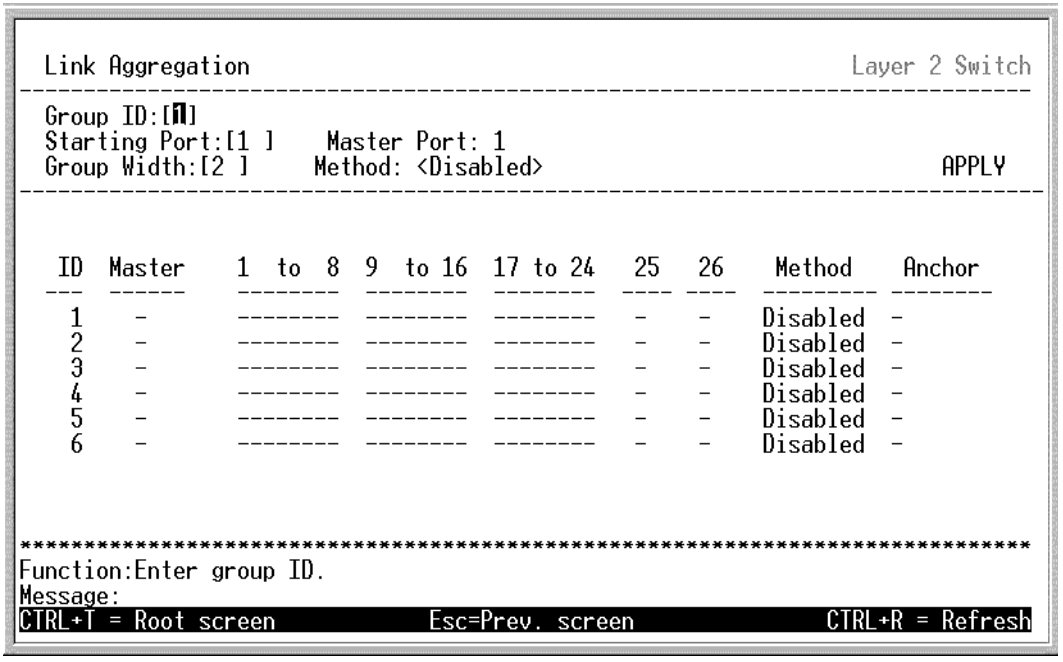

Figure 6-48: Link Aggregation Settings Menu

1 Type an port group identification number in the Group ID:[1] field, up to six groups can be configured, allowable Group ID entries numbers 1 - 6.

2 Toggle the Master Module:<Base Unit> field to configure a group either on the Base Unit (the 24 Fast Ethernet ports) or on the Slot-1 module (the 2 Gigabit Ethernet ports).

The two Gigabit ports can only be grouped together as a link aggregation group. They cannot be combined with Fast Ethernet ports in a link aggregation group.

3 Specify the Group Width:[2]. This is the number of ports, in sequential order from the base port, that will be included in the link aggregation group.

4 Toggle the Method:<Disabled> field to select Enabled or Disabled —this is used to turn a link aggregation group on or off. This is useful for diagnostics, to quickly

isolate a bandwidth intensive network device or to have an absolute backup aggregation group that is not under automatic control.

5 Highlight Apply and press **Enter** to make the link aggregation group configuration active. Use Save Changes from the Main Menu to enter the configuration into NV-RAM.

A Appendix A - Technical Specfications

Table A-1: General Specifications

	General
Network Cables:	
10BaseT:	2-pair UTP Cat. 3,4,5 (100 m max.) EIA/TIA-568 100-ohm STP (100 m max.)
100BaseTX:	2-pair UTP Cat. 5 (100 m max.) EIA/TIA-568 100-ohm STP (100 m max.)
1000BaseTX:	4-pair UTP Cat. 5+ (100 m max.)
Fiber Optic:	IEC 793-2:1992 Type A1a - 50/125um multimode Type A1b - 62.5/125um multimode Both types use MTRJ or SC optical connector
Number of Ports:	24 x 10/100 Mbps ports 2 Gigabit Ethernet

Table A-2: Physical and Environmental Specifications

	Physical and Environmental
AC inputs:	100 - 240 VAC, 50/60 Hz (internal universal power supply)
Power Consumption:	400 watts maximum
DC fans:	4 built-in 40 x 40 x 10mm fans + 2 built-in 40 x 40 x 20mm fans in power supply
Operating Temperature:	0 to 60 degrees Celsius
Storage Temperature:	-25 to 70 degrees Celsius
Humidity:	Operating: 10% to 90% RH non-condensing; Storage: 5% to 95% RH non-condensing
Dimensions:	441 mm x 367 mm x 44 mm (1U), 19 inch rack-mount width
Weight:	7 kg
EMI:	FCC Class A, CE Class A, VCCI Class A, BSMI Class A, C-Tick Class A
Safety:	UL, UL/CUL, CE Mark, TUV/GS

Table A-3: Performance Specifications

	Performance
Transmission Method:	Store-and-forward
RAM Buffer:	16 MB per device
Forwarding Address Table:	8K MAC address max. per device
Packet Filtering/ Forwarding Rate:	Full-wire speed for all connections. 148,800 pps per port (for 100Mbps)
MAC Address Learning:	Automatic update.
Forwarding Table Age Time:	Max age:10–1,000,000 seconds. Default = 300.

B Appendix B - Runtime Switching Software Defaults

Table 2-1: Runtime Switching Software Default Settings

Load Mode	Network Ethernet
Configuration update	Disabled
Firmware update	Disabled
Configuration file name	No Default
Firmware file name	No Default
Out-of-band baud rate	9600
Rs232 mode	Console
Ip address	10.0.0.0
Subnet mask	255.0.0.0
Default router	0.0.0.0
Bootp service	Disabled
TFTP server IP address	0.0.0.0
IGMP time out	260 secs
IGMP snooping state	Disabled
Console time out	15 min
User name	“admin”
Password	“admin”

Table 2-1: Runtime Switching Software Default Settings

Device STP	Disabled
Port STP	Enabled
Port enable	Enabled
Bridge max age	20 secs
Bridge hello time	2 sec
Bridge forward delay	15 sec
Bridge priority	32768
Port STP cost	10Mbps = 100 100Mbps = 19 1000Mbps = 4
Port STP priority	128
Forwarding table aging time	300 secs
Auto-negotiation	Enabled
Flow control	No flow control for 10/100 Mbps ports in full-duplex mode. Gigabit ports will react to "PAUSE" frames received, but will not send "PAUSE" frames when congested.
Backpressure	Always Enabled for 10/100 Mbps ports in half-duplex mode.
Port priority	Normal
Broadcast storm rising action	Do nothing
Upper threshold for base ports	128 Kbps
Upper threshold for module ports	128 Kbps
Community string	"public", "private"
SNMP VLAN(802.1Q)	1
Default port VID	1
Ingress rule checking	Disabled

Index