Part No. 217327-A
February 2005

4655 Great America Parkway
Santa Clara, CA 95054

Application Guide for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3

217327-A

NORTEL

Copyright © 2005 Nortel Networks

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document a re believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel Networks Inc.

The software described in this document is furnished under a license agreement and may be used only in accordance with the terms of that license. The software license agreement is included in this document.

Trademarks

Nortel, the Nortel logo, the Globemark, Unified Networks, BayStack, Autotopology, Optivity and Passport are trademarks of Nortel Networks.
Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation.
Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.
Java is a trademark of Sun Microsystems, Inc.
Cisco is a trademark of Cisco Systems, Inc.

The asterisk after a name denotes a trademarked item.

Restricted rights legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of conditions

In the interest of improving internal design, operational function, and/or reliability, Nortel Networks Inc. reserves the right to make changes to the products described in this document without notice.

Nortel Networks Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

International regulatory statements of conformity

This is to certify that the Nortel Ethernet Routing Switch 3510-24T was evaluated to the international regulatory standards for electromagnetic compliance (EMC) and safety and were found to have met the requirements for the following international standards:

- EMC - Electromagnetic Emissions – CISPR 22, Class A
- EMC - Electromagnetic Immunity – CISPR 24
- Electrical Safety – IEC 60950, with CB member national deviations

Further, the equipment has been certified as compliant with the national standards as detailed below.

National electromagnetic compliance (EMC) statements of compliance

FCC statement (USA only)

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the Federal Communications Commission (FCC) rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy. If it is not installed and used in accordance with the instruction manual, it may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to take whatever measures may be necessary to correct the interference at their own expense.

ICES statement (Canada only)

Canadian Department of Communications Radio Interference Regulations

This digital apparatus (Nortel Ethernet Routing Switch 3510-24T) does not exceed the Class A limits for radio-noise emissions from digital apparatus as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Règlement sur le brouillage radioélectrique du ministère des Communications

Cet appareil numérique (Nortel Ethernet Routing Switch 3510-24T) respecte les limites de bruits radioélectriques visant les appareils numériques de classe A prescrites dans le Règlement sur le brouillage radioélectrique du ministère des Communications du Canada.

CE marking statement (Europe only)

EN 55 022 statements

This is to certify that the Nortel Ethernet Routing Switch 3510-24T is shielded against the generation of radio interference in accordance with the application of Council Directive 89/336/EEC. Conformity is declared by the application of EN 55 022 Class A (CISPR 22).

Caution: This device is a Class A product. In a domestic environment, this device can cause radio interference, in which case the user may be required to take appropriate measures.

EN 55 024 statement

This is to certify that the Nortel Ethernet Routing Switch 3510-24T switches is shielded against the susceptibility to radio interference in accordance with the application of Council Directive 89/336/EEC. Conformity is declared by the application of
EN 55 024 (CISPR 24).

EC Declaration of Conformity

This product conforms to the provisions of the R&TTE Directive 1999/5/EC.

VCCI statement (Japan/Nippon only)

This is a Class A product based on the standard of the Voluntary Control Council for Interference (VCCI) for information technology equipment. If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

BSMI statement for Nortel Ethernet Routing Switch 3510-24T (Taiwan only)

This is a Class A product based on the standard of the Bureau of Standards, Metrology and Inspection (BSMI) CNS 13438, Class A.

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

MIC notice for Nortel Ethernet Routing Switch 3510-24T (Republic of Korea only)

This device has been approved for use in Business applications only per the Class A requirements of the Republic of Korea Ministry of Information and Communications (MIC). This device may not be sold for use in a non-business application.

Observe the Regulatory Marking label on the bottom surface of the chassis for specific certification information pertaining to this model. Each model in the Nortel Ethernet Routing Switch Series which is approved for shipment to/ usage in Korea is labeled as such, with all appropriate text and the appropriate MIC reference number.

National safety statements of compliance

CE marking statement (Europe only)

EN 60 950 statement

This is to certify that the Nortel Ethernet Routing Switch 3510-24T is in compliance with the requirements of EN 60 950 in accordance with the Low Voltage Directive. Additional national differences for all European Union countries have been evaluated for compliance.

NOM statement Nortel Ethernet Routing Switch 3510-24T (Mexico only)

The following information is provided on the devices described in this document in compliance with the safety requirements of the Norma Oficial Méxicana (NOM):

Exporter: Nortel Networks, Inc.
4655 Great America Parkway
Santa Clara CA 95054 USA

Importer: Nortel Networks de México, S.A. de C.V.
Avenida Insurgentes Sur #1605
Piso 30, Oficina
Col. San Jose Insurgentes
Deleg-Benito Juarez
México D.F. 03900

Tel: 52 5 480 2100

Fax: 52 5 480 2199

Input: Nortel Ethernet Routing Switch 3510-24T

100 - 240 VAC 50/60 Hz 1.3A max

Información NOM (unicamente para México)

La información siguiente se proporciona en el dispositivo o en los dispositivos descritos en este documento, en cumplimiento con los requisitos de la Norma Oficial Méxicana (NOM):

Exportador: Nortel Networks, Inc.
4655 Great America Parkway
Santa Clara, CA 95054 USA

Importador: Nortel Networks de México, S.A. de C.V.
Avenida Insurgentes Sur #1605
Piso 30, Oficina
Col. San Jose Insurgentes
Deleg-Benito Juarez
México D.F. 03900

Tel: 52 5 480 2100

Fax: 52 5 480 2199

Embarcar a: Nortel Ethernet Routing Switch 3510-24T

100 - 240 VAC 50/60 Hz 1.3A max

Nortel Networks Inc. software license agreement

This Software License Agreement (“License Agreement”) is between you, the end-user (“Customer”) and Nortel Networks Corporation and its subsidiaries and affiliates (“Nortel Networks”). PLEASE READ THE FOLLOWING CAREFULLY. YOU MUST ACCEPT THESE LICENSE TERMS IN ORDER TO DOWNLOAD AND/OR USE THE SOFTWARE. USE OF THE SOFTWARE CONSTITUTES YOUR ACCEPTANCE OF THIS LICENSE AGREEMENT. If you do not accept these terms and conditions, return the Software, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

“Software” is owned or licensed by Nortel Networks, its parent or one of its subsidiaries or affiliates, and is copyrighted and licensed, not sold. Software consists of machine-readable instructions, its components, data, audio-visual content (such as images, text, recordings or pictures) and related licensed materials including all whole or partial copies. Nortel Networks grants you a license to use the Software only in the country where you acquired the Software. You obtain no rights other than those granted to you under this License Agreement. You are responsible for the selection of the Software and for the installation of, use of, and results obtained from the Software.

1. Licensed Use of Software. Nortel Networks grants Customer a nonexclusive license to use a copy of the Software on only one machine at any one time or to the extent of the activation or authorized usage level, whichever is applicable. To the extent Software is furnished for use with designated hardware or Customer furnished equipment (“CFE”), Customer is granted a nonexclusive license to use Software only on such hardware or CFE, as applicable. Software contains trade secrets and Customer agrees to treat Software as confidential information using the same care and discretion Customer uses with its own similar information that it does not wish to disclose, publish or disseminate. Customer will ensure that anyone who uses the Software does so only in compliance with the terms of this Agreement. Customer shall not a) use, copy, modify, transfer or distribute the Software except as expressly authorized; b) reverse assemble, reverse compile, reverse engineer or otherwise translate the Software; c) create derivative works or modifications unless expressly authorized; or d) sublicense, rent or lease the Software. Licensors of intellectual property to Nortel Networks are beneficiaries of this provision. Upon termination or breach of the license by Customer or in the event designated hardware or CFE is no longer in use, Customer will promptly return the Software to Nortel Networks or certify its destruction. Nortel Networks may audit by remote polling or other reasonable means to determine Customer’s Software activation or usage levels. If suppliers of third party software included in Software require Nortel Networks to include additional or different terms, Customer agrees to abide by such terms provided by Nortel Networks with respect to such third party software.

2. Warranty. Except as may be otherwise expressly agreed to in writing between Nortel Networks and Customer, Software is provided “AS IS” without any warranties (conditions) of any kind. NORTEL NETWORKS DISCLAIMS ALL WARRANTIES (CONDITIONS) FOR THE SOFTWARE, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABLITITY AND FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OF NON-INFRINGEMENT. Nortel Networks is not obligated to provide support of any kind for the Software. Some jurisdictions do not allow exclusion of implied warranties, and, in such event, the above exclusions may not apply.

3. Limitation of Remedies. IN NO EVENT SHALL NORTEL NETWORKS OR ITS AGENTS OR SUPPLIERS BE LIABLE FOR ANY OF THE FOLLOWING: a) DAMAGES BASED ON ANY THIRD PARTY CLAIM; b) LOSS OF, OR DAMAGE TO, CUSTOMER’S RECORDS, FILES OR DATA; OR c) DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS OR SAVINGS), WHETHER IN CONTRACT, TORT OR OTHERWISE (INCLUDING NEGLIGENCE) ARISING OUT OF YOUR USE OF THE SOFTWARE, EVEN IF NORTEL NETWORKS, ITS AGENTS OR SUPPLIERS HAVE BEEN ADVISED OF THEIR POSSIBILITY. The forgoing limitations of remedies also apply to any developer and/or supplier of the Software. Such developer and/or supplier is an intended beneficiary of this Section. Some jurisdictions do not allow these limitations or exclusions and, in such event, they may not apply.

4. General

a) If Customer is the United States Government, the following paragraph shall apply: All Nortel Networks Software available under this License Agreement is commercial computer software and commercial computer software documentation and, in the event Software is licensed for or on behalf of the United States Government, the respective rights to the software and software documentation are governed by Nortel Networks standard commercial license in accordance with U.S. Federal Regulations at 48 C.F.R. Sections 12.212 (for non-DoD entities) and 48 C.F.R. 227.7202 (for DoD entities).

b) Customer may terminate the license at any time. Nortel Networks may terminate the license if Customer fails to comply with the terms and conditions of this license. In either event, upon termination, Customer must either return the Software to Nortel Networks or certify its destruction.

c) Customer is responsible for payment of any taxes, including personal property taxes, resulting from Customer's use of the Software. Customer agrees to comply with all applicable laws including all applicable export and import laws and regulations.

d) Neither party may bring an action, regardless of form, more than two years after the cause of the action arose.

e) The terms and conditions of this License Agreement form the complete and exclusive agreement between Customer and Nortel Networks.

f) This License Agreement is governed by the laws of the country in which Customer acquires the Software. If the Software is acquired in the United States, then this License Agreement is governed by the laws of the state of New York.

Revision History

Date Revised	Version	Reason for revision
February 2005	1.0	Nortel Ethernet Routing Switch 3510-24T.

Contents

Chapter 6

Appendix A

Appendix B

Appendix C

Figures

Tables

Preface

This guide describes the Nortel* Ethernet Routing Switch 3510-24T features and its uses.

Ethernet Routing Switch 3510-24T provides low-cost 10/100/1000-Mbps switching to the desktop and scalable Gigabit aggregation switching. These products provide high-bandwidth Gigabit edge connectivity, increase the server and desktop connections speeds, and significantly improve bandwidth.

The Ethernet Routing Switch 3510-24T is in a standalone switch configuration mode.

This chapter covers the following topics:

- "Before you Begin", next
- "Related Publications" on page 28
- "How to Get Help" on page 29

Before you Begin

This guide is intended for network managers and administrators with the following background:

- Basic knowledge of Networks, Ethernet bridging, and IP routing
- Familiarity with networking concepts and terminology
- Specific knowledge about the networking devices, protocols, topologies, and interfaces that comprise your network
- Experience with Windowing systems, Graphical User Interfaces (GUIs), or Web browsers

Related Publications

For more information about using the Ethernet Routing Switch 3510-24T, refer to the following publications:

- *Release Notes for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3* (part number 217331-A)

  Documents important changes about the software and hardware that are not covered in other related publications.

- *Installing the Nortel Ethernet Routing Switch 3510-24T* (part number 217326-A)

  Describes how to install the Ethernet Routing Switch 3510-24T.

- *Switch Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3* (part number 217329-A)

  Describes how to use the Java-based device-level software management application.

- *Web Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3* (part number 217328-A)

  Describes how to use the Web-based management tool to configure switch features.

- *NCLI Configuration Guide for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3* (part number 217330-A)

  Describes how to use Nortel Command Line Interface (NCLI) commands to configure and manage the Ethernet Routing Switch 3510-24T.

You can print selected technical manuals and release notes free, directly from the Internet. Go to the www.nortel.com/support. Find the product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Use Adobe* Acrobat Reader* to open the manuals and release notes, search for the sections you need, and print them on most standard printers. Go to Adobe Systems at the www.adobe.com to download a free copy of Adobe Acrobat Reader.

How to Get Help

If you have purchased a service contract for your Nortel product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you have purchased a Nortel service program, contact Nortel Technical Support. To obtain contact information online, go to the www.nortel.com/support, then click on Technical Support.

From the Technical Support page, you can open a Customer Service Request online or find the telephone number for the nearest Technical Solutions Center. If you are not connected to the Internet, you can call 1-800-4NORTEL (1-800-466-7835) to learn the telephone number for the nearest Technical Solutions Center.

An Express Routing Code (ERC) is available for many Nortel products and services. When you use an ERC, your call is routed to a technical support person who specializes in supporting that product or service. To locate an ERC for your product or service, go to www.nortel.com/support and click Express Routing Codes located on the right side at the bottom of the page.

Chapter 1
The Nortel Ethernet Routing Switch 3510-24T

This chapter introduces the Ethernet Routing Switch 3510-24T and covers the following topics:

General Description

This guide describes the Ethernet Routing Switch 3510-24T.

The Ethernet Routing Switch 3510-24T provides low-cost 10/100/1000-Mbps switching to the desktop and scalable Gigabit aggregation switching. These products provide high-bandwidth Gigabit edge connectivity, increase the server and desktop connections speeds, support IP routing, and significantly improves bandwidth.

The Ethernet Routing Switch 3510-24T has 24 10/100/1000-BASE-TX ports, as well as 4 ports for SFP GBICs. Ports 21-24 are shared ports. The last 4 10/100/1000-BASE-TX ports are ports 21, 22, 23 and 24, as are the 2 SFP GBICs ports; thus they cannot be active simultaneously. When an SFP GBIC is active in ports 21, 22, 23 and 24, the corresponding 10/100/1000-BASE-TX port numbers 21, 22, 23 and 24 are unavailable.

The Ethernet Routing Switch 3510-24T Differentiated Services (DiffServ) network architecture offers varied levels of service for different types of data traffic. DiffServ lets you designate a specific level of performance on a per-packet basis. Packet classification is based on Layer 2/3/4 information and implements an action list based on the classification. The packet classification occurs across multiple layers at wire speed to provide security, filtering, and Quality of Service (QoS) at the edge of the network.

Physical Description

Figure 1 depicts the front panel of the Ethernet Routing Switch 3510-24T, and Figure 2 shows the front and side views of the Ethernet Routing Switch 3510-24T with rack mounting brackets.

Figure 1 Ethernet Routing Switch 3510-24T

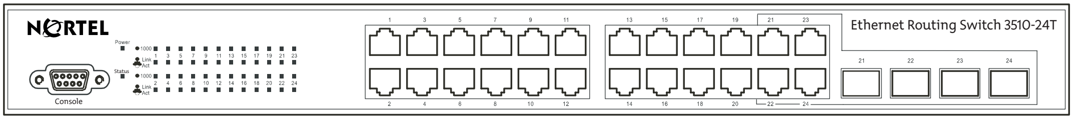

Figure 2 Ethernet Routing Switch 3510-24T

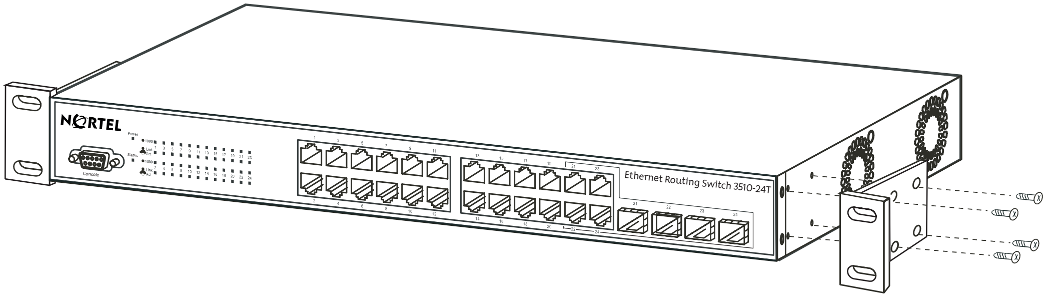

Front Panel

Figure 3 shows the front-panel configuration for the Ethernet Routing Switch 3510-24T. Descriptions of the front-panel components follow the figures (Table 1).

For descriptions of the back-panel Ethernet Routing Switch 3510-24T components, see “Back Panel” on page 38.

Figure 3 Ethernet Routing Switch 3510-24T front panel

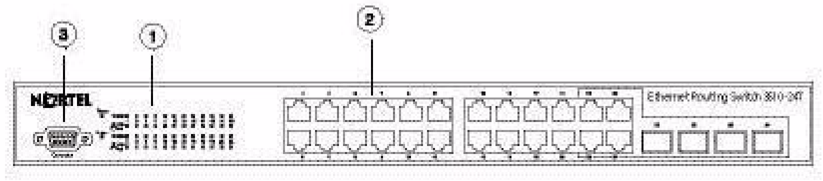

Table 1 Ethernet Routing Switch 3510-24T front-panel description

1	Switch LED’s
2	Port Connectors
3	Console port

Console Port

The console port allows you to access the Console Interface (CI) screens and customize your network using the supplied menus and screens (see Chapter 3).

The console port is a DB-9, RS-23. You must use a male-female serial port connector to connect a management station or console/terminal to the Ethernet Routing Switch 3510-24T by using a straight-through DB-9 to DB-9 standard serial port cable. You must use a VT100/ANSI-compatible terminal (for cursor control and to enable cursor and functions keys) to use the console port.

Note: The console port is configured as a data communications equipment (DCE) connector. Ensure that your RS-232 cable pinouts are configured for DCE connections (see Appendix E).

The console port default settings are:

- 9600 baud
- 8 data bits
- 1 stop bit
- no parity as the communications format
- flow control set to enabled (Xon/Xoff supported)

Small Form Factor Pluggable Gigabit Interface Converter

Small Form Factor Pluggable Gigabit Interface Converters (SFP GBICs) are hot-swappable input/output products to allow Gigabit Ethernet ports to link with Short Wavelength (SX), Long Wave length (LX), and Coarse Wavelength Division Multiplexed (CWDM) fiber optic networks.

Port Connectors

The Ethernet Routing Switch 3510-24T uses 10/100/1000BASE-TX RJ-45 (8-pin modular) port connectors.

The 10/100/1000BASE-TX port connectors feature auto-MDI/X (media-dependent interface-crossover). These ports connect over straight-through cables to the Network Interface Card (NIC) in a node or server, similar to a conventional Ethernet repeater hub. However, with this feature and auto-negotiation enabled, you can still use straight-through cables while connecting to an Ethernet hub or switch.

The Ethernet Routing Switch 3510-24T features auto-polarity. With autonegotiation enabled, auto-polarity automatically reverses the polarity of a pair of pins from positive to negative or negative to positive. This corrects the polarity of the received data if the port detects that the polarity of the data has been reversed due to a wiring error.

For details on pin assignments, see Appendix C.

The Ethernet Routing Switch 3510-24T uses autosensing ports designed to operate at 10 Mbps, 100 Mbps, OR 1000 Mbps (1 GB) depending on the connecting device. These ports support the IEEE 802.3u, 802.3z for 1000SX, or 802.3ab for 1000TX autonegotiation standard, which means that when a port is connected to another device that also supports the IEEE 802.3u, 802.3z for SFP GBIC, or 802.3ab for 1000TX standard, the two devices negotiate the best speed and duplex mode.

Note: In autonegotiation mode, the Ethernet Routing Switch 3510-24T automatically provides the proper MDI/MDI-X connection on the RJ-45 ports, thereby eliminating the need for crossover cables. When autonegotiation is disabled, the RJ-45 ports provide an MDI-X connection, which allows end-station equipment to be connected using straight-through cables. To connect other MDI-X port devices, such as another switch or a hub, a crossover cable must be used.

The 10/100/1000BASE-TX switch ports support half as well as full-duplex modes operation at 10 Mbps and 100 Mbps.

The 10/100/1000BASE-TX RJ-45 ports can connect to 10 Mbps or 100 Mbps or 1000 Mbps (1 GB) Ethernet segments or nodes.

Note: You can use Category 5 copper unshielded twisted pair (UTP) cable when you are running at speeds of 10 Mbps, 100 Mbps, or 1000 Mbps. You can use Category 3 cable when you are running at speeds of 10 Mbps.
If you use Category 3 cable, you must disable autonegotiation when the device connected to the Ethernet Routing Switch 3510-24T supports a 100-Mbps or 1000-Mbps connection. If the connected device supports only 10-Mbps connection, autonegotiation may be enabled.

Note: IEEE 1000BASE-TX requires operating in full-duplex mode with auto-negotiation enabled.

See Appendix C for more information about the RJ-45 port connectors.

LED Display Panel

Figure 4 shows the Ethernet Routing Switch 3510-24T LED display panel. See Table 2 and Table 3 for a description of the Ethernet Routing Switch 3510-24T LEDs.

Figure 4 Ethernet Routing Switch 3510-24T LED display panel

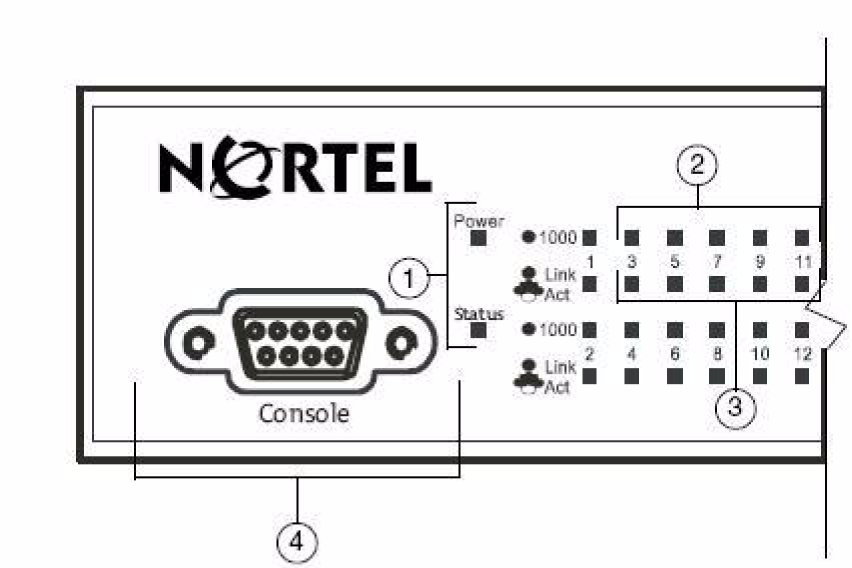

1 = Switch LEDs
2 = 10/100/1000-Mbps LEDs
3 = Link/Activity LEDs
4 = Console Port

Table 2 Ethernet Routing Switch 3510-24T LED descriptions

Label	State	Meaning
Power	On	The switch is connected to AC power and is receiving power.
	Off	The switch is not connected to AC power, or the AC power is not supplying power.
Status	Steady	The power-on self-test is complete, and the switch is operating normally.
	Blinking	A nonfatal error occurred during the self-test.
	Off	The switch failed the self-test.

Table 3 10/100/1000 Port LEDs on the Ethernet Routing Switch 3510-24T

Label	Color/State	Meaning
1000	On	The port is set to operate at 1000 Mbps.
	Off	The port is set to operate at 10 or 100 Mbps (and LNK/ACT is green). When the LED is off, refer to the LNK/ACT Section.

Table 3 10/100/1000 Port LEDs on the Ethernet Routing Switch 3510-24T (continued)

Label	Color/State	Meaning
LNK/ACT	Steady	The link is good.
	Blinking	There is activity on this port. The blinking rate indicates the level of activity.
	Slow blinking	This port has been disabled by software.
	Off	This port has no link or activity.

LED indications during software download

When you download software to the Ethernet Routing Switch 3510-24T, the port LEDs light one after another in a chasing pattern (except the LEDs on ports 11, 12, 23, and 24 on a Ethernet Routing Switch 3510-24T.

While downloading the image, the pattern is fast, and then the pattern slows as the switch erases the flash memory. The pattern moves very fast as the switch programs the new image into the switch's memory. When the process is complete, the port LEDs are no longer lit and the switch resets.

Back Panel

The switch back panel is shown in Figure 5. See Table 4 for a description of the Ethernet Routing Switch 3510-24T back panel.

Figure 5 Ethernet Routing Switch 3510-24T back panel

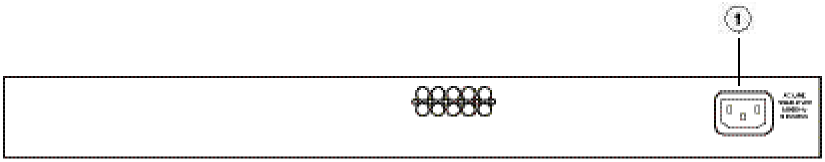

Table 4 Ethernet Routing Switch 3510-24T back-panel descriptions

1	AC Power Receptacle

Cooling Fans

Two cooling fans are located on the right side and one fan on the back panel of the Ethernet Routing Switch 3510-24T to provide cooling of the internal components. (See Figure 1 on page 32.) When you install the switch, be sure to allow enough space on both sides of the switch for adequate air flow. See *Installing the Nortel Ethernet Routing Switch 3510-24T* for detailed information.

AC Power Receptacle

The AC power receptible accepts the AC power cord (supplied). For installation outside of North America, make sure that you have the proper power cord for your region. Your cord must have a CEE-22 standard V female connector on one end. It must also meet the IEC 320-030 specifications.

Table 5 lists the specifications for international power cords.

Table 5 International power cord specifications

Country/Plug description	Specifications	Typical plug
Continental Europe: • CEE7 standard VII male plug • Harmonized cord (HAR marking on the outside of the cord jacket to comply with the CENELEC Harmonized Document HD-21)	220 or 230 VAC 50 Hz Single phase	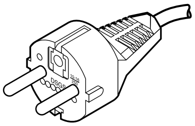 228FA
U.S./Canada/Japan: • NEMA5-15P male plug • UL recognized (UL stamped on cord jacket) • CSA certified (CSA label secured to the cord)	100 or 120 VAC 50–60 Hz Single phase	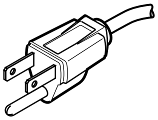 227FA

Table 5 International power cord specifications (continued)

Country/Plug description	**Specifications**	**Typical plug**
United Kingdom: • BS1363 male plug with fuse • Harmonized cord	240 VAC 50 Hz Single phase	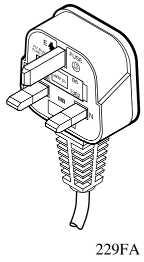
Australia: • AS3112-1981 Male plug	240 VAC 50 Hz Single phase	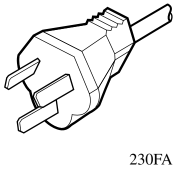

Features

The Ethernet Routing Switch 3510-24T provides wire-speed switching that allows high-performance and low-cost connections to full-duplex and half-duplex 10/100/1000 Mbps ethernet Local Area Networks (LANs). The Ethernet Routing Switch 3510-24T provides the following features:

- Hardware
  - 10/100/1000BASE-TX front-panel switching ports
  - 4 SFP GBIC ports (shared)
- Software
  - Policy-enabled networks—uses Quality of Service (QoS)
  - Virtual LAN (VLANs)
  - Automatic PVID
  - Multiple Spanning Tree Protocol Groups (MSTGs)
  - Security—RADIUS security, EAPOL security, MAC-based security, MAC DA-based security, IP address-based security
  - Flash memory storage
  - MultiLink Trunking
  - Port mirroring
  - Autosensing and autonegotiation
  - ASCII configuration file
  - Support for 47 RMON alarms
  - Support for jumbo packets
  - Support for IGMP
  - SNMPv3 Support
  - SNMP MIB support
  - SNMP trap support

Flash Memory Storage

Switch software image storage

The Ethernet Routing Switch 3510-24T uses flash memory to store the switch software image. The flash memory allows you to update the software image with a newer version without changing the switch hardware (see Chapter 3). An in-band connection between the switch and the TFTP load host is required to download the software image.

Configuration parameters storage

All configuration parameters are stored in flash memory. These parameters are updated every 60 seconds (if a change occurs) or whenever a reset command is executed.

➡ **Note:** Do not power off the switch within 60 seconds of changing any configuration parameters. Powering down the switch within 60 seconds of changing configuration parameters can cause the changed configuration parameters to be lost.

Policy-enabled Networking

The Ethernet Routing Switch 3510-24T enables system administrators to implement classes of service and assign priority levels to different types of traffic. You can configure policies to monitor the characteristics of traffic (for example, its source, destination, and protocol) and perform a controlling action on the traffic when certain user-defined characteristics match.

Differentiated Services (DiffServ) is a network architecture that lets service providers and enterprise network environments offer varied levels of service for different types of data traffic. Instead of using the “best-effort” service model to ensure data delivery, DiffServ’s Quality of Service (QoS) lets you designate a specific level of performance on a packet-by-packet basis. If you have applications that require high performance and reliable service, such as voice and video over IP, you can use DiffServ to give preferential treatment to this data over other traffic.

Use the Web-based management system, the NCLI, or DM to configure QoS features. For detailed information on DiffServ, QoS, and policy-enabled networking refer to Chapter 4.

To configure this feature using the Web-based management system, refer to *Web Management for the Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3*. To use Device Manager (DM) to configure QoS, refer to *Switch Management for the Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3*. And, to configure this feature using NCLI commands, refer to *NCLI Configuration Guide for the Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3*.

Virtual Local Area Networks

In a traditional shared-media network, traffic generated by a station is transmitted to all other stations on the local segment. Therefore, for any given station on the shared ethernet, the local segment is the collision domain. This is because the traffic on the segment has the potential to cause an ethernet collision. The local segment is also the broadcast domain because all broadcasts are sent to all stations on the local segment. Although ethernet switches and bridges divide a network into smaller collision domains, they do not affect the broadcast domain. In simple terms, a Virtual Local Area Network (VLAN) provides a mechanism to fine-tune broadcast domains.

Note: You cannot use a value for a VLAN ID (VID) that is already used for Multiple Spanning Tree Group Protocol (MSTG) Bridge Protocol Data Units (BPDUs).

Note: For information on configuring VLANs, Spanning Tree Groups (STGs), and MultiLink Trunking (MLTs), refer to Chapter 2.

The Ethernet Routing Switch 3510-24T allows you to create two types of VLANs:

- IEEE 802.1Q port-based VLANs

Port-based VLANs filter on the 802.1Q value of the packet. Tagged packets ingress the device containing an 802.1Q value. Untagged packets assume the PVID value assigned to the ingressing port as the 802.1Q value. The packets are forwarded only to those ports with VLAN membership lists containing the same 802.1Q value as the packet.

Automatic PVID automatically sets the PVID when you configure a port-based VLAN. When the port is added to the VLAN, the PVID value will be the same value as the first VLAN ID you associated with this port. The user can also manually change the PVID value.

The default global setting for AutoPVID is enabled.

- Protocol-based VLANs

  A protocol-based VLAN is a VLAN in which you assign your switch ports as members of a broadcast domain, based on the protocol information within the packet. Protocol-based VLANs can localize broadcast traffic and assure that only the protocol-based VLAN ports are flooded with the specified protocol type packets. The maximum number of available protocols is 7.

The order in which the rules for VLAN classification are applied are:

1. Is the packet tagged?
2. Does the packet belong in a protocol-based VLAN?

If none of the criteria applies, the packet belongs to the VLAN identified by the PVID of the ingress port. See Chapter 2, for more information.

Each VLAN uses its own forwarding database and operates in Independent VLAN Learning (IVL) mode. Forwarding information is not shared between VLANs.

The Ethernet Routing Switch 3510-24T supports up to 256 VLANs (port- or protocol-based), including VLAN #1 which is always port-based. The 256 VLANs is a standalone Ethernet Routing Switch 3510-24T.

When a switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belong to one broadcast domain. You can assign different ports (and therefore the devices attached to these ports) to different broadcast domains. This feature allows network flexibility because you can reassign VLANs to accommodate network moves, additions, and changes, eliminating the need to change physical cabling.

Refer to Chapter 2, for more information on VLANs. For information on configuring VLANs using the CI menus, refer to Chapter 3. To configure this feature using the Web-based management system, refer to *Web Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3*. To use Device Manager (DM) to configure VLANs, refer to S*witch Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3*. And, to configure this feature using CLI commands, refer to *NCLI Configuration Guide for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3*.

Multiple Spanning Tree Protocol Groups

The Ethernet Routing Switch 3510-24T supports the Spanning Tree Protocol (STP) as defined in IEEE 802.1D. As defined in the IEEE 802.1D standard, the Spanning Tree Protocol detects and eliminates logical loops in a bridged or switched network. When multiple paths exist, the spanning tree algorithm configures the network in such a way that a bridge or switch uses only the most efficient path. If that path fails, the protocol automatically reconfigures the topology to select a new active path. The STG or bridge group forms a loop-free topology that includes one or more virtual LANs (VLANs). The Ethernet Routing Switch 3510-24T supports multiple instances (8) of STGs running simultaneously.

The Ethernet Routing Switch 3510-24T supports a maximum of 256 VLANs. With a maximum of 8 STGs, each STG can have 32 VLANs.

Security

The Ethernet Routing Switch 3510-24T provides the following levels of security:

- RADIUS-based security — limits administrative access to the switch through user authentication
- MAC address-based security — limits access to the switch based on allowed source and destination MAC addresses
- EAPOL-based security (IEEE 802.1X) — allows the exchange of authentication information between any end station or server connected to the switch and authentication server (such as a RADIUS server)
- IP manager list — limits access to management features of the switch based on the management station's IP address
- SNMPv3 — allows access using password authentication (MD5)

Example of RADIUS, EAPOL, and MAC Addressed-based Security

Figure 6 shows a typical campus configuration using the RADIUS-based and MAC address-based security features for the Ethernet Routing Switch 3510-24T. This example assumes that the switch, the teachers' offices, classrooms, and the library are physically secured. The student dormitory may also be physically secured.

Figure 6 Ethernet Routing Switch 3510-24T security feature

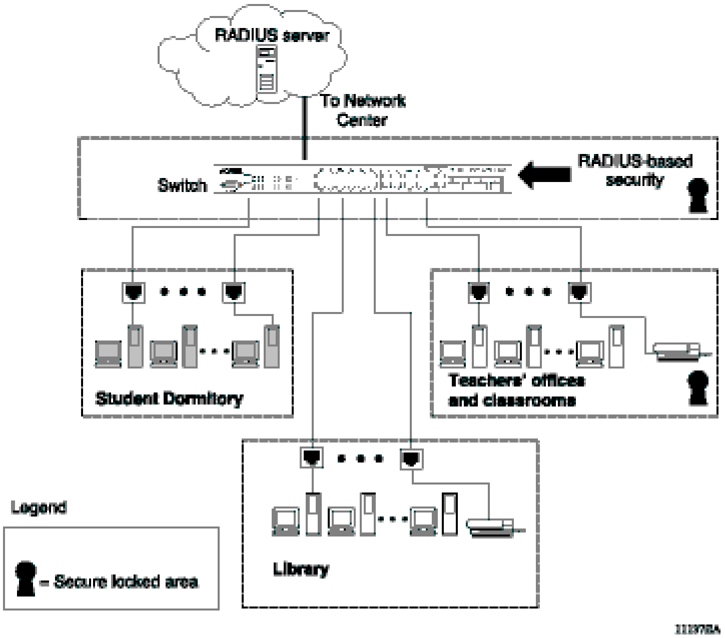

In this configuration example, the following security measures are implemented:

- The switch
  - — RADIUS-based security is used to limit administrative access to the switch through user authentication (see "RADIUS-based Network Security" on page 47).
  - — MAC address-based security is used to allow up to 448 authorized stations (MAC addresses) access to one or more switch ports (see "MAC Address-based Security" on page 49).

  - — The switch is located in a locked closet, accessible only by authorized Technical Services personnel.

- Student dormitory

  Dormitory rooms are typically occupied by two students and has been prewired with two RJ-45 jacks. Only students who are authorized (as specified by the MAC address-based security feature) can access the switch on the secured ports.

- Teachers’ offices and classrooms

  The PCs that are located in the teachers’ offices and in the classrooms are assigned MAC address-based security that is specific for each classroom and office location. The security feature logically locks each wall jack to the specified station and prevents unauthorized access to the switch, thereby preventing anyone from connecting a personal laptop PC into the wall jack. The printer is assigned to a single station and is allowed full bandwidth on that switch port.

  It is assumed that all PCs are password protected and that the classrooms and offices are physically secured.

- Library

  The wall jacks in the library are set up so that the PCs can be connected to any wall jack in the room. This arrangement allows the PCs to be moved anywhere in the room. The exception is the printer, which is assigned to a single station with full bandwidth to that port.

  It is assumed that all PCs are password protected and that access to the library is physically secured.

RADIUS-based Network Security

The RADIUS-based security feature allows you to set up network access control, using the Remote Authentication Dial-In User Services (RADIUS) security protocol. The RADIUS-based security feature uses the RADIUS protocol to authenticate local console and telnet logins.

Three retries for each server(primary and secondary) is allowed; the timeout between each retry is 2 seconds.

You will need to set up specific user accounts (user names and passwords, and Service-Type attributes) on your RADIUS server before the authentication process can be initiated.

To provide each user with an appropriate level of access to the switch, set the following username attributes on your RADIUS server:

- Read-write access—Set the Service-Type field value to Administrative.
- Read-only access—Set the Service-Type field value to NAS-Prompt.

The system can use the local password if the RADIUS server is unavailable to authenticate the user for administrative access. This option is disabled by default.

RADIUS Password Fallback

The RADIUS password fallback feature allows you to login to the switch by using the local password. This option is disabled by default.

The RADIUS server password can be turned on via NCLI using the command ‘radius server password fallback’. You can use the ‘no radius-server’ command to disable this feature, along with the rest of the RADIUS configuration.

RADIUS Access Challenge

This release of Ethernet Routing Switch 3510-24T provides support for RADIUS access challenge as specified in RFC 2138. No additional configuration of the switch is required, as the RADIUS access challenge feature is always enabled.

The RADIUS access challenge feature provides security of authentication by challenging users with more levels of challenges and passwords.

The following depicts the authentication process:

1. Telnet into the switch.
    - The switch will prompt you for the RADIUS Authentication Information.
2. Enter your username and password at the prompt.
    - This request is sent to the RADIUS server
    - The RADIUS server may proxy to another authentication server

- Switch receives an Access-Challenge with State and Challenge Reply Attributes from the RADIUS server
- Switch displays the Challenge Screen with reply information for the user

3 Enter Challenge Response which may be the next password.

- Original username and new information are sent to the server
- Authentication procedure similar to Step 2 is followed once again
- More challenges cause repeat of Step 2
- Access-Accept or Access-Reject complete the Authentication cycle

MAC Address-based Security

The MAC address-based security feature is based on Nortel BaySecure* LAN Access for Ethernet, a real-time security system that safeguards ethernet networks from unauthorized surveillance and intrusion.

The MAC address-based security feature allows you to set up network access control, based on source MAC addresses of authorized stations.

You can:

- Create a list of up to 10 MAC Destination Addresses (DAs) that you want to filter. All packets with the specified DAs are dropped. The packet with the specified MAC DA will be dropped regardless of the ingress port, source address (SA) intrusion, or VLAN membership. (You must use the Web-based Management System or the NCLI Management System for this feature.)
- Create a list of up to 448 MAC source addresses (SAs) and specify which SAs are authorized to connect to your switch. The 448 MAC SAs can be configured within a single standalone switch.
  - Specify which of your switch ports can be accessed by each MAC SA.

    The options for allowed port access include: NONE, ALL, and single or multiple ports that are specified in a list, for example, 1-4,6,9.
  - Specify optional actions to be exercised by your switch if the software detects an SA security violation.

    The response can be to send a trap, turn on Destination Address (DA) filtering for the specified SAs, disable the specific port, or any combination of these three options.

You can configure the Ethernet Routing Switch 3510-24T to drop all packets with specified MAC destination addresses (DA). You can enter up to 10 specific MAC DAs you want filtered.

➡ **Note:** You must use either the CLI or the Web-based Management System to configure MAC DA filtering.

EAPOL-based Security

The Ethernet Routing Switch 3510-24T provides support for security based on the Extensible Authentication Protocol over LAN (EAPOL), which uses the EAP as described in the IEEE 802.1X to allow you to set up network access control on internal LANs.

Extensible Authentication Protocol (EAP) allows the exchange of authentication information between any end station or server connected to the switch and an authentication server (such as a RADIUS server). The EAPOL-based security feature operates in conjunction with a RADIUS-based server to extend the benefits of remote authentication to internal LAN clients.

The following example illustrates how the Ethernet Routing Switch 3510-24T, configured with the EAPOL-based security feature, reacts to a new network connection:

- The switch detects a new connection on one of its ports.
  - The switch requests a user ID from the new client.
  - EAPOL encapsulates the user ID and forwards it to the RADIUS server.
  - The RADIUS server responds with a request for the user’s password.
- The new client forwards a password to the switch, within the EAPOL packet.
  - The switch relays the EAPOL packet to the RADIUS server.
  - If the RADIUS server validates the password, the new client is allowed access to the switch and the network.

Some components and terms used with EAPOL-based security are:

- Supplicant—the device applying for access to the network.
- Authenticator—software with the sole purpose of authorizing a supplicant that is attached to the other end of a LAN segment.

- Authentication Server—a RADIUS server that provides authorization services to the Authenticator.
- Port Access Entity (PAE)—a software entity associated with each port that supports the Authenticator or Supplicant functionality. In the preceding example, the Authenticator PAE resides on the switch.
- Controlled Port—any switch port with EAPOL-based security enabled.

The Authenticator communicates with the Supplicant using an encapsulation mechanism known as EAP over LANs (EAPOL).

The Authenticator PAE encapsulates the EAP message into a RADIUS packet before sending the packet to the Authentication Server. The Authenticator facilitates the authentication exchanges that occur between the Supplicant and the Authentication Server by encapsulating the EAP message to make it suitable for the packet's destination.

The Authenticator determines the controlled port's operational state. After the RADIUS server notifies the Authenticator PAE about the success or failure of the authentication, it changes the controlled port's operational state accordingly.

The Authenticator PAE functionality is implemented for each controlled port on the switch. At system initialization, or when a supplicant is initially connected to the switch's controlled port, the controlled port's state is set to Blocking. During that time, EAP packets are processed by the authenticator.

When the Authentication server returns a "success" or "failure" message, the controlled port's state is changed accordingly. If the authorization is successful, the controlled port's operational state is set to Forwarding. Otherwise, the controlled port's state depends on the Operational Traffic Control field value in the EAPOL Security Configuration screen.

The Operational Traffic Control field can have one of the following two values:

- Incoming and Outgoing—If the controlled port is unauthorized, frames are not transmitted through the port; all frames received on the controlled port are discarded. The controlled port's state is set to Blocking.
- Incoming—If the controlled port is unauthorized, frames received on the port are discarded, but the transmit frames are forwarded through the port.

EAPOL dynamic VLAN assignment

If EAPOL-based security is enabled on a port, and then the port is authorized, the EAPOL feature dynamically changes the port's VLAN configuration according to preconfigured values, and assigns a new VLAN. The new VLAN configuration values are applied according to previously stored parameters (based on the user_id) in the Authentication server.

The following VLAN configuration values are affected:

- Port membership
- PVID
- Port priority

When the EAPOL-based security is disabled on a port that was previously authorized, the port's VLAN configuration values are restored directly from the switch's non-volatile random access memory (NVRAM).

The following exceptions apply to dynamic VLAN assignments:

- The dynamic VLAN configuration values assigned by EAPOL are **not** stored in the switch's NVRAM.
- If an EAPOL connection is active on a port, any changes to the port membership, PVID, or port priority will not be saved to NVRAM.
- When EAPOL is enabled on a port, and you configure values other than VLAN configuration values, those values are applied and stored in NVRAM.

You set up your Authentication server (RADIUS server) for EAPOL dynamic VLAN assignments. The Authentication server allows you to configure user-specific settings for VLAN memberships and port priority.

When you log on to a system that has been configured for EAPOL authentication, the Authentication server recognizes your user ID and notifies the switch to assign preconfigured (user-specific) VLAN membership and port priorities to the switch. The configuration settings are based on configuration parameters that were customized for your user ID and previously stored on the Authentication server.

To set up the Authentication server, set the following "Return List" attributes for all user configurations (refer to your Authentication server documentation):

- VLAN membership attributes (automatically configures PVID)
  - Tunnel-Type: value 13, Tunnel-Type-VLAN
  - Tunnel-Medium-Type: value 6, Tunnel-Medium-Type-802
  - Tunnel-Private-Group-Id: ASCII value 1 to 4094 (this value is used to identify the specified VLAN)
- Port priority (vendor-specific) attributes
  - Vendor Id: value 562, Nortel vendor Id
  - Attribute Number: value 1, Port Priority
  - Attribute Value: value 0 (zero) to 7 (this value is used to indicate the port priority value assigned to the specified user)

System requirements

The following are minimum system requirements for the EAPOL-based security feature:

- At least one Ethernet Routing Switch 3510-24T
- RADIUS server (Microsoft Windows XP Server)
- Client software that supports EAPOL (Microsoft Windows XP Client)

You must specify the Microsoft 2001 IAS server as the primary RADIUS server for these devices.

EAPOL-based security configuration rules

The following configuration rules apply to your Ethernet Routing Switch 3510-24T when using EAPOL-based security:

- Before configuring your switch, you must configure the Primary RADIUS Server and Shared Secret fields.
- You cannot configure EAPOL-based security on ports that are currently configured for:
  - Shared segments
  - MultiLink Trunking
  - MAC address-based security
  - IGMP (Static Router Ports)
  - Port mirroring

- You can connect only a single client on each port that is configured for EAPOL-based security. (If you attempt to add additional clients to a port, that port goes to Blocking mode.)

RADIUS-based security uses the RADIUS protocol to authenticate local console, Telnet, and EAPOL-authorized logins. Refer to "RADIUS-based Network Security" on page 47 for more information on using the RADIUS protocol.

IP Manager List

You can limit access to the management features of the Ethernet Routing Switch 3510-24T by defining the IP addresses allowed access to the switch. The features provided by the IP manager list are:

- Definitions of up to 50 IP addresses and masks allowed
- Options to enable or disable access for Telnet, SNMP, and the Web-based management system

You cannot change the Telnet access field while connected to the switch through Telnet. Please use a non-Telnet connection to modify the Telnet access field.

Note: To avoid locking a user out of the switch, Nortel recommends that you configure *ranges* of IP addresses that you allow access.

When you configure the access, you are setting access for the *next* session. The current session any user has open is unaffected.

Overview of Security

Table 6 through Table 9 display the types of security available on the Ethernet Routing Switch 3510-24T.

Table 6 EAPOL Security

EAPOL	Description
Description	Extensible Authentication Protocol Over LAN (Ethernet) - set up network access control on internal LANs.
What is being secured	User access to the network

Table 6 EAPOL Security

EAPOL	Description
Per Port or Per Switch	User authentication per port
Layer	Layer 2
Level of Security	Network access encryption
Violations	Switch blocks a port if intruder is seen on that port. Admin has to re-enable port
Requirements for Setup	Radius Server configuration on the switch. EAP-Radius server needs to be accessible from the switch.
Configuring using interfaces	Device Manger (DM), Nortel Command Line (NCLI), Web-based management system.
Restrictions and Limitations	Not allowed—Shared segments and ports configured for MultiLink Trunking, MAC address-based security, IGMP (static router ports), or port mirroring
Reference	IEEE802.1X, RFC 2284
Comments	

Table 7 MAC Security

MAC Security	Description
Description	The MAC address-based security feature allows a user to set up network access control, based on source MAC addresses of authorized stations.
What is being secured	Access to the network or specific subnets or hosts.
Per Port or Per Switch	Per port
Layer	Layer 2
Level of Security	Forwarding
Violations	SA filtering, DA filtering, Port Partitioning, SNMP Trap
Requirements for Setup	Not applicable
Configuring using interfaces	Web, Console, NCLI, ASCII configuration file, SNMP.
Restrictions and Limitations	
Reference	s5sbs103 MIB
Comments	

Table 8 Password Authentication Security

Password Authentication	**Description**
Description	Security feature
What is being secured	User access to a switch
Per Port or Per Switch	For Radius authentication: - Radius server needs to be accessible from switch. - The Radius client from switch should be provisioned with Radius server IP and UDP Port and a shared secret.
Layer	Not applicable
Level of Security	Provides Read Only / Read Write access. The access rights are checked against Local Password / Radius Server.
Violations	Not applicable
Requirements for Setup	For Radius authentication: - Radius server needs to be accessible from switch. - The Radius client from switch should be provisioned with Radius server IP and UDP Port and a shared secret.
Configuring using interfaces	Console, web, NCLI, ASCII configuration file.
Restrictions and Limitations	Not applicable

Table 9 IP Manager Security

IP Manager	**Description**
Description	IP Manager is an extension of Telnet. It provides an option to enable/disable access for TELNET (Telnet On/Off), SNMP (SNMP On/Off) and Web Page Access (Web On/Off) with or without a list of 10 IP Addresses and masks.
What is being secured	User access to the switch via telnet, SNMP, or Web.
Per Port or Per Switch	Per switch
Layer	IP
Level of Security	Access
Violations	User is not allowed to access the switch.
Requirements for Setup	Optional IP Addresses/Masks, Individual Access (enable/disable) for TELNET, SNMP or Web Page
Configuring using interfaces	Web, console, and CLI
Restrictions and Limitations	

MultiLink Trunking

The MultiLink Trunking feature allows you to group multiple ports, two to four together, when forming a link to another switch or server, thus increasing aggregate throughput of the interconnection between two devices, up to 8 GB in full-duplex mode.The Ethernet Routing Switch 3510-24T can be configured with up to six MultiLink Trunks. For more information about the MultiLink Trunking feature, refer to Chapter 2.

Note: For information on configuring VLANs, STGs, and MLTs, refer to Chapter 2.

Port Mirroring

The port mirroring feature (sometimes referred to as *conversation steering*) allows you to designate a single switch port as a traffic monitor for a specified port. You can specify *port-based* monitoring for ingress and egress at a specific port. You can also attach a probe device (**such as a Nortel StackProbe*, or equivalent**) to the designated monitor port.

For more information about the port mirroring feature, refer to Chapter 2.

Note: Use the console interface (CI) menus, the CLI, or the Web-based management system to configure port mirroring.

Autosensing, Autonegotiation, Auto-MDI/X, and Autopolarity

The Ethernet Routing Switch 3510-24T is a autosensing and autonegotiating device:

- The term *autosense* refers to a port's ability to *sense* the speed of an attached device.

- The term *autonegotiation* refers to a standardized protocol (IEEE 802.3u or 802.3z or 802.3ab) that exists between two IEEE -capable devices. Autonegotiation allows the switch to select the best of both speed and duplex modes.
- The term *autopolarity* refers to the ability of the port to compensate for positive and negative signals being reversed on the receive cables.
- The term *auto-MDI/X* refers to automatic detection of transmit and receive twisted pairs.

Autosensing is used when the attached device is not capable of autonegotiation or is using a form of autonegotiation that is not compatible with the IEEE 802.3z autonegotiation standard. In this case, because it is not possible to sense the duplex mode of the attached device, the Ethernet Routing Switch 3510-24T reverts to half-duplex mode.

When autonegotiation-capable devices are attached to the Ethernet Routing Switch 3510-24T, the ports negotiate down from 1000 Mbps speed and full-duplex mode until the attached device acknowledges a supported speed and duplex mode.

Auto-MDI/X detects receive and transmit twisted pairs automatically. When auto-MDI/X is active, any straight or crossover category 5 cable can be used to provide connection to a port. If autonegotiation is disabled, then auto-MDI/X is not active.

The Ethernet Routing Switch 3510-24T features auto-polarity. With autonegotiation enabled, auto-polarity automatically reverses the polarity of a pair of pins from positive to negative or negative to positive. This corrects the polarity of the received data if the port detects that the polarity of the data has been reversed due to a wiring error. If autonegotiation is disabled, auto-polarity is not active.

The Ethernet Routing Switch has four shared front-panel ports: 21, 22, 23 and 24 on the Ethernet Routing Switch 3510-24T. If you insert an SFP GBIC into one on these ports, that port handles gigabit Ethernet speed only. With no optional SFP GIBCs inserted, these ports function as the other 10/100/1000 front-panel ports. The autonegotiation configuration settings are set only once for each port, so these settings are the same for these two sockets that share a port. Autonegotiation is enabled or disabled based on the configuration for these shared ports. However, the speed and duplex settings are ignored on that port once you insert an SFP GBIC.

➡ **Note:** Ensure that you have both sides of the link configured identically when you are using the SFP GBIC, or you may lose connectivity. This applies to all SFP GBIC ports.

For more information about autosensing and autonegotiation modes, see Chapter 6.

Custom Autonegotiation Advertisements

The Custom Autonegotiation Advertisements feature (CANA) allows control of the speed and duplex settings that each ethernet port of the device will advertise as part of the autonegotiation process. Without CANA, a port with autonegotiation enabled will advertise all speed and duplex modes that the switch supports and attempt to establish a link at the highest common speed and duplex setting. Using CANA, the port can be configured to advertise only certain speed and duplex settings, thereby allowing links to be established only at these settings, regardless of the highest common supported operating mode. The CANA feature also allows control over the IEEE802.3x flow control settings advertised by the port, as part of the autonegotiation process. Flow control advertisements can be set to Symmetric, Asymmetric, or Disabled if neither is selected.

In some situations, you may not want a port to advertise all speed and duplex modes supported, as in the following examples:

- If a network can support only 10 Mbps connection, a port can be configured to advertise only 10 Mbps capabilities. Devices using autonegotiation to connect to this port would connect at 10 Mbps, even if both devices are capable of higher speeds.
- If a port is configured to advertise only 100 Mbps full-duplex capability, the link will go active only if the link partner is also capable of autonegotiating a 100 Mbps full duplex connection. This can prevent mismatched speed or duplex settings if autonegotiation is disabled on the link partner.

- For testing or network troubleshooting, it may be useful to configure a link to autonegotiate at a particular speed or duplex mode.

Configuring CANA Using the CLI

This section describes configuring CANA using the CLI and includes the following topics:

- "Configuring CANA", next
- "Viewing Current Autonegotiation Advertisements" on page 61
- "Viewing Hardware Capabilities" on page 62
- "Setting Default Advertisements" on page 64

Configuring CANA

Use the `auto-negotiation-advertisements` command to configure CANA.

To configure port 5 to advertise the operational mode of 10 Mbps and full duplex enter the following command line:

```
auto-negotiation-advertisements port 5 10-full
```

Figure 7 shows sample output for the `auto-negotiation-advertisements` command.

Figure 7 auto-negotiation-advertisements command sample output

```
3510-24T#show auto-negotiation-advertisements port 5
Port Autonegotiation Advertised Capabilities
---- ----------------------------------------------------------------
5    10Full 10Half 100Full 100Half 1000Full                     Pause
```

Viewing Current Autonegotiation Advertisements

To view the autonegotiation advertisements for the device, enter the following command line:

```
show auto-negotiation-advertisements [port <portlist>]
```

Figure 8 and Figure 9 show sample output for the `show auto-negotiation-advertisements` command.

Port 5 has been configured to only advertise an operational mode of 10 Mbps full duplex.

Figure 8 show auto-negotiation-advertisements command sample output

```
3510-24T#show auto-negotiation-advertisements
Port Autonegotiation Advertised Capabilities
---- ----------------------------------------------------------------
1    10Full 10Half 100Full 100Half 1000Full                     Pause
2    10Full 10Half 100Full 100Half 1000Full                     Pause
3    10Full 10Half 100Full 100Half 1000Full                     Pause
4    10Full 10Half 100Full 100Half 1000Full                     Pause
5    10Full 10Half 100Full 100Half 1000Full                     Pause
6    10Full 10Half 100Full 100Half 1000Full                     Pause
7    10Full 10Half 100Full 100Half 1000Full                     Pause
8    10Full 10Half 100Full 100Half 1000Full                     Pause
9    10Full 10Half 100Full 100Half 1000Full                     Pause
10   10Full 10Half 100Full 100Half 1000Full                     Pause
11   10Full 10Half 100Full 100Half 1000Full                     Pause
12   10Full 10Half 100Full 100Half 1000Full                     Pause
13   10Full 10Half 100Full 100Half 1000Full                     Pause
14   10Full 10Half 100Full 100Half 1000Full                     Pause
15   10Full 10Half 100Full 100Half 1000Full                     Pause
16   10Full 10Half 100Full 100Half 1000Full                     Pause
17   10Full 10Half 100Full 100Half 1000Full                     Pause
18   10Full 10Half 100Full 100Half 1000Full                     Pause
19   10Full 10Half 100Full 100Half 1000Full                     Pause
20   10Full 10Half 100Full 100Half 1000Full                     Pause
----More (q=Quit, space/return=Continue)----
```

Figure 9 show auto-negotiation-advertisements command sample output

```
3510-24T#show auto-negotiation-advertisements port 5
Port Autonegotiation Advertised Capabilities
---- ----------------------------------------------------------------
5    10Full 10Half 100Full 100Half 1000Full                     Pause
```

Viewing Hardware Capabilities

To view the operational capabilities of the device, enter the following command line:

```
show auto-negotiation-capabilities [port <portlist>]
```

Figure 10 and Figure 11 shows the sample output for the `show auto-negotiation-capabilities` command.

Figure 10 show auto-negotiation-capabilities command sample output

```
3510-24T#show auto-negotiation-capabilities
Port Autonegotiation Capabilities
---- ----------------------------------------------------------------
1    10Full 10Half 100Full 100Half 1000Full                     Pause
2    10Full 10Half 100Full 100Half 1000Full                     Pause
3    10Full 10Half 100Full 100Half 1000Full                     Pause
4    10Full 10Half 100Full 100Half 1000Full                     Pause
5    10Full 10Half 100Full 100Half 1000Full                     Pause
6    10Full 10Half 100Full 100Half 1000Full                     Pause
7    10Full 10Half 100Full 100Half 1000Full                     Pause
8    10Full 10Half 100Full 100Half 1000Full                     Pause
9    10Full 10Half 100Full 100Half 1000Full                     Pause
10   10Full 10Half 100Full 100Half 1000Full                     Pause
11   10Full 10Half 100Full 100Half 1000Full                     Pause
12   10Full 10Half 100Full 100Half 1000Full                     Pause
13   10Full 10Half 100Full 100Half 1000Full                     Pause
14   10Full 10Half 100Full 100Half 1000Full                     Pause
15   10Full 10Half 100Full 100Half 1000Full                     Pause
16   10Full 10Half 100Full 100Half 1000Full                     Pause
17   10Full 10Half 100Full 100Half 1000Full                     Pause
18   10Full 10Half 100Full 100Half 1000Full                     Pause
19   10Full 10Half 100Full 100Half 1000Full                     Pause
20   10Full 10Half 100Full 100Half 1000Full                     Pause
----More (q=Quit, space/return=Continue)----
```

Figure 11 show auto-negotiation-capabilities command sample output

```
3510-24T#show auto-negotiation-capabilities port 5
Port Autonegotiation Capabilities
---- ----------------------------------------------------------------
5    10Full 10Half 100Full 100Half 1000Full                     Pause
```

Setting Default Advertisements

To set default autonegotiation advertisements for the device, enter the following command line in the interface configuration mode:

```
default auto-negotiation-advertisements [port <portlist>]
```

or

```
no auto-negotiation-advertisements [port <portlist>]
```

To set default advertisements for port 5 of the device, enter the following command line:

```
default auto-negotiation-advertisements port 5
```

or

```
no auto-negotiation-advertisements port 5
```

Figure 12 and Figure 13 shows the sample output from the `default auto-negotiation-advertisements` command.

Figure 12 default auto-negotiation-advertisements command sample output

```
3510-24T(config-if)#default auto-negotiation-advertisements port 5
```

Figure 13 no auto-negotiation-advertisements command sample output

```
3510-24T(config-if)#no auto-negotiation-advertisements port 5
```

Configuring CANA Using the Console/Menu Interface

CANA cannot be configured from the console/menu interface, but the port configuration screen of the console/menu interface will indicate if a port has been configured for CANA by displaying custom in the Autonegotiation column. Port 5 in Figure 14 has been configured for CANA.

Figure 14 Port Configuration Screen

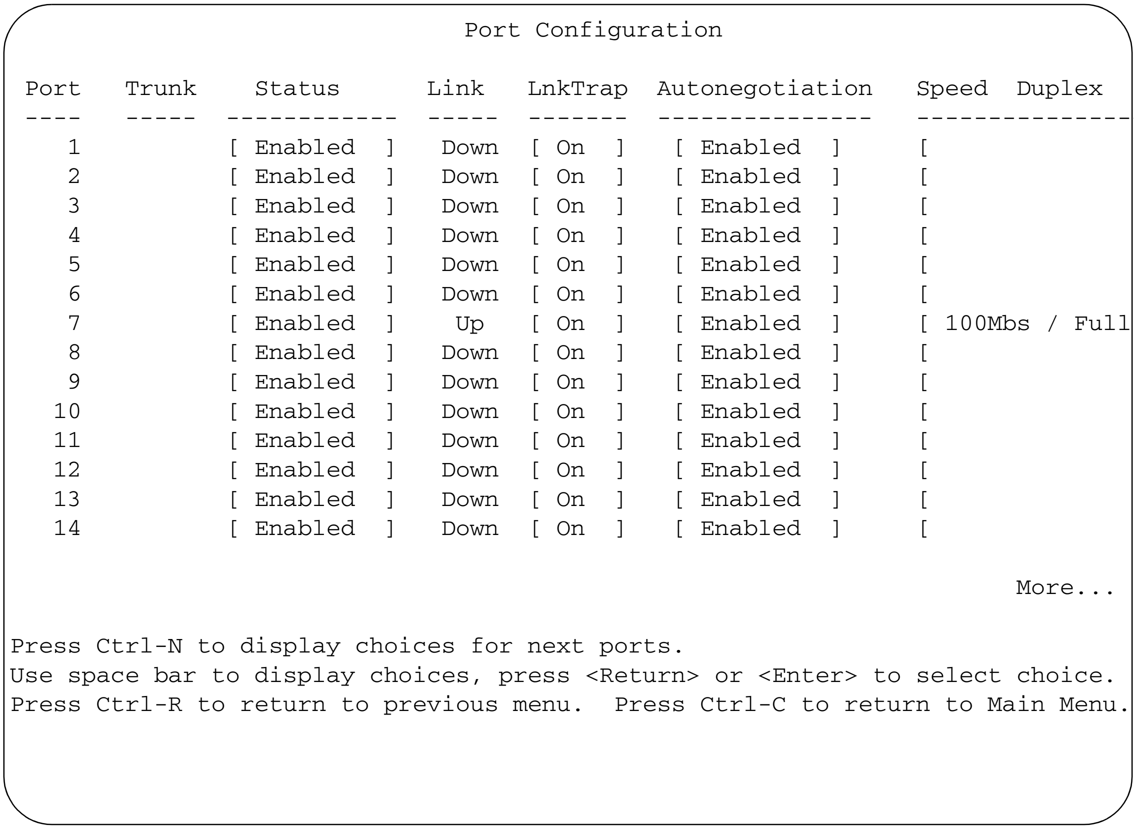

```
                          Port Configuration

 Port    Trunk    Status       Link    LnkTrap  Autonegotiation    Speed  Duplex
 ----    -----  ------------  -----   -------  ---------------    ----------------
    1           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
    2           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
    3           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
    4           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
    5           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
    6           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
    7           [ Enabled  ]    Up    [ On   ]   [ Enabled  ]      [ 100Mbs / Full
    8           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
    9           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
   10           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
   11           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
   12           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
   13           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [
   14           [ Enabled  ]   Down   [ On   ]   [ Enabled  ]      [

                                                                         More...

Press Ctrl-N to display choices for next ports.
Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Autonegotiation can be disabled and enabled for a CANA-configured port using the console/menu interface, but the custom speed and duplex advertisement cannot be changed. You must use the CLI or Device Manager to change the CANA configuration values.

ASCII Configuration File

The Ethernet Routing Switch 3510-24T can download a user-editable ASCII configuration file from a TFTP server. You can load the ASCII configuration file automatically at boot time or on demand using the management systems (console menus or CLI). Once downloaded, the configuration file automatically configures the switch according to the Command Line Interface (CLI) commands in the file. This feature allows the flexibility of generating command configuration files that can be used on several switches with minor modifications. (The maximum size for an ASCII configuration file is 100 KBs; larger configuration files must be split into multiple files.)

Use a text editor to edit the ASCII configuration; the command format is the same as that of the CLI.

You can initiate the ASCII configuration file download using CLI commands only while connected to the base unit, and the ASCII configuration script will execute to completion. When you initiate downloading the ASCII configuration file from the console menu interface, the console does not display output. For this reason, it is important that you review the commands in the file to ensure accuracy and completeness.

For information on setting the parameters for the ASCII configuration file feature, refer to Chapter 3.

Sample ASCII Configuration File

This section shows a sample ASCII configuration file. This file is an example only and shows a basic configuration for a standalone Ethernet Routing Switch 3510-24T that includes Multi-Link Trunking, VLANs, port speed and duplex, and SNMP configurations.

```
! -------------------------------------------------------
! example script to configure different features from CLI
! -------------------------------------------------------
!
enable
configure terminal
!
!
```

```
! ---------------------------------------------------------
! add several MLTs and enable
! ---------------------------------------------------------
mlt 3 name lag3 enable member 13-14
mlt 4 name lag4 enable member 15-16
mlt 5 name lag5 enable member 17-18
!
!
! ---------------------------------------------------------
! add vlans and ports
! ---------------------------------------------------------
!
! create vlan portbased
vlan create 100 name vlan100 type port
!
! add Mlts created above to this VLAN
vlan members add 100 17
!
! create vlan ip protocol based
vlan create 150 name vlan150 type protocol-ipEther2
!
! add ports to this VLAN
! in this case all ports
vlan members add 150 ALL
vlan ports ALL priority 3
!
! igmp
! you could disable proxy on vlan 100
vlan igmp 100 proxy disable
!
! ---------------------------------------------------------
! Examples of changing interface parameters
! ---------------------------------------------------------
! change speed of port 3
interface Fastethernet 3
speed 10
duplex half
exit
!
! change speed of port 4
interface Fastethernet 4
```

```
speed auto
duplex auto
!
!
! -------------------------------------------------------
! SNMP configuration
! -------------------------------------------------------
snmp host 192.168.100.125 private
snmp community private
!
!
exit
end
! -------------------------------------------------------
! Finished
! -------------------------------------------------------
```

Note: To add comments to the ASCII configuration file, add an exclamation point (!) to the beginning of the line.

Default Management System: NCLI or CI menus

You can set the default management interface when you connect to the Ethernet Routing Switch 3510-24T console port or Telnet to the switch either through NCLI or the console interface (CI) menus. This selection is stored in NVRAM.

On system startup, the banner displays and instructs the user to enter Ctrl+Y. After entering these characters, the system will display either the menus or the Nortel Command Line Interface (NCLI) prompt, depending on which is set using this command.

When using the console port, you must logout for the new mode to display. When using Telnet, all subsequent Telnet sessions display the selection.

Port Naming

You can name, or specify a text string for each port. This feature provides easy identification of the connected users.

Note: You must use either the CLI, DM, or the Web-based management system to name ports.

Port Error summary

You can view all ports that have an error. If a particular port has no errors, it will not be displayed. You must use the Web-based management system for this feature.

BootP mode

The BootP mode with the Ethernet Routing Switch 3510-24T are:

- BootP or Last Address modes
- BootP When Needed
- BootP Always
- BootP Disabled

You can retrieve the ASCII configuration file name and configuration server address using BootP. The Ethernet Routing Switch 3510-24T has a unique 48-bit hardware address, or MAC address, that is printed on a label on the back panel. You use this MAC address when you configure the network BootP server to recognize the Ethernet Routing Switch 3510-24T BootP requests. A properly configured BootP server enables the switch to automatically learn its assigned IP address, subnet mask and the IP address of the default router (default gateway).

Defaulting to BootP-when-needed

The BootP default value is BootP-when-needed. This allows a defaulted unit, or a new unit, to power on and get it’s IP configuration from the BootP server.

➡ **Note:** If an IP address is assigned to the device and the BootP process times out, the BootP retains the default mode of BootP-when-needed. However, if the device does not have an assigned IP address and the BootP process times out, the BootP mode automatically changes to BootP disabled. But this change to BootP disabled is not stored, and the BootP reverts to the default value of BootP-when-needed after rebooting the device.

Additionally, the switch could also receive a complete configuration from an ASCII config file, or even a fresh agent download using the “Image If Newer” feature.

When you upgrade, the switch retains the previous BootP value. When you default the switch after the upgrade, the system moves to the default value of BootP-when-needed.

The CLI command “default ip bootp server” will now set the BootP mode to “BootP When Needed.”

For more information and for an example of a BootP configuration file, see Appendix E.

Configuration and Switch Management

You must assign an IP address to the switch depending on the mode of operation. You can set the IP address by using the console port or BootP, which resides on the switch. You can manage the switch using:

- Console/menu interface

The console interface (CI) allows you to configure and manage the switch locally or remotely. Access the CI menus and screens locally through a console terminal attached to your Ethernet Routing Switch 3510-24T, remotely through a dial-up modem connection, or in-band through a Telnet session.

For information about the console/menu interface, refer to Chapter 3

- Web-based management

  You can manage the network from the World Wide Web. Access the Web-based graphical user interface (GUI) through the HTML-based browser located on your network. The GUI allows you to configure, monitor, and maintain your network through Web browsers. You can also download software using the Web.

  For information about Web-based management, refer to *Web Management for Nortel Ethernet Routing Switch 3510-24T* .

- Java-based Device Manager

  Device Manager is a Java-based set of graphical network management applications used to configure and manage a Ethernet Routing Switch 3510-24T.

  For more information, refer to *Switch Management for Nortel Ethernet Routing Switch 3510-24T.*

- Nortel Command Line Interface (NCLI)

  The CLI is used to automate general management and configuration of the Ethernet Routing Switch 3510-24T. Use the CLI through a Telnet connection or through the serial port on the console.

  For complete information on using the CLI, refer to *NCLI Configuration Guide for Nortel Ethernet Routing Switch 3510-24T.*

- Any generic SNMP-based network management software.

  You can use any generic SNMP-based network management software to configure and manage a Ethernet Routing Switch 3510-24T.

SNMP Support

The Ethernet Routing Switch 3510-24T allows you to configure SNMPv3 in Device Manager, Web-based Management, or by using NCLI commands.

The SNMP agent supports exchanges using SNMPv, SNMPv2c and SNMPv3. Support for SNMPv2c introduces a standards-based GetBulk retrieval capability using SNMPv1 communities. SNMPv3 support introduces industrial grade user authentication and message security. This includes MD5.

SNMP MIB Support

The Ethernet Routing Switch 3510-24T supports an SNMP agent with industry-standard MIBs, as well as private MIB extensions, which ensures compatibility with existing network management tools. The switch supports the MIB-II (RFC 1213), Bridge MIB (RFC 1493), and the RMON MIB (RFC 1757), which provide access to detailed management statistics. With SNMP management, you can configure SNMP traps (on individual ports) to generate automatically for conditions such as an unauthorized access attempt or changes in a port's operating status. Table 10 to Table 12 list supported SNMP MIBs.

Table 10 SNMP Standard MIB support

MIB name	RFC	File name
RMON-MIB	2819	rfc2819.mib
RFC1213-MIB	1213	rfc1213.mib
IF-MIB	2863	rfc2863.mib
SNMPv2-MIB	1907	rfc1907.mib
EtherLike-MIB	2665	rfc2665.mib
ENTITY-MIB	2737	rfc2737.mib
BRIDGE-MIB	n/a	draft-ietf-bridge-bridgemib-smiv2-04.txt
P-BRIDGE-MIB	2674	rfc2674-p.mib
Q-BRIDGE-MIB	2674	rfc2674-q.mib
IEEE8021PAE-MIB	n/a	eapol-d10.mib
SMIv2-MIB	2578	rfc2578.mib
SMIv2-TC-MIB	2579	rfc2579.mib
SNMPv2-MIB	3418	rfc3418.mib

Table 10 SNMP Standard MIB support (continued)

MIB name	RFC	File name
SNMP-FRAMEWORK-MIB	3411	rfc3411.mib
SNMP-MPD-MIB	3412	rfc3412.mib
SNMP-NOTIFICATION-MIB	3413	rfc3413-notif.mib
SNMP-TARGET-MIB	3413	rfc3413-tgt.mib
SNMP-USER-BASED-MIB	3414	rfc3414.mib
SNMP-VIEW-BASED-ACM-MIB	3415	rfc3415.mib
SNMP-COMMUNITY-MIB	2576	rfc2576.mib

Table 11 SNMP proprietary MIB support

MIB name	File name
S5-AGENT-MIB	s5age153.mib
S5-CHASSIS.MIB	s5cha135.mib
S5-CHASSIS-TRAP.MIB	s5ctr121.trp
S5-ETHERNET-TRAP.MIB	s5etr113.trp
RAPID-CITY-MIB	xlr30.mib
S5-SWITCH-BAYSECURE-MIB	s5sbs103.mib
BN-IF-EXTENSIONS-MIB	s5ifx103.mib
BN-LOG-MESSAGE-MIB	bnlog002.mib
S5-ETH-MULTISEG-TOPOLOGY-MIB	s5emt104.mib
NTN-QOS-POLICY-EVOL-PIB	pibNtnEvol_003.mib
BAY-STACK-NOTIFICATIONS-MIB	bsn002.mib

Table 12 Application and related MIBs

Application	Related MIBs	File name
RMON-MIB	RMON-MIB	rfc2819.mib
MLT	RAPID-CITY-MIB (rcMlt group)	xlr30.mib
Policy management	NTN-QOS-POLICY-EVOL-PIB	pibNtnEvol_003.mib

Table 12 Application and related MIBs (continued)

Application	Related MIBs	File name
SNMPv3	SNMP-FRAMEWORK-MIB	rfc3411.mib
	SNMP-MPD-MIB	rfc3412.mib
	SNMP-NOTIFICATION-MIB	rfc3413-notif.mib
	SNMP-TARGET-MIB	rfc3413-tgt.mib
	SNMP-USER-BASED-SM-MIB	rfc3414.mib
	SNMP-VIEW-BASED-ACM-MIB	rfc3415.mib
	SNMP-COMMUNITY-MIB	rfc2576.mib
BaySecure	S5-SWITCH-BAYSECURE-MIB	s5bss103.mib
IP multicast (IGMP snooping/ proxy)	RAPID-CITY-MIB (rcVlanIgmp group)	xlr30.mib
System log	BN-LOG-MESSAGE-MIB	bnlog002.mib
Autotopology	S5-ETH-MULTISEG-TOPOLOGY-MIB	s5emt104.mib
VLAN	RAPID-CITY-MIB (rcVlan group)	xlr30.mib
Spanning Tree	BRIDGE-MIB, RAPID-CITY-MIB (rcStg group)	draft-jetf-bridge-bridgemib-smiv2 -04.txt; xlr30.mib
EAPOL	IEEE8021PAE-MIB	eapol-d10.mib
MIB-2	RFC1213-MIB	rfc1213.mib

SNMP Trap Support

The Ethernet Routing Switch 3510-24T supports an SNMP agent with industry-standard SNMPv1 traps, as well as private SNMPv1 trap extensions.

Table 13 describes the SNMP Traps

Table 13 Supported SNMP traps

Trap name	**Configurable**	**Sent when**
RFC 1215 (industry standard):		
linkUp	Per port	A port's link state changes to up.
linkDown	Per port	A port's link state changes to down.
authenticationFailure	System wide	There is an SNMP authentication failure.
coldStart	Always on	The system is powered on.
warmStart	Always on	The system restarts due to a management reset.
s5CtrMIB (Nortel proprietary traps):		
s5CtrProblem	Always on	• Base unit fails • AC power fails or is restored • Fan fails or is restored
s5EtrSbsMacAccessViolation	Always on	A MAC address security violation is detected.
entConfigChange	Always on	Any hardware change— GBIC inserted or removed.
risingAlarm fallingAlarm	Always on	An RMON alarm threshold is crossed.
bsnConfigurationSavedToNvram	Always on	Each time the system configuration is saved to NVRAM.
bsnEapAccessViolation	Always on	An EAP access violation occurs.

Supported Standards and RFCs

This section lists the standards and RFCs supported by the Ethernet Routing Switch 3510-24T.

Standards

The following IEEE Standards contain information related to the Ethernet Routing Switch 3510-24T:

- IEEE 802.1D (Standard for Spanning Tree Protocol)
- IEEE 802.3 (Ethernet)
- IEEE 802.1Q (VLAN Tagging)
- IEEE 802.1p (Prioritizing)
- IEEE 802.1X (EAPOL)
- IEEE 802.3u (Fast Ethernet)
- IEEE 802.3z (Gigabit Ethernet)
- IEEE 802.3ab (Gigabit Ethernet over Copper)
- IEEE 802.3x (Flow Control)—the Ethernet Routing Switch 3510-24T supports asymmetric flow control

RFCs

For more information about networking concepts, protocols, and topologies, consult the following RFCs:

- RFC 791 (IP)
- RFC 894 (IP over Ethernet)
- RFC 792 (ICMP)
- RFC 793 (TCP)
- RFC 783 (TFTP)
- RFC 826 (ARP)
- RFC 768 (UDP)
- RFC 854 (Telnet)
- RFC 951 (BootP)
- RFC 1213 (MIB-II)
- RFC 1493 (Bridge MIB)
- RFC 2863 (Interfaces Group MIB)
- RFC 2665 (Ethernet MIB)
- RFC 2737 (Entity MIBv2)

- RFC 2819 (RMON MIB)
- RFC 1757 (RMON)
- RFC 1271 (RMON)
- RFC 1157 (SNMP)
- RFC 1112 (IGMPv1)
- RFC 2236 (IGMPv2)
- RFC 1945 (HTTP v1.0)
- RFC 2138 (RADIUS)
- RFC 2674 (Q-BRIDGE-MIB)
- RFC 2571 (SNMP Frameworks)
- RFC 2573 (SNMPv3 Applications)
- RFC 2574 (SNMPv3 USM)
- RFC 2575 (SNMPv3 VACM)
- RFC 2572 (SNMP Message Processing)

Chapter 2
Network Configuration

Use the Ethernet Routing Switch 3510-24T to connect workstations, personal computers (PCs), and servers to each other by connecting these devices directly to the switch, through a shared media hub connected to the switch or by creating a virtual LAN (VLAN) through the switch.

This chapter covers the following important information on configuring networks:

- ‘Network Configuration Examples”, next
- “IP Routing” on page 84
- “VLAN Workgroups” on page 102
- “Spanning Tree Protocol Groups” on page 115
- “Internet Group Management Protocol Snooping” on page 120
- “MultiLink Trunks” on page 127
- “Port Mirroring” on page 135

Network Configuration Examples

This section provides four network configuration examples using Ethernet Routing Switch 3510-24T. In these examples, the packet classification feature can be used to prioritize the traffic of the network to ensure uninterrupted traffic of critical applications.

- “High-bandwidth Desktop Switch Configuration”, next
- “High-bandwidth Server Configuration” on page 81
- “OEL2 Aggregation” on page 82
- “Layer 2 Aggregator” on page 83

High-bandwidth Desktop Switch Configuration

Figure 15 shows the Ethernet Routing Switch 3510-24T used as a desktop switch, where desktop workstations are connected directly to Ethernet Routing Switch 3510-24T ports. A Passport 8600 provides high-capacity and low latency connections to the rest of the network. Users can transfer files to and from the network with much greater speed. Configuring a high-bandwidth desktop configuration requires only three major steps:

1. Configure the multi-link transfer (MLT) ports that link to the Passport 8600
2. Configure the MLT ports on the Passport 8600 that attach to the Ethernet Routing Switch 3510-24T.
3. Attach one or more high-speed workstations to the Ethernet Routing Switch 3510-24T.

Figure 15 Ethernet Routing Switch 3510-24T used as a desktop switch

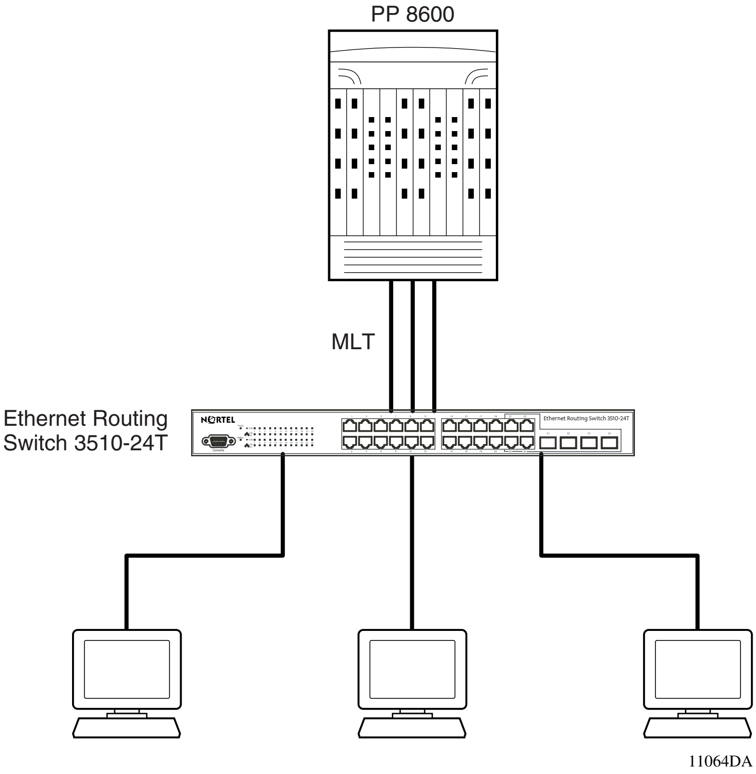
PP 8600
MLT
Ethernet Routing Switch 3510-24T
NORTEL
Ethernet Routing Switch 3510-24T
11064DA

High-bandwidth Server Configuration

Figure 16 shows an example of the Ethernet Routing Switch 3510-24T used to service a group of servers, where the servers are connected directly to Ethernet Routing Switch 3510-24T ports.

A Passport 8600 provides high-capacity and low latency connections to the rest of the network. The Ethernet Routing Switch 3510-24T provides up to four gigabit links for each server, and can balance the high-speed server connections with multi-gigabit links back to the network. The Ethernet Routing Switch 3510-24T also provides configuration of multiple 10/100/1000 Mbps link. This allows for the evolution of connections from a single 10 Mbps connection to a multi-gigabit connection without requiring another switch.

Configuring a high-bandwidth server configuration requires only four major steps:

1. Configure the network servers.
2. Configure the multi-link transfer (MLT) ports on the Ethernet Routing Switch 3510-24T that link to the network servers.
3. Configure the MLT ports that link to the Passport 8600.
4. Configure the MLT ports on the Passport 8600 that attach to the Ethernet Routing Switch 3510-24T.

Figure 16 Ethernet Routing Switch 3510-24T used in a high-bandwidth server configuration

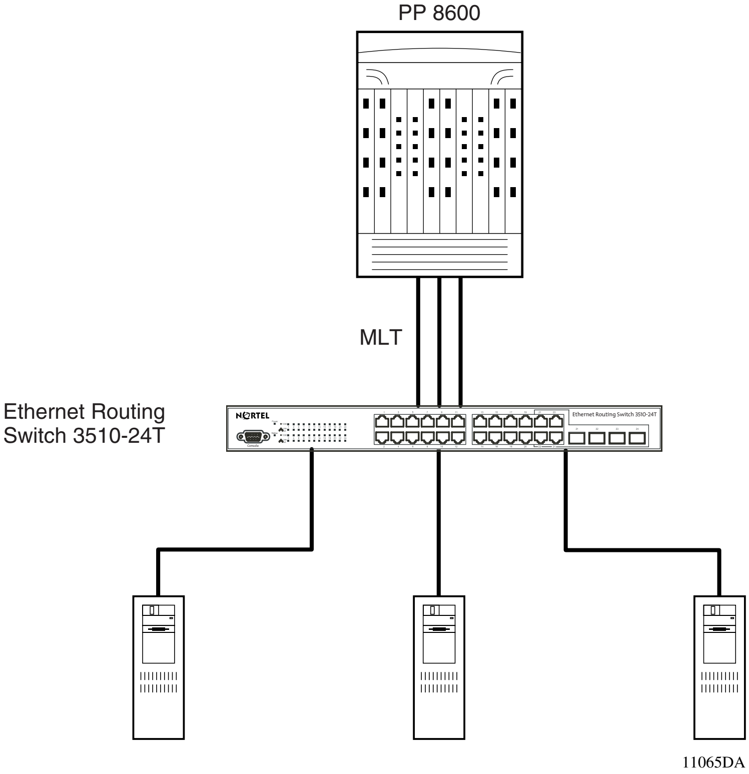
PP 8600
MLT
Ethernet Routing Switch 3510-24T
NORTEL
Ethernet Routing Switch 3510-24T
11065DA

OEL2 Aggregation

Figure 17 shows an example of the Ethernet Routing Switch 3510-24T used to aggregate the uplink connection from OPTera* Metro 1200 Ethernet Service Modules (OM 1200 ESM) at one site to a Passport 8600 at another site. Inexpensive copper connections can be used to connect the OM 1200 ESM units to the Ethernet Routing Switch 3510-24T at one site, while small form factor pluggable gigabit interface connectors (SFP GBICs) connect the Ethernet Routing Switch 3510-24T to the Passport 8600 at the other site.

Configuring the OEL2 aggregation requires four major steps:

1 Configure the OM 1200 ESM units

2 Configure the multi-link transfer (MLT) ports that link the OM 1200 ESM units to the Ethernet Routing Switch 3510-24T.

3 Configure the MLT ports on the Ethernet Routing Switch 3510-24T that link to the Passport 8600.

4 Configure the MLT ports on the Passport 8600 that link to the Ethernet Routing Switch 3510-24T.

Figure 17 Ethernet Routing Switch 3510-24T used in an OEL2 Aggregation

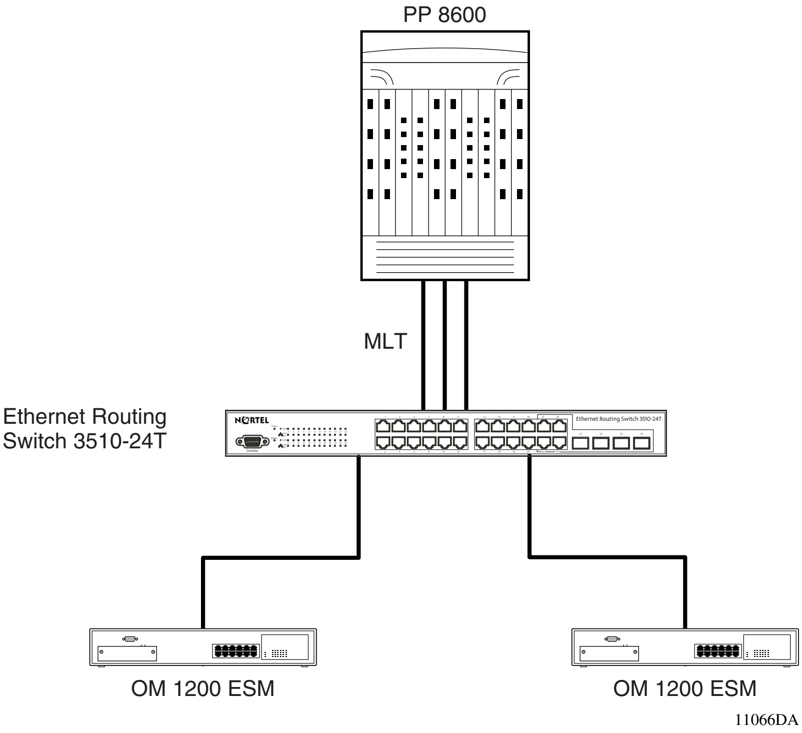

Layer 2 Aggregator

Figure 18 shows an example of the Ethernet Routing Switch 3510-24T used to aggregate the uplink connection from several Business Policy Switch 2000 (BPS 2000) switches to a Passport 8600.

Configuring the Ethernet Routing Switch 3510-24T as a layer 2 aggregator requires three major steps:

1 Attach the BPS 2000 switches to tagged VLAN ports on the Ethernet Routing Switch 3510-24T.

2 Configure the multi-link transfer (MLT) ports on the Ethernet Routing Switch that connect to the Passport 8600.

3 Configure the MLT ports on the Passport 8600 that connect to the Ethernet Routing Switch 3510-24T.

Figure 18 Layer 2 Aggregator

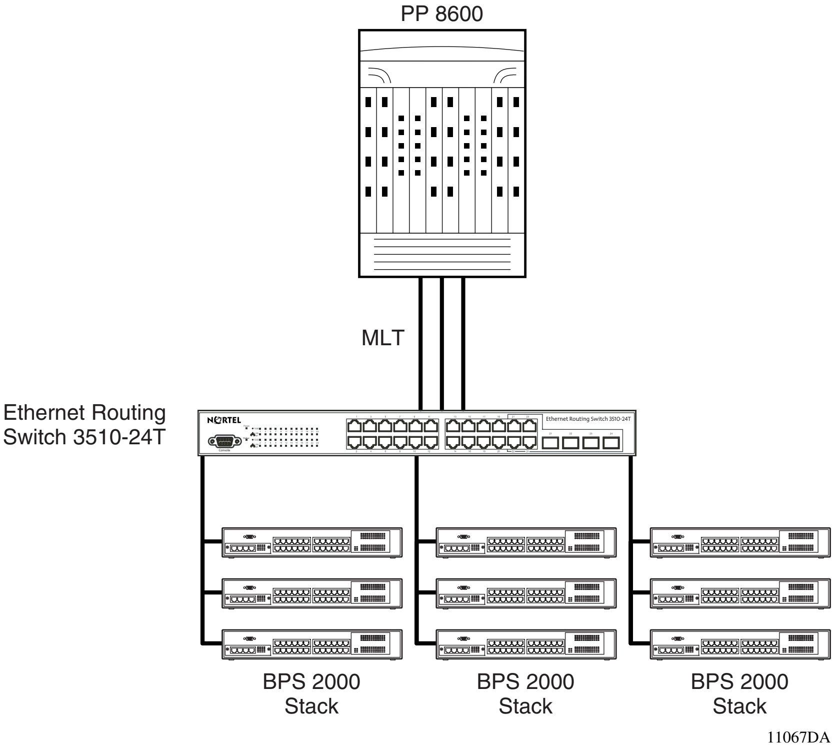

IP Routing

With software release 4.0.3, the Ethernet Routing Switch 3510-24T has introduced IP routing functionality.

To configure IP routing on the Ethernet Routing Switch 3510-24T, use VLANs to create virtual router interfaces by assigning an IP address to the VLAN. This section discusses this concept in depth.

This section covers the following topics:

- "IP Addressing", next
- "IP Routing Using VLANs" on page 88
- "Management VLAN" on page 89
- "CLI commands for the management VLAN static routes" on page 91
- "Setting IP Routing" on page 92
- "Static Routes" on page 92
- "IP Connectivity" on page 92

IP Addressing

An IP version 4 (IPv4) address consists of 32 bits expressed in a "dotted-decimal" format (x.x.x.x). The IPv4 address space is divided into "classes," with classes A, B, and C reserved for unicast addresses and accounting for 87.5 percent of the 32-bit IP address space. Class D is reserved for multicast addressing. Table 14 lists the breakdown of IP address space by address range and mask.

Table 14 IP addresses

Class	Address range	Mask	Number of addresses
A	1.0.0.0 - 126.0.0.0	255.0.0.0	126
B	128.0.0.0 - 191.0.0.0	255.255.0.0	127 * 255
C	192.0.0.0 - 223.0.0.0	255.255.255.0	31 * 255 * 255
D	224.0.0.0 - 239.0.0.0		

To express an IP address in dotted-decimal notation, you convert each octet of the IP address to a decimal number and separate the numbers by decimal points. For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted-decimal notation as 128.32.10.167.

Each IP address class, when expressed in binary, has a different boundary point between the network and host portions of the address as illustrated in Figure 19. The network portion is a network number field from 8 through 24 bits. The remaining 8 through 24 bits identify a specific host on the network.

Figure 19 Network and host boundaries in IP address classes

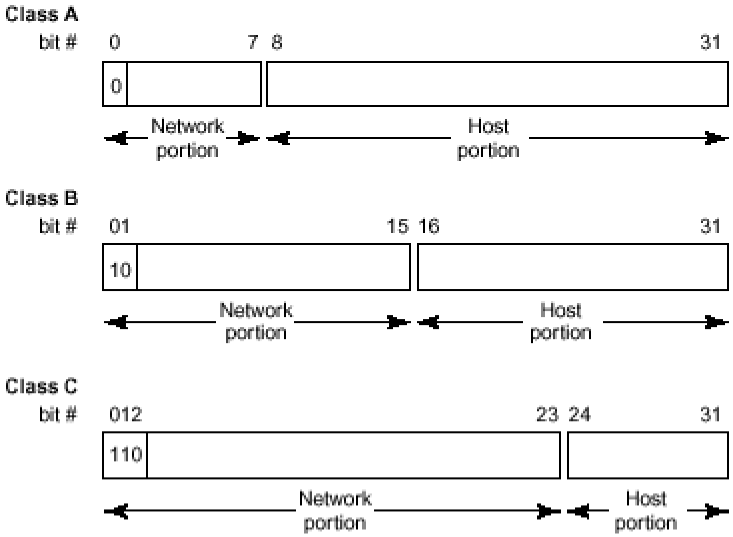

This section includes the following topic:

- “Subnet Addressing”, next

Subnet Addressing

The concept of subnetworks (or subnets) is an extension of the IP addressing scheme. It allows an organization to use one IP address range for multiple networks. Subnets are two or more physical networks that share a common network-identification field (the network portion of the 32-bit IP address).

You can create a subnet address by increasing the network portion to include a subnet address, thus decreasing the host portion of the IP address. For example, in the address 128.32.10.0, the network portion is 128.32, while the subnet is found in the first octet of the host portion (10). A subnet mask is applied to the IP address and identifies the network and host portions of the address.

Table 15 illustrates how subnet masks used with Class B and Class C addresses can create differing numbers of subnets and hosts. This example shows the use of the zero subnet, which is permitted on a Ethernet Routing Switch 3510-24T.

Table 15 Subnet masks for Class B and Class C IP addresses

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts per subnet
Class B			
2	255.255.192.0	2	16,382
3	255.255.224.0	6	8,190
4	255.255.240.0	14	4,094
5	255.255.248.0	30	2,046
6	255.255.252.0	62	1,022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1,022	62
11	255.255.255.224	2,046	30
12	255.255.255.240	4,094	14
13	255.255.255.248	8,190	6
14	255.255.255.252	16,382	2
Class C			
1	255.255.255.128	0	126
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

Variable-length subnet masking (VLSM) is the ability to divide your intranet into pieces that match your requirements. Routing is based on the longest subnet mask or network that matches.

IP Routing Using VLANs

The Ethernet Routing Switch 3510-24T supports wire-speed IP routing between virtual LANs (VLANs). This type of routing is also referred to as virtual routing. When you create a virtual router interface for a specified VLAN, you associate a specific IP address with the specific VLAN. In this release, the Ethernet Routing Switch 3510-24T supports static routing, in which you manually enter the identifiers of the devices you are routing between.

This virtual router interface does not have an association with any specified port or set of ports (it is called a virtual router interface because it is not associated with any particular port). The VLAN IP address can be reached through any of the ports in the VLAN you specified as a virtual router interface, and the assigned IP address is the gateway through which packets are routed out of that VLAN. Routed traffic can be forwarded to another VLAN within the switch.Ethernet Routing Switch 3510-24T.

Once routing is enabled on two VLANs by assigning IP addresses, you can route between those two VLANs (refer to Figure 20).

Figure 20 IP routing with VLANs

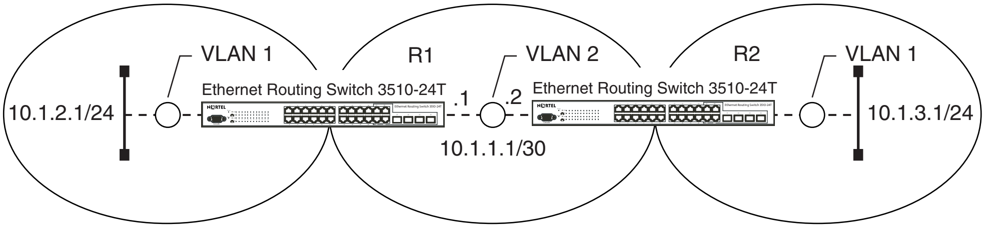

You enable or disable IP routing globally on the Ethernet Routing Switch 3510-24T. The default value is IP routing disabled.

Note: You can configure all the parameters for Ethernet Routing Switch 3510-24T routing prior to actually enabling routing on the switch.

There is no longer a one-to-one correspondence between the physical port and the router interface, because a given port can belong to multiple VLANs. The VLANs may be configured for routing on the switch.

As with any IP address, virtual router interface addresses are also used for device management. For management over IP, you can use any virtual router interface IP address to access the switch as long as routing is enabled. When you use the Ethernet Routing Switch 3510-24T *without* routing enabled, the management VLAN is reachable only through the switch IP address. With IP routing enabled on the switch, any of the virtual router IP interfaces can be used for management over IP.

Once you enable routing on the Ethernet Routing Switch 3510-24T, the management VLAN behaves like all the other routable VLANs (including following the rules to avoid duplicate IP addresses, outlined below). The IP address is reachable through any virtual router interface, as long as a route is available. All virtual router interfaces can be used as the management VLAN over IP.

This feature uses the following MIB:

- rcIpAddrTable

Management VLAN

In the previous releases,the Management VLAN is the only VLAN that can be used to carry the management traffic, including telnet, web, SNMP, BootP, and TFTP for the switch. The Management VLAN always exists on the switch and cannot be removed. All IP settings, including switch IP address,subnet mask, and default gateway apply only to the Management VLAN.

In this release, all VLANs that are virtual routing IP interfaces, are capable of carrying Management traffic when IP routing is enabled. The Management VLAN is said to be in L2 mode if routing is disabled. When routing is enabled, the Management VLAN is said to be in L3 mode.

Default Gateway

When IP routing is first enabled after the device is booted with factory defaults, the default gateway is added to the Static Route table as the default route, if no other default route is present.

Management VLAN Static Routes

Ethernet Routing Switch 3510-24T supports up to four static routes on the management VLAN. These routes can be used if the management station cannot be reached via the management VLAN. For example, the routes can be used in the following situations:

- when the management station is not directly attached to the management VLAN, or
- the default router is not reachable from the management VLAN

L2 Mode vs. L3 Mode

When the Ethernet Routing Switch 3510-24T is configured to route IP traffic between different VLANs, the switch is considered running in L3 mode, otherwise the switch is running in L2 mode.

The L3 manager determines in which mode the switch should be running. This is determined based on the user settings and events. But the general rule is to select:

- L3 mode: if routing is turned on globally for the switch.
- L2 mode: if routing is turned off globally for the switch.

Routing and Management

In L3 mode, the Management VLAN as well as all other L3 VLANs have the capability to route and carry the management traffic. In this release of Ethernet Routing Switch 3510-24T, the settings apply to all L3 VLANs or only to the Management VLAN. Table 16 on page 91, shows all possible settings and default settings for each type of VLAN.

Table 16 VLAN settings

VLAN \ Feature	**Routing (default)**	**Management (default)**	**Default Route**
Management VLAN (L2 mode)	Off	On	Y (Management VLAN only)
Management VLAN (L3 mode)	On/off (on)	On	N
L3 VLAN	On/off (on)	On/off (on)	Y (global)

CLI commands for the management VLAN static routes

Table 17 on page 91, describes the CLI commands for the management VLAN static routes.

Table 17 CLI commands for the management VLAN static routes

Command	**Definition**
show ip mgmt route	To display all configured (up to 4) static routes for the management VLAN
< destination IP > < destination subnet mask > < gateway IP>	To configure a static route for the management VLAN: ip mgmt route
< destination IP > < destination subnet mask > < gateway IP>	To remove a static route from the management VLAN: no ip mgmt route

Avoiding Duplicate IP Addresses

The Ethernet Routing Switch 3510-24T has built-in safeguards that block the issuing of duplicate IP addresses.

The system allows the use of an existing IP address under the following condition:

- When unit boots up: If the unit was the designated Base Unit (BU)-that is, selected by hardware switch on the unit, either on the back or UI button on the front. If you want the switch to use the IP address that the system is blocking, then you will have to specify this to the system by issuing an CLI command.

Setting IP Routing

To set IP routing (or L3 VLANs), take the following steps:

1. Enable IP routing globally.
2. Enable IP routing for the specific VLAN.
3. Assign an IP address to the specific VLAN.

You can use any of the Ethernet Routing Switch 3510-24T management systems to set the IP routes: the Web-based management system, or the Device Manager (DM), or the Nortel Command Line Interface (NCLI). Refer to the document on each management system for detailed instructions on configuring IP routes using that specific system.

Static Routes

Once you create routable VLANs by assigning IP addresses to the VLAN, you can set up static routes, which allow you to manually create specific routes to a destination IP address. Additionally, you can use a static default route to specify a route to all networks for which there are no explicit routes in the Forwarding Information Base or the routing table. This static default route is, by definition, a route to the network address 0.0.0.0, which conforms to IEEE RFC 1812.

IP Connectivity

The following sections describe various protocols used for enhanced and resilient IP connectivity on the Ethernet Routing Switch 3510-24T. The following topics are covered:

- “Address Resolution Protocol”, next
- “DHCP/BootP Relay” on page 94

Address Resolution Protocol

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In the situation where the station knows only the network host's IP address, the Address Resolution Protocol (ARP) enables the network station to determine a network host's physical address and bind the 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts.

If a network station wants to send a packet to a host but knows only the host's IP address, the network station uses ARP to determine the host's physical address as follows:

1. The network station broadcasts a special packet, called an ARP request, that asks the host at the specified IP address to respond with its physical address.
2. All network hosts receive the broadcast message.
3. Only the specified host responds with its hardware address.
4. The network station then maps the host's IP address to its physical address and saves the results in an address resolution table for future use.
5. The network station's ARP table displays the association of the known MAC addresses to IP addresses.

Note: The default timeout value for ARP entries is 6 hours.

You can create static ARP entries and delete individual ARP entries.

This feature uses the following MIB:

- ipNetToMediaTable

DHCP/BootP Relay

The Dynamic Host Configuration Protocol (DHCP) is an extension of the Bootstrap Protocol (BootP) and provides host configuration information to the workstations on a dynamic basis. To lower administrative overhead, network managers prefer to configure a small number of DHCP servers in a central location. It is necessary for routers to support the BootP/DHCP relay function so that hosts can access configuration information from servers several router hops away.

Differences between DHCP and BootP

The following differences between DHCP and BootP are specified in RFC 2131 and include functions that BootP does not address:

- DHCP defines mechanisms through which clients can be assigned a network address for a finite lease (allowing for reuse of IP addresses).
- DHCP provides the mechanism for clients to acquire all of the IP configuration parameters needed to operate.

DHCP uses the BootP message format defined in RFC 951. The remainder of the options field consists of a list of tagged parameters that are called “options” (RFC 2131).

Summary of DHCP Relay Operation

BootP/DHCP clients (workstations) generally use UDP/IP broadcasts to determine their IP addresses and configuration information. If such a host is on a network or a subnet segment (or VLAN) that does not include a DHCP server, the UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. The Ethernet Routing Switch 3510-24T can be configured to overcome this issue by forwarding the broadcasts to the server. The router interfaces can be configured to forward DHCP broadcasts to other locally connected network segments or directly to the server’s IP address. DHCP must be enabled on a per VLAN basis.

Figure 21 shows an end station connected to subnet 1, corresponding to VLAN 1. The Ethernet Routing Switch 3510-24T connects two subnets by means of the virtual routing function. When the end station generates a DHCP request as a limited UDP broadcast to the IP address of all 1s (that is, 255.255.255.255) with the DHCP relay function configured, the Ethernet Routing Switch 3510-24T forwards DHCP requests to subnet 2 or to the host address of the DHCP server, depending on the configuration.

Figure 21 DHCP operation

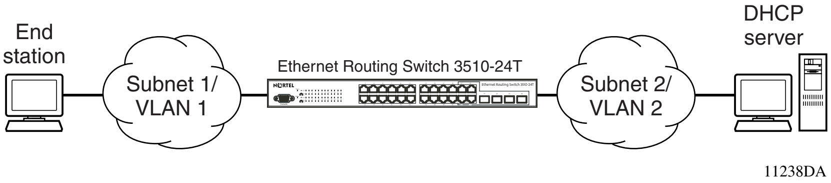

Forwarding DHCP Packets

In the example shown in Figure 22, the *agent address* is 10.10.1.254. To configure the Ethernet Routing Switch 3510-24T to forward DHCP packets from the end station to the server, use 10.10.2.1 as the *server address*.

Figure 22 Forwarding DHCP packets

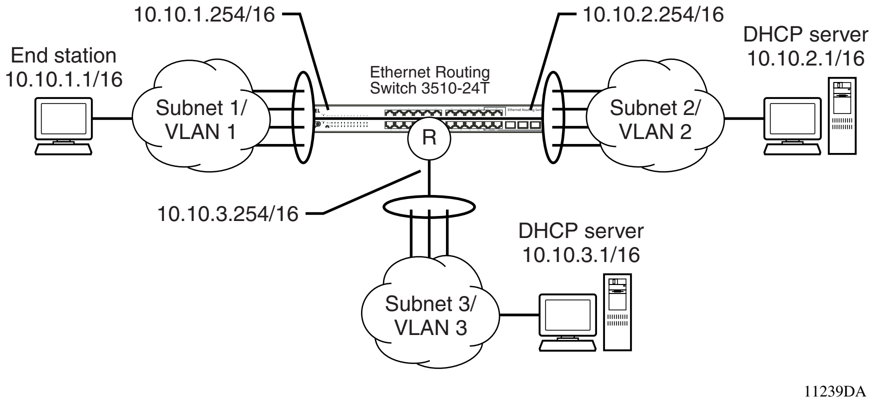

All BootP broadcast packets, including DHCP packets that appear on the VLAN 1 router interface (10.10.1.254), will be forwarded to the DHCP server. In this case, the DHCP packets will be forwarded as unicast to the DHCP server's IP address.

To forward BootP/DHCP packets as broadcast packets to VLAN 2, specify the IP address of the switch VLAN 2 router interface (10.10.2.254) as the server address.

Multiple BootP/DHCP Servers

Most enterprise networks use multiple BootP/DHCP servers for fault tolerance. The Ethernet Routing Switch 3510-24T allows you to configure the switch to forward BootP/DHCP requests to multiple servers. You can configure up to 10 servers to receive copies of the forwarded BootP/DHCP messages.

If a DHCP client is connected to a routable interface, to configure DHCP requests to be sent to 10 different routable interfaces or 10 different server IP addresses, enable DHCP on the client (agent address) and then enable DHCP from the client to each of the interfaces or IP addresses (server addresses).

In the example shown in Figure 23, two DHCP servers are located on two different subnets. To configure the Ethernet Routing Switch 3510-24T to forward the copies of the BootP/DHCP packets from the end station to both servers, specify the switch (10.10.1.254) as the agent address. Then enable DHCP to each of the DHCP servers by entering 10.10.2.1 and 10.10.3.1 as the server addresses.

Figure 23 Configuring multiple BootP/DHCP servers

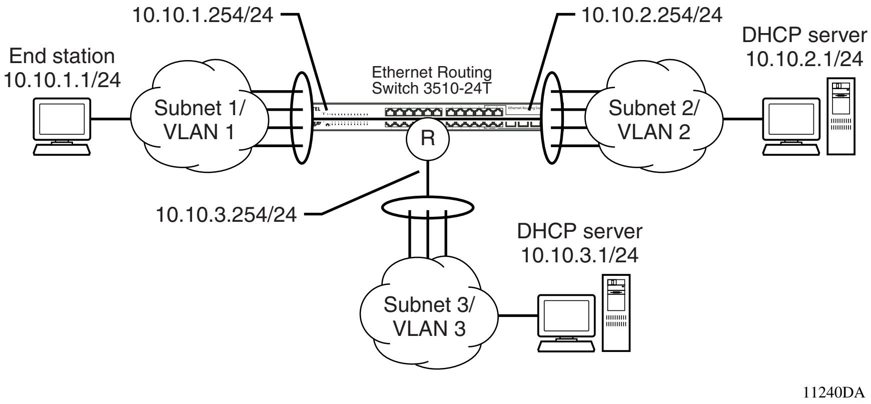

Note: You must configure DHCP both globally on the switch as well as on each virtual router interface VLAN.

Setting DHCP

To set DHCP, take the following steps:

1 Enable IP routing on the Ethernet Routing Switch 3510-24T and on the target VLAN interface.

2 Enable DHCP globally.

Note: DHCP is enabled by default.

3 Set the DHCP forwarding paths, using the VLAN IP as the starting point, or agent IP.

4 Set the mode for each DHCP forwarding path.

5 Enable DHCP for the specific VLAN.

6 Enable the DHCP broadcast message for the specific VLAN.

You use any of the Ethernet Routing Switch 3510-24T management systems to set DHCP: the Web-based management system, or the Device Manager (DM), or the Nortel Command Line Interface (NCLI). Refer to the document on each management system for detailed instructions on configuring IP routes using that specific system.

Dynamic Host Configuration Protocol Relay

Dynamic Host Configuration Protocol (DHCP) is a mechanism to assign network IP address to clients who request an address. It is built on top of the existing BOOTP protocol and can be specified for DHCP, BOOTP, or both.

The DHCP relay feature relays client requests to DHCP servers on different L3 VLANs. It also relays server replies back to the clients.

DHCP relay can be configured through CLI, SNMP and JDM. DHCP can only be configured on the base unit from CLI, like all L3 commands. There are three parts in the DHCP relay configurations. They are:

- global DHCP enable/disable
- interface configurations
- forward path configurations

To relay DHCP messages, you must create two L3 VLANs and assign IP addresses to them. The client and server must reside on different L3 VLANS to use DHCP relay. IP routing and global DHCP relay must be enabled on both the client as well as server.

Global DHCP relay Configuration:

This configuration enables or disables DHCP relay for the entire unit. Once DHCP relay is disabled, the switch will not relay DHCP/BOOTP across L3 VLANs. However the settings will still be configurable.

CLI commands

Table 18 describes the DHCP relay commands.

Table 18 DHCP relay commands

Command	Definition
3510-24T (config)# show ip dhcp-relay	shows global DHCP relay state
3510-24T (config)#no ip dhcp-relay	disables DHCP relay globally
3510-24T (config)# ip dhcp-relay	enables DHCP relay globally

Interface DHCP relay Configurations:

These configurations are associated with the L3 VLAN that the client or server resides. IP routing must be enabled and a valid IP address must be assigned to the L3 VLAN, before it generates the default settings for DHCP relay.

Table 19 describes the interface DHCP relay commands. To change the interface DHCP relay configurations you must be in the interface config prompt 3510-24T(config-if)#.

Table 19 Interface DHCP relay commands

Command	Definition
3510-24T (config-if)# show vlan dhcp-relay	shows vlan dhcp relay state
3510-24T (config-if)# ip dhcp-relay max-hop 16	sets max-hop to 16
3510-24T (config-if)# ip dhcp-relay min-sec 30	sets min-sec to 30
3510-24T (config-if)# ip dhcp-relay mode dhcp	sets mode to dhcp
3510-24T (config-if)# no ip dhcp-relay	disables ip dhcp-relay
3510-24T (config-if)# ip dhcp-relay broadcast	enables broadcast for this interface

Figure 24 shows the output of the show vlan dhcp-relay command.

Figure 24 show vlan dhcp-relaycommand output

```
3510-24T(config)#show vlan dhcp-relay

IfIndex   MIN_SEC  ENABLED  MODE  ALWAYS_BROADCAST
---------------------------------------------------
```

Figure 20 describes the parameters and variables for the ip dhcp relay command.

Table 20 Interface DHCP relay parameters

Parameter	**Definition**
IfIndex	Interface index of the vlan.
max-hop	The maximum number of hops before the DHCP message times out. Default = 4, range 1-16.
min-sec	Minimum number of seconds to wait before forwarding dhcp message. Default = 0 (no wait), range 0-65535.
enable	Enable/Disable relay for this interface. Default = TRUE, range TRUE/ FALSE
mode	bootp: Relay bootp messages only for this interface. dhcp: Relay DHCP messages only. both: Relay both bootp and DHCP messages for this interface. Default = both, range bootp/dhcp/both.
broadcast	Always broadcast the dhcp message when relaying it to the server of client for this interface. Default = Disabled (Unicast), range Enable/ Disable.

DHCP relay forward path configurations:

These configurations are made per interface IP address and server IP address. Table 21 describes the ip dhcp relay fwd-path command.

Table 21 dhcp relay fwd path commands

Command	**Definition**
3510-24T (config)# show ip dhcp-relay fwd-path	shows ip dhcp-relay fwd-path
3510-24T (config)# ip dhcp-relay fwd-path <agent IP> <server IP> mode bootp-dhcp	creates interface IP and server IP path with mode DHCP & BootP
3510-24T (config)# ip dhcp-relay fwd-path <agent IP> <server IP> disable	disables the interface/server pair, enable = false
3510-24T (config)# no dhcp-relay fwd-path <agent IP> <server IP>	deletes the interface/server pair
3510-24T (config-if)# no ip dhcp-relay	disables ip dhcp-relay
3510-24T (config-if)# ip dhcp-relay broadcast	enables broadcast for this interface

Figure 25 shows the output of the show ip dhcp-relay fwd-path command.

Figure 25 show ip dhcp-relay fwd-path command

```
----------------------------------------------------
3510-24T(config)#show ip dhcp-relay fwd-path

=============================================================
                     DHCP
=============================================================
INTERFACE         SERVER             ENABLE     MODE
-------------------------------------------------------------
```

Table 22 describes the parameters and variables for the dhpc-relay fwd-path command.

Table 22 DHCP relay forward path configuration

Parameter	Definition
interface	IP address of the interface.
server	IP address of the server. One interface can have multiple DHCP servers.
enable	Enable/Disable relay between the interface and server. Default = True, range True/False.
mode	Same as Interface DHCP relay configuration mode except it applies to this interface/server pair not for entire interface. Default = DHCP & BOOTP, range DHCP, BOOTP, DHCP & BOOTP.

DHCP relay can be managed from SNMP and JDM. The MIB that DHCP relay uses is rcIpConfDhcpTable and rcIpDhcpForwardTable.

DHCP relay uses a hardware resource that is shared by QoS. When DHCP relay is enabled globally, QoS filter manager will not be able to use precedence 11 for configurations. For filter manager to be able to use this resource, DHCP relay must be disabled for the entire unit.

VLAN Workgroups

The Ethernet Routing Switch 3510-24T support up to 256 VLANs.

Ports are grouped into broadcast domains by assigning them to the same VLAN. Frames received in one VLAN can only be forwarded within that VLAN, and multicast frames and unknown unicast frames are flooded only to ports in the same VLAN.

Setting up virtual LANs (VLANs) is a way to segment networks to increase network capacity and performance without changing the physical network topology (Figure 26). With network segmentation, each switch port connects to a segment that is a single broadcast domain. When a switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belong to one broadcast domain.

The Ethernet Routing Switch 3510-24T allows you to assign ports to VLANs using the console interface (CI) menus, Telnet, Web-based management, CLI, or an appropriate SNMP-based application, such as the Device Manager. You can assign different ports (and therefore the devices attached to these ports) to different broadcast domains. This feature allows network flexibility because you can reassign VLANs to accommodate network moves, additions, and changes, eliminating the need to change physical cabling.

Figure 26 Port-based VLAN example

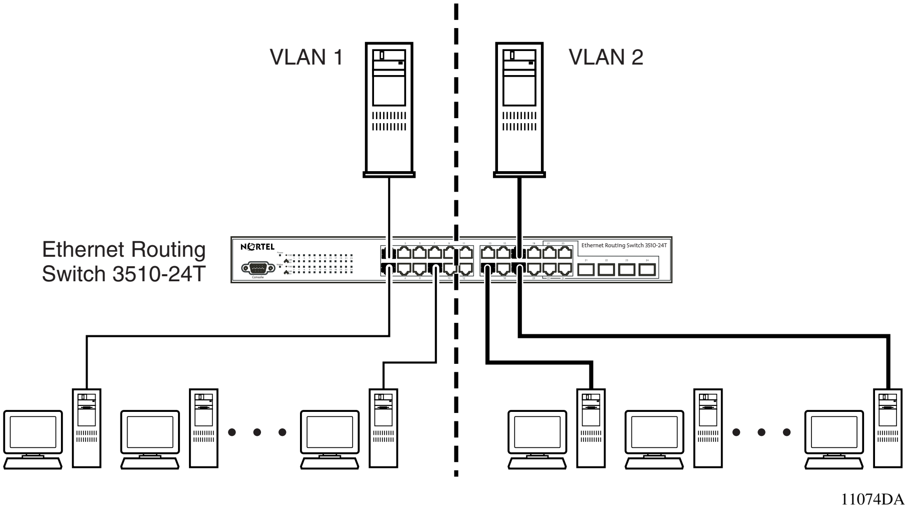
VLAN 1
VLAN 2
Ethernet Routing Switch 3510-24T
NORTEL
Ethernet Routing Switch 3510-24T
11074DA

IEEE 802.1Q Tagging

The Ethernet Routing Switch 3510-24T operates in accordance with the IEEE 802.1Q tagging rules. Important terms used with the 32-bit 802.1Q tagging feature are:

- VLAN identifier (VID)—the 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN. When other types of VLANs are enabled, this default value can be overridden by the values enabled in the management interfaces.

Note: You cannot use a value for a VID that is already used for tagged BPDUs in multiple Spanning Tree Groups (STGs).

- Port VLAN identifier (PVID)—a classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID of 3 (PVID =3) assigns all untagged frames received on this port to VLAN 3.
- Tagged frame—a frame that contains the 32-bit 802.1q field (VLAN tag). this field identifies the frame as belonging to a specific VLAN.
- Untagged frame— a frame that does not carry any VLAN tagging information in the frame header.
- VLAN port members—a group of ports that are all members of a particular VLAN. A port can be a member of one or more VLANs.
- Untagged member—a port that has been configured as an untagged member of a specific VLAN. When an untagged frame exits the switch through an untagged member port, the frame header remains unchanged. When a tagged frame exits the switch through an untagged member port, the tag is stripped and the tagged frame is changed to an untagged frame.
- Tagged member—a port that has been configured as a tagged member of a specific VLAN. When an untagged frame exits the switch through a tagged member port, the frame header is modified to include the 32-bit tag associated with the ingress port PVID. When a tagged frame exits the switch through a tagged member port, the frame header remains unchanged (original VID remains).
- User priority—a three-bit field in the header of a tagged frame. The field is interpreted as a binary number, therefore has a value of 0 - 7. This field allows the tagged frame to carry the user-priority across bridged LANs where the individual LAN segments may be unable to signal priority information.
- Port priority—the priority level assigned to *untagged* frames received on a port. This value becomes the user priority for the frame. *Tagged* packets get their user priority from the value contained in the 32-bit 802.1Q frame header.
- Unregistered packet—a tagged frame that contains a VID where the receiving port is not a member of that VLAN.
- Filtering database identifier (FID)—the specific filtering/forwarding database within the Ethernet Routing Switch 3510-24T that is assigned to each VLAN. Each VLAN has its own filtering database, which is called independent VLAN learning (IVL). IVLs can have duplicate MAC addresses in different VLANs.

The default configuration settings for the Ethernet Routing Switch 3510-24T have all ports set as untagged members of VLAN 1 with all ports configured as PVID = 1. Every VLAN is assigned a unique VLAN identifier (VID) that distinguishes it from all other VLANs. In the default configuration example shown in Figure 27, all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID =1). Untagged packets enter and leave the switch unchanged.

Figure 27 Default VLAN settings

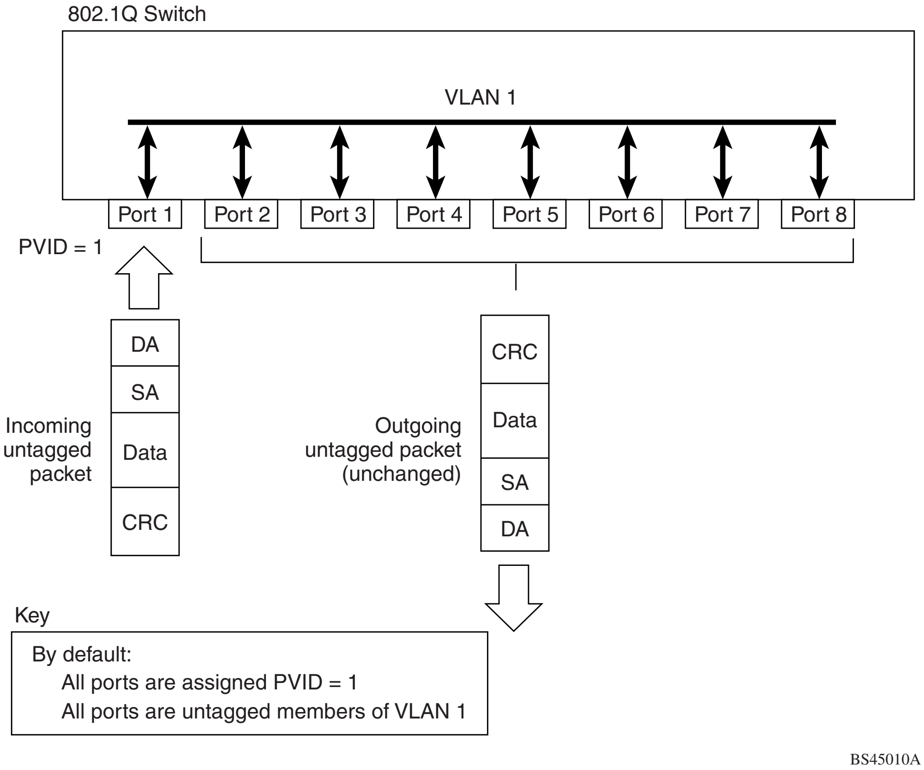

You can configure switch ports to transmit frames tagged on some VLANs, and untagged on other VLANs.

When you configure VLANs, you configure the egress tagging of each switch port as *Untag All, Untag PVID Only, Tag All* or *Tag PVID Only* (see Figure 28 through Figure 35).

In Figure 28, untagged incoming packets are assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

Figure 28 Port-based VLAN assignment

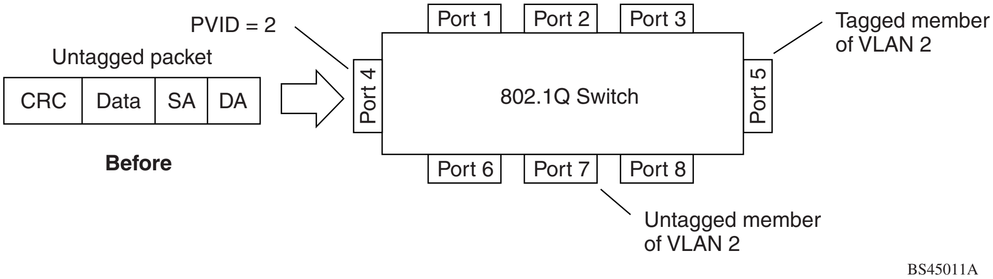

As shown in Figure 29, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 29 802.1Q tagging (after port-based VLAN assignment)

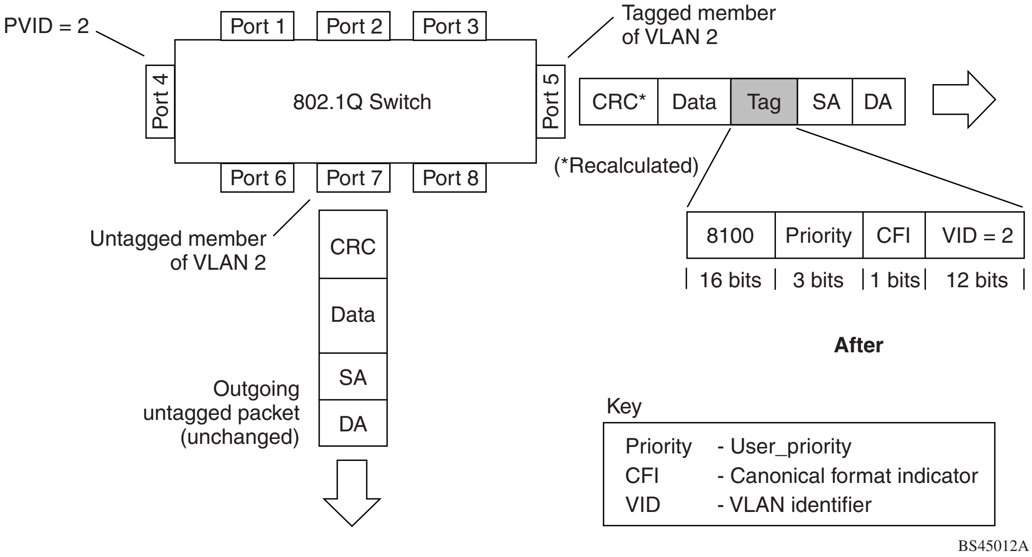

In Figure 30, untagged incoming packets are assigned to VLAN 3 (policy VLAN = 3, PVID = 2). Port 5 is configured as a *tagged* member of VLAN 3, and port 7 is configured as an *untagged* member of VLAN 3.

Figure 30 Policy-based VLAN assignment

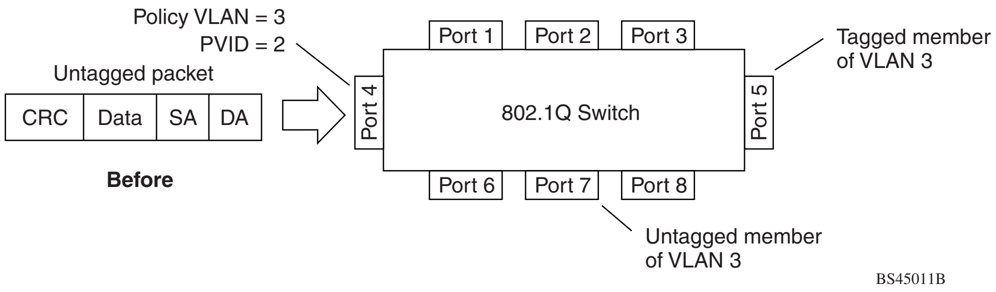

As shown in Figure 31, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of VLAN 3. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of VLAN 3.

Figure 31 802.1Q tagging (after policy-based VLAN assignment)

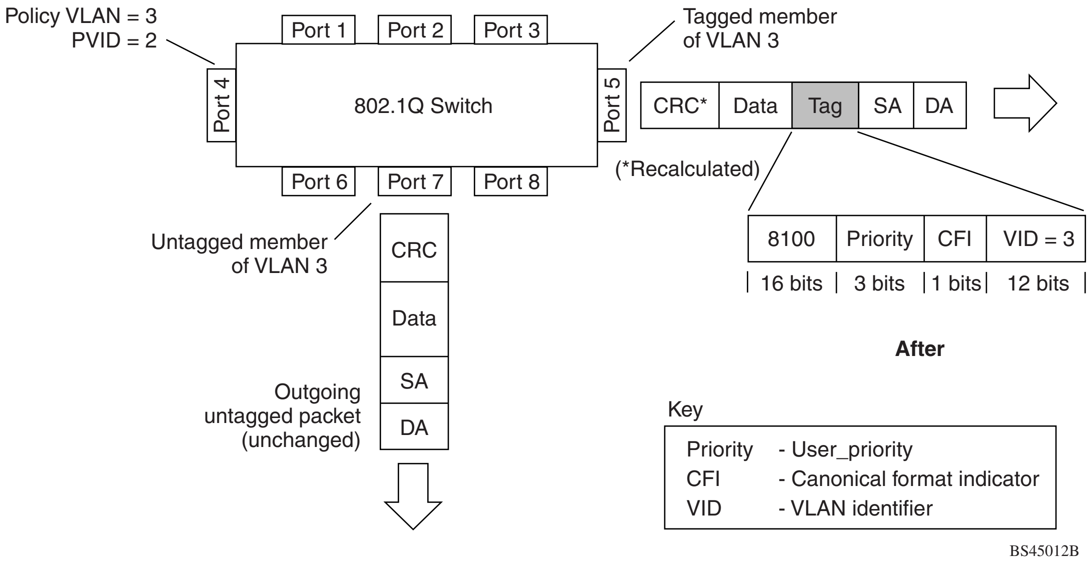

In Figure 32, tagged incoming packets are assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

Figure 32 802.1Q tag assignment

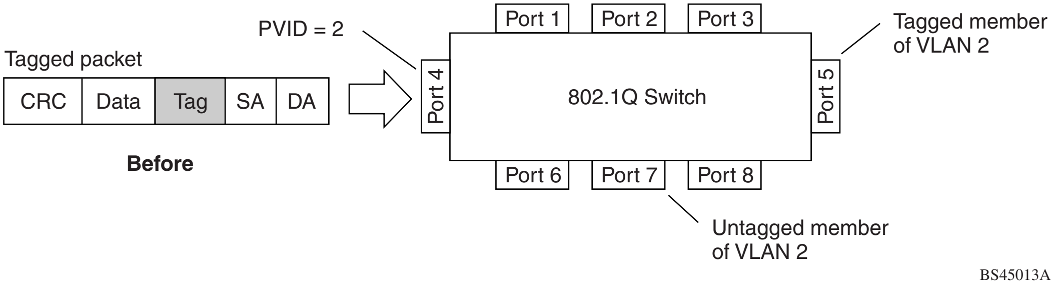

As shown in Figure 33, the tagged packet remains unchanged as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 33 802.1Q tagging (after 32-bit 802.1Q tag assignment)

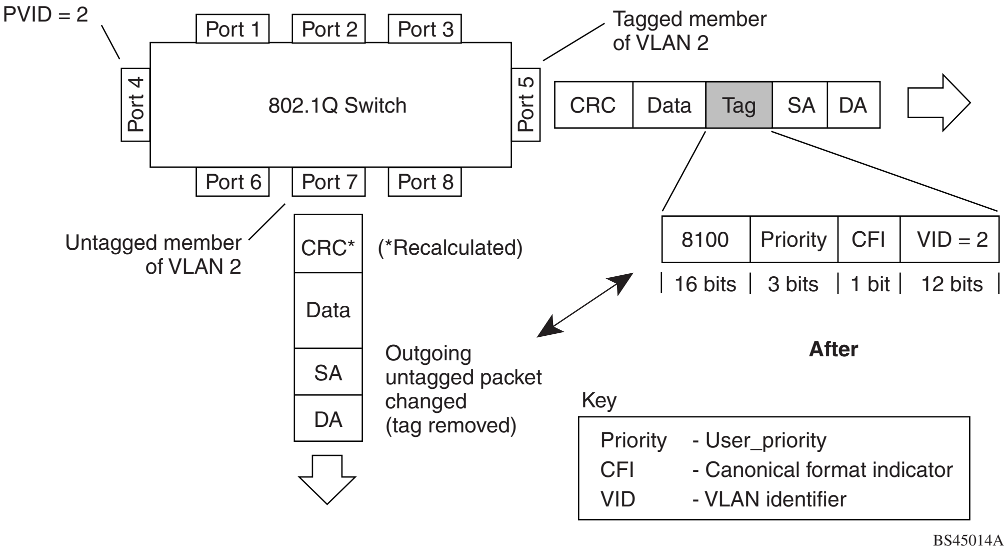

In Figure 34, untagged incoming packets are assigned directly to PVID = 2. Port 5 is configured as a *tagged* member of PVID 2, and port 7 is configured as an *untagged* member of PVID 2.

Figure 34 802.1Q tag assignment

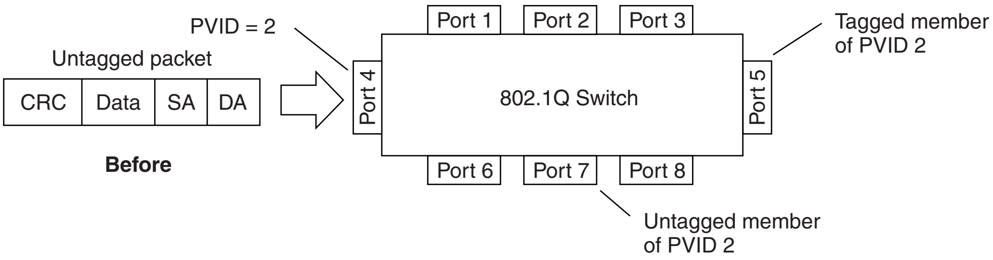

BS45011C

As shown in Figure 35, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of PVID 2. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of PVID 2.

Figure 35 802.1Q tagging (after 30-bit 802.1Q tag assignment)

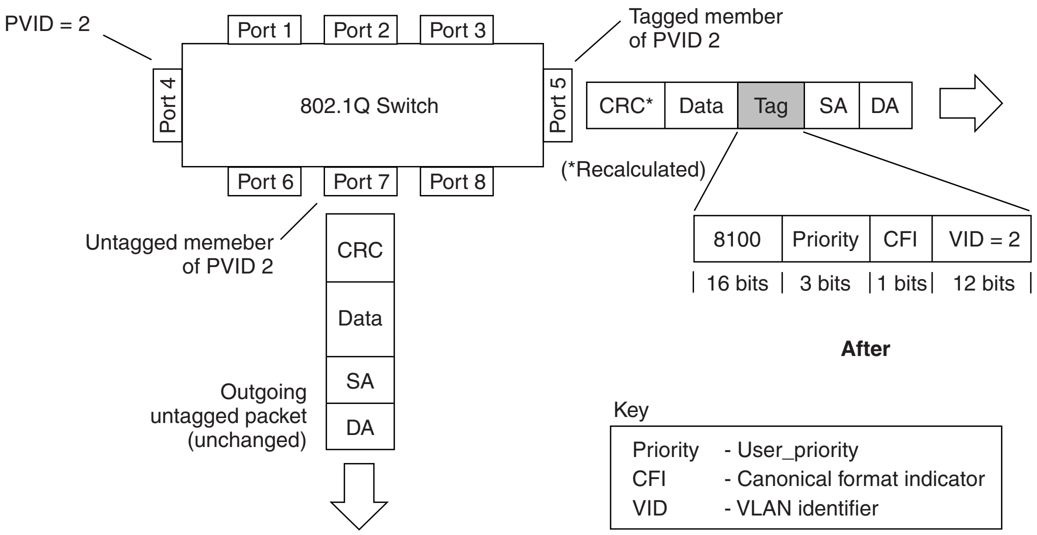

BS45012C

VLANs Spanning Multiple Switches

You can use VLANs to segment a network within a switch. When you connect multiple switches, it is possible to connect users of one VLAN with users of that same VLAN in another switch. However, the configuration guidelines depend on whether both switches support 32-bit 802.1Q tagging.

With 32-bit 802.1Q tagging enabled on a port for a VLAN, all frames leaving the port for that VLAN are *marked* as belonging to that specific VLAN. You can assign specific switch ports as members of one or more VLANs that span multiple switches, without interfering with the Spanning Tree Protocol.

VLANs Spanning Multiple 802.1Q Tagged Switches

Figure 36 shows VLANs spanning two Ethernet Routing Switch 3510-24T. The 32-bit 802.1Q tagging is enabled on S1, port 14 and on S2, port 13 for VLAN 1 and VLAN 2. Both ports are tagged members of VLAN 1 and VLAN 2.

Figure 36 VLANs spanning multiple 802.1Q tagged switches

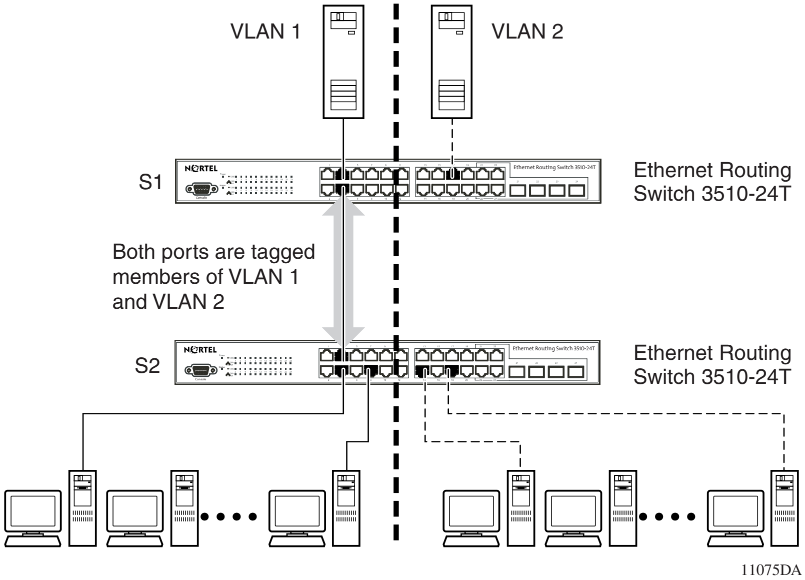
VLAN 1
VLAN 2
S1
Ethernet Routing Switch 3510-24T
Both ports are tagged members of VLAN 1 and VLAN 2
S2
Ethernet Routing Switch 3510-24T
11075DA

Because there is only one link between the two switches, the Spanning Tree Protocol (STP) treats this configuration as any other switch-to-switch connection. For this configuration to work properly, both switches must support the 32-bit 802.1Q tagging protocol.

VLANS Spanning Multiple Untagged Switches

Figure 37 shows VLANs spanning multiple untagged switches. In this configuration, Switch S2 does not support 32-bit 802.1Q tagging and you must use a single switch port on each switch for each VLAN.

For this configuration to work properly, you must set spanning tree participation to Disabled (the STP is not supported across multiple LANs).

Figure 37 VLANs spanning multiple untagged switches

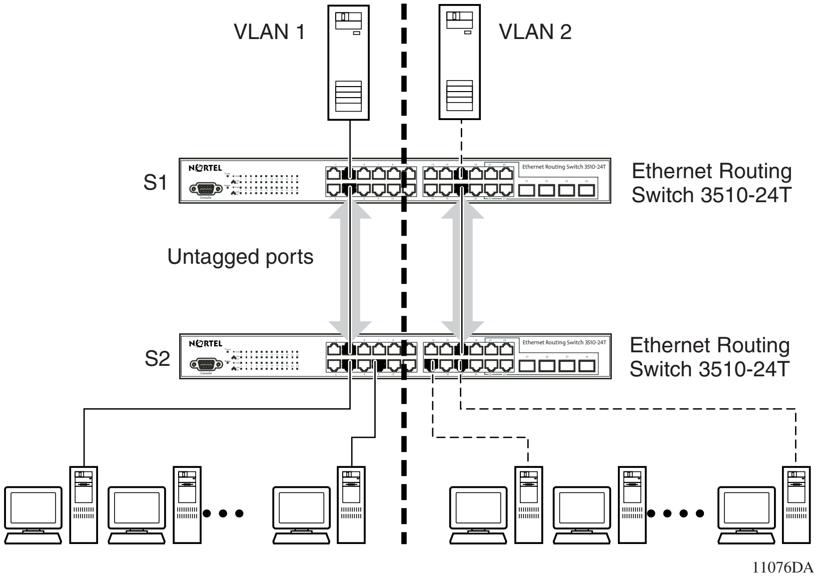
VLAN 1
VLAN 2
S1
Ethernet Routing Switch 3510-24T
Untagged ports
S2
Ethernet Routing Switch 3510-24T
11076DA

When the STP is enabled on these switches, only one link between the pair of switches will be forwarding traffic. Because each port belongs to only one VLAN at a time, connectivity on the other VLAN is lost. Exercise care when configuring the switches to ensure that the VLAN configuration does not conflict with spanning tree configuration.

To connect multiple VLANs across switches with redundant links, you must disable the STP on all participating switch ports. Figure 38 shows possible consequences of enabling the STP when using VLANs between untagged (non-802.1Q tagged) switches.

Figure 38 Possible problems with VLANs and Spanning Tree Protocol

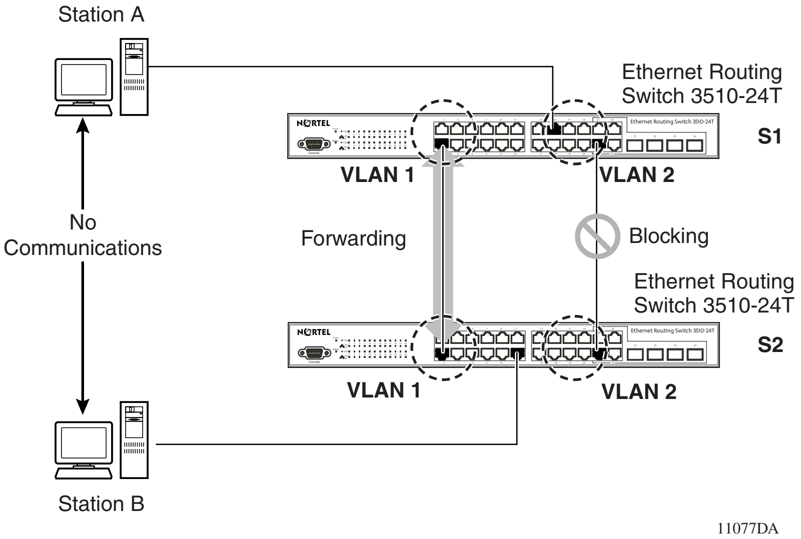

As shown in Figure 38, with STP enabled, only one connection between Switch S1 and Switch S2 is forwarding at any time. Communications failure occurs between VLAN 2 of S1 and VLAN 2 of S2, blocking communications between Stations A and B.

The STP selects the link connecting VLAN 1 on Switches S1 and S2 as the forwarding link based on port speed, duplex-mode, and port priority. Because the other link connecting VLAN 2 is in Blocking mode, stations on VLAN 2 in Switch S1 cannot communicate with stations in VLAN 2 on Switch S2. With multiple links only one link will be forwarding.

VLAN Workgroup Summary

This section summarizes the VLAN workgroup examples discussed in the previous sections of this chapter.

As shown in Figure 39, Switch S1 (Ethernet Routing Switch 3510-24T) is configured with multiple VLANs:

- Ports 17, 20, 25, and 26 are in VLAN 1.
- Ports 16, 18, 19, 21, and 24 are in VLAN 2.

- Port 22 is in VLAN 3.

Because S4 does not support 30-bit 802.1Q tagging, a single switch port on each switch must be used for each VLAN (see "VLANS Spanning Multiple Untagged Switches").

The connection to S2 requires only one link between the switches because S1 and S2 are both Ethernet Routing Switch 3510-24T that support 32-bit 802.1Q tagging (see "VLANs Spanning Multiple 802.1Q Tagged Switches").

Figure 39 VLAN configuration spanning multiple switches

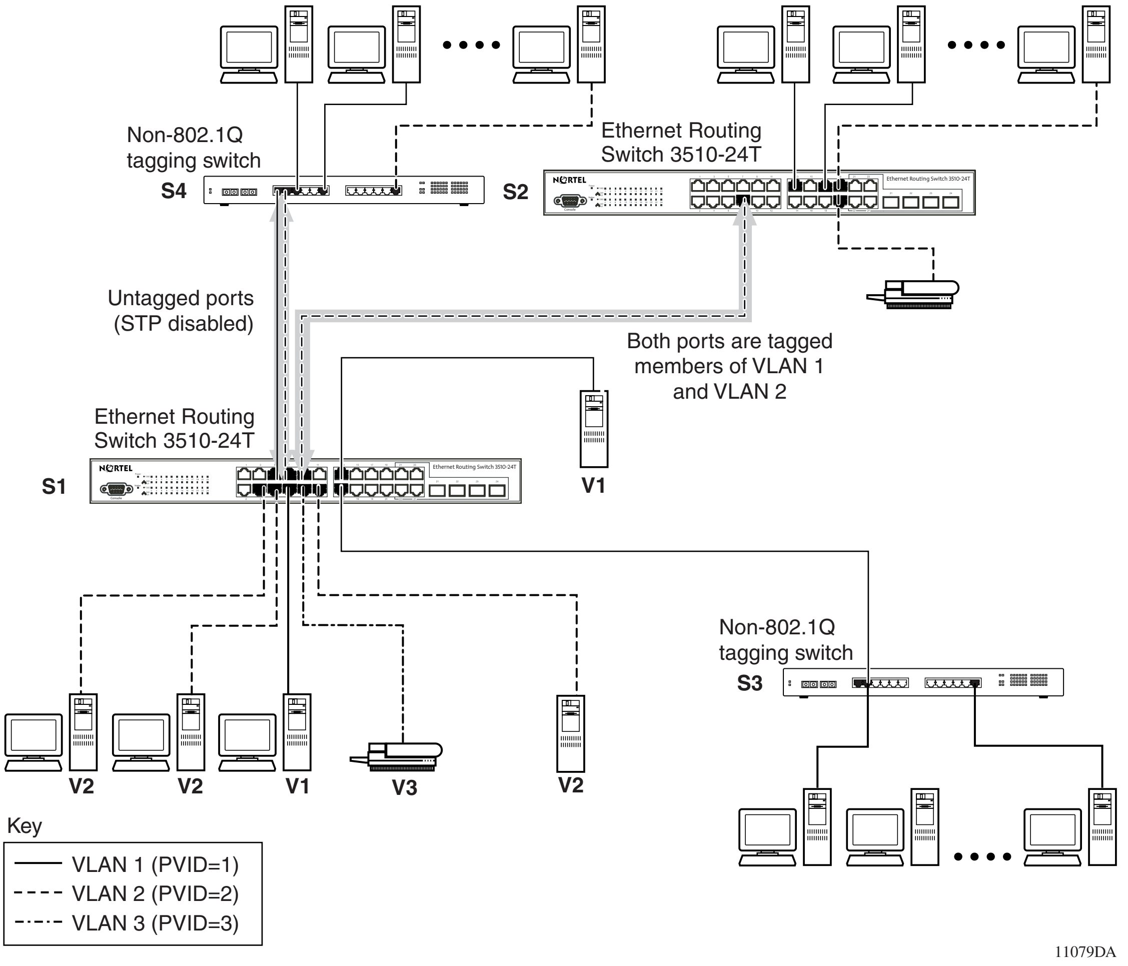

VLAN Configuration Rules

VLANs operate according to specific configuration rules. When creating VLANs, consider the following rules that determine how the configured VLAN reacts in any network topology:

- All ports that are involved in port mirroring must have memberships in the same VLANs. If a port is configured for port mirroring, the port's VLAN membership cannot be changed.
- If a port is a trunk group member, all trunk members are added or deleted from the VLAN.
- All ports involved in trunking and port mirroring must have the same VLAN configuration. If a port is on a trunk with a mirroring port, the VLAN configuration cannot be changed.
- VLANs are not dependent on Rate Limiting settings.
- If a port is an IGMP member on any VLAN, and is removed from a VLAN, the port's IGMP membership is also removed.
- If a port is added to a different VLAN, and it is already configured as a static router port, the port is configured as an IGMP member on that specific VLAN.

See Appendix B for configuration flowcharts that can help you use this feature.

Spanning Tree Protocol Groups

The Ethernet Routing Switch 3510-24Tsupports the Spanning Tree Protocol (STP) as defined in IEEE 802.1D. As defined in the IEEE 802.1D standard, the Spanning Tree Protocol detects and eliminates logical loops in a bridged or switched network. When multiple paths exist, the spanning tree algorithm configures the network so that a bridge or switch uses only the most efficient path. If that path fails, the protocol automatically reconfigures the network to make another path become active, thus sustaining network operations.

The Ethernet Routing Switch 3510-24T supports multiple Spanning Tree Groups (STGs). The Ethernet Routing Switch 3510-24T supports a maximum of 8 STGs, all in one standalone switch. Multiple STGs provide multiple data paths, which can be used for load-balancing and redundancy. You can enable load balancing between two Ethernet Routing Switch 3510-24T switches using multiple STGs by configuring each path with a different VLAN and then assigning each VLAN to a separate STG**.** Each STG is independent. Each STG sends its own Bridge Protocol Data Units (BPDUs), and each STG must be independently configured.

The STG, or bridge group, forms a loop-free topology that includes one or more virtual LANs (VLANs). The Ethernet Routing Switch 3510-24T supports multiple instances (8) of STGs running simultaneously.

The Ethernet Routing Switch 3510-24T supports a maximum of 256 VLANs. With a maximum of 8 STGs, on average, each STG will have 32 VLANs.

In the default configuration of the Ethernet Routing Switch 3510-24T, a single STG with the ID of 1 includes all ports on the switch. It is called the default STG. Although ports can be added to or deleted from the default STG, the default STG (STG1) itself cannot be deleted from the system. Also, you cannot delete the default VLAN (VLAN1) from STG1.

The tagging for the BPDUs from STG1, or the default STG, is user-configurable (as are tagging settings for all STGs). However, by default STG1 sends out only untagged BPDUs in order to operate with all devices that support only one instance of STP. (The default tagging of STG2 through STG8 is tagged.) The tagging setting for each STG is user-configurable.

Note: If the STG is tagging a BPDU, the BPDU packet is tagged *only* on a tagged port.

All other STGs, except the Default STG, must be created by the user. To become active, each STG must be enabled by the user after creation. Each STG will be assigned an ID number from 2 to 8 (the Default STG is assigned the ID number 1). You can assign ports or VLANs to an active STG. However, a port that is not a member of a VLAN will not be allowed to join an STG.

When you create an STG, all ports belonging to any assigned VLAN are automatically added to the STG. However, once the STG is created and the VLANs assigned, any *subsequent* ports you add to the VLAN must be manually added to the STG. Ports added to VLANs after they are part of an STG are not automatically added to the STG; the VLAN ports are automatically added *only* upon creation of the STG.

When you no longer need a particular STG, disable and delete that particular one. The procedure is to disable the STG, delete all VLAN and port memberships, and then delete the STG.

You can configure a unique multicast address for STGs 1 to 4.

Note: If you configure a unique multicast address for an STG, each device in that STG must also be configured with the same spanning tree multicast address.

STG Configuration Guidelines

This section provides important information on configuring STGs:

- An STG must be created in the following order:
  - Create the STG
  - Add the existing VLAN and port memberships
  - Enable the STG
- When you create a VLAN, that VLAN automatically belongs to STG 1, the default STG. If you want the VLAN in another STG, you must move the VLAN by assigning it to another STG.
- You move a newly created VLAN to an existing STG by following this order:
  - Create the VLAN
  - Add the VLAN to an existing STG
- When you add ports to a VLAN that is in an existing STG, you must *manually* add that port to the STG.
- You cannot delete or move VLAN1 from STG1.
- VLANs must be contained within a single STG; a VLAN cannot span multiple STGs. By confining VLANs within a single STG, you avoid problems with spanning tree blocking ports and causing a loss of connectivity within the VLAN. When a VLAN spans multiple switches, the VLAN must be within the same Spanning Tree Group (have the same STG ID) across all the switches.
- All members of a particular MultiLink Trunking (MLT) group must be assigned to the same STG; that is, they can belong to one and only one STG.
- A port that is not a member of any VLAN cannot be added to any STG. The port must be added to a VLAN, and that VLAN added to the desired STG.
- Tagged ports can belong to more than one STG, but untagged ports can belong to only one STG.
- When a tagged port belongs to more than one STG, the egress BPDUs are tagged to distinguish the BPDUs of one STG from those of another STG.

- Because some STP-compliant devices do not support tagging, you can configure whether to send tagged or untagged BPDUs, even from tagged ports, with the Ethernet Routing Switch 3510-24T. The VLAN ID for the tagged BPDUs will be 4000+STG ID.
  - The default VLAN ID for tagged BPDUs is as follows:

    4001—STG1
    4002—STG2
    4003—STG3
    4004—STG4
    4005—STG5
    4006—STG6
    4007—STG7
    4008—STG8
  - You can select a VLAN ID for tagged BPDUs for each STG. Valid VLAN IDs are 1 to 4094.
  - Tagged BPDUs cannot use the same VID as an active VLAN.
- An untagged port cannot span multiple STGs.
- When you remove a port from VLAN that belongs to an STG, that port will also be removed from the STG. However, if that port belongs to another VLAN in the same STG, the port remains in the STG.

  As an example, assume that port 1 belongs to VLAN1, and VLAN1 belongs to STG1. When you remove port 1 from VLAN1, port 1 is also removed from STG1.

  However, if port 1 belongs to both VLAN1 and VLAN2 and both VLANs belong to STG1, removing port 1 from VLAN1 does *not* remove port 1 from STG1 because VLAN2 is still a member of STG1.
- An STG cannot be deleted until you disable it. Additionally, you cannot delete an STG while it contains VLAN members, so you must first delete the VLANs from the STG.
- You can configure a unique multicast address for STG 1 to 4 *only.*

Spanning Tree Fast Learning

Spanning Tree Fast Learning is an enhanced port mode supported by the Ethernet Routing Switch 3510-24T. If you enable Spanning Tree Fast Learning on a port with no other bridges, the port is brought up more quickly following the switch initialization or a spanning tree change. The port goes through the normal blocking and learning states before the forwarding state, but the hold times for these states is the bridge hello timer (2 seconds by default) instead of the bridge forward delay timer (15 seconds by default). If the port sees a BPDU it will revert to regular behavior.

The port set with Fast Learning can forward data immediately, as soon as the switch learns that the port is enabled.

Fast Learning is intended for access ports where only one device is connected to the switch (as in workstations with no other spanning tree devices). It may not be desirable to wait the usual 30 to 35 seconds for spanning tree initialization and bridge learning.

Note: Use Spanning Tree Fast Learning with caution. This procedure is contrary to that specified in the IEEE 802.1D standard for Spanning Tree Protocol (STP), in which a port enters the blocking state following the initialization of the bridging device or from the disabled state when the port is enabled through configuration.

Internet Group Management Protocol Snooping

The Ethernet Routing Switch 3510-24T can sense Internet Group Management Protocol (IGMP) host membership reports from attached stations and use this information to set up a dedicated path between the requesting station and a local IP Multicast router. After the pathway is established, the Ethernet Routing Switch 3510-24T blocks the IP Multicast stream from exiting any other port that does not connect to another host member, thus conserving bandwidth. The following section describes how Ethernet Routing Switch 3510-24T provide the same benefit as IP Multicast routers, but in the local area.

IGMP is used by IP Multicast routers to learn about the existence of host group members on their directly attached subnets (see RFC 2236). The IP Multicast routers get this information by broadcasting IGMP queries and listening for IP hosts reporting their host group memberships. This process is used to set up a client/server relationship between an IP Multicast source that provides the data streams and the clients that want to receive the data.

By default, unknown multicast traffic is flooded to all ports in a VLAN. In situations where there is a multicast transmitter that is not doing IGMP and there are no multicast receivers, the traffic transmitted by the transmitter is flooded.

The NCLI commands for IGMP allow you to send all unknown multicast traffic to IGMP static router ports only. This traffic will not be forwarded to dynamically discovered mrouter ports. If you want to forward unknown unicast traffic to certain ports only, you can set those ports as static mrouter ports.

- When disabled, the Ethernet Routing Switch 3510-24T treats unknown multicast traffic like broadcast traffic (flood). This is the default behavior.
- User setting for the unknown multicast no flood feature will be stored in NVRAM.
- Nortel recommends that you enable this feature when IGMP snooping is enabled.

Note: You have a maximum of 256 multicast groups with the Ethernet Routing Switch 3510-24T.

Figure 40 shows how IGMP is used to set up the path between the client and server. As shown in this example, the IGMP host provides an IP Multicast stream to designated routers that forward the IP Multicast stream on their local network only if there is a recipient.

The client/server path is set up as follows:

1. The designated router sends out a host membership query to the subnet and receives host membership reports from end stations on the subnet.
2. The designated routers then set up a path between the IP Multicast stream source and the end stations.
3. Periodically, the router continues to query end stations on whether or not to continue participation.
4. As long as any client continues to participate, all clients, including nonparticipating end stations on that subnet, receive the IP Multicast stream.

Note: Although the nonparticipating end stations can filter the IP Multicast traffic, the IP Multicast traffic still exists on the subnet and consumes bandwidth.

IP Multicast can be optimized in a LAN by using IP Multicast filtering switches, such as the Ethernet Routing Switch 3510-24T.

As shown in Figure 40, a non-IP Multicast filtering switch causes IP Multicast traffic to be sent to all segments on the local subnet.

Figure 40 IP Multicast propagation with IGMP routing

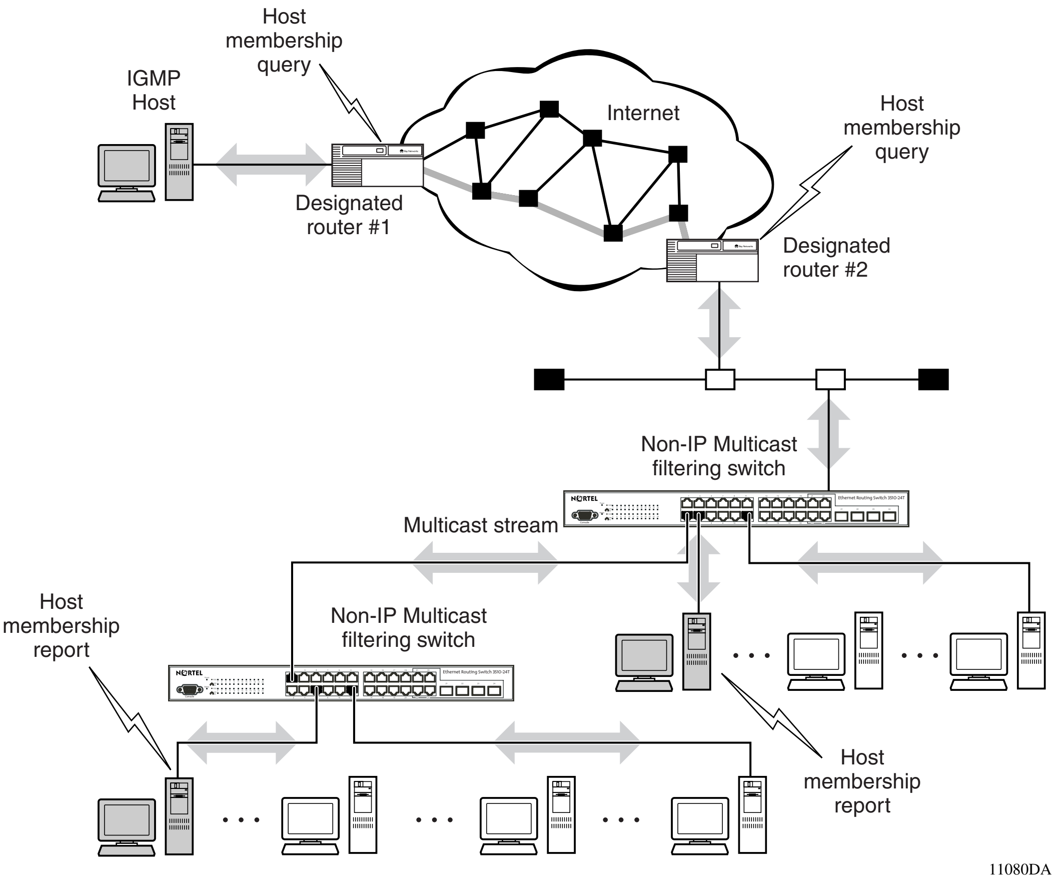

The Ethernet Routing Switch 3510-24T can automatically set up IP Multicast filters so the IP Multicast traffic is only directed to the participating end nodes (see Figure 41).

In Figure 41, switches S1 to S4 represent a LAN connected to an IP Multicast router. The router periodically sends Host Membership Queries to the LAN and listens for a response from end stations. All of the clients connected to switches S1 to S4 are aware of the queries from the router.

One client, connected to S2, responds with a host membership report. Switch S2 intercepts the report from that port, and generates a *proxy* report to its upstream neighbor, S1. Also, two clients connected to S4 respond with host membership reports, causing S4 to intercept the reports and to generate a *consolidated proxy report* to its upstream neighbor, S1.

Figure 41 Ethernet Routing Switch 3510-24T filtering IP multicast streams (1 of 2)

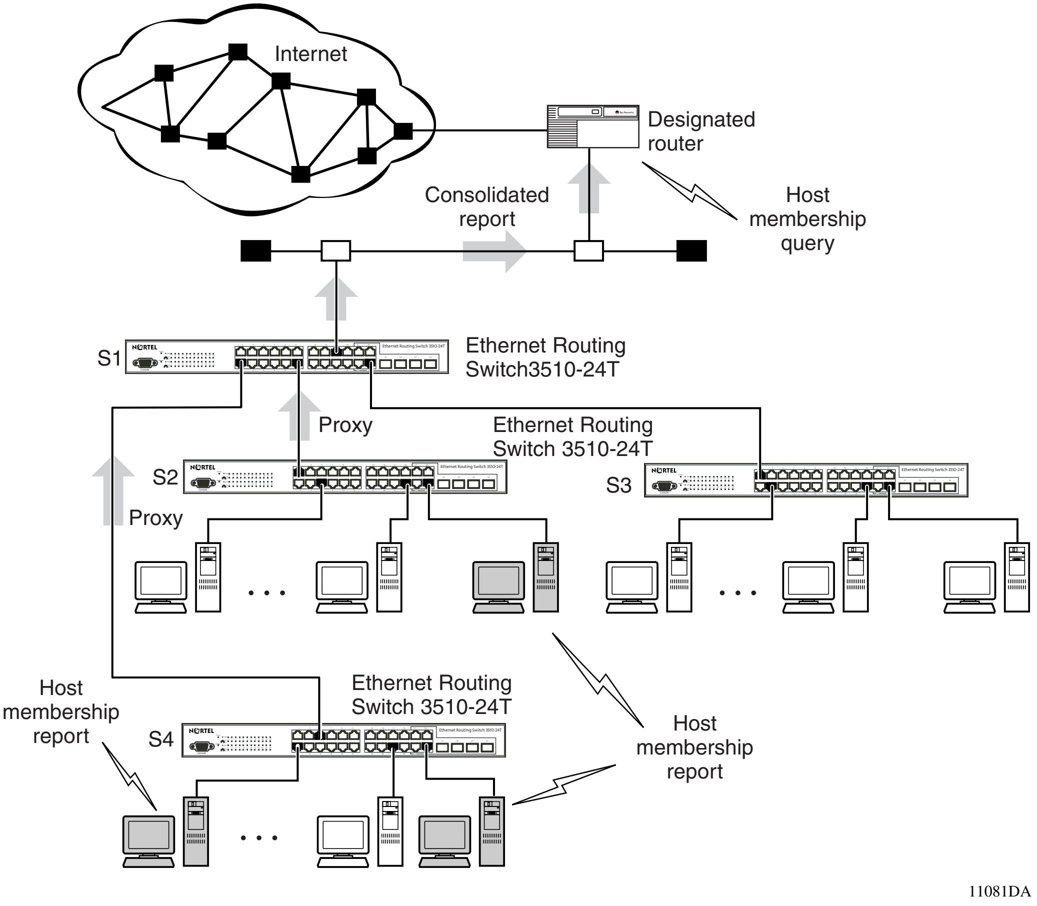

Switch S1 treats the consolidated proxy reports from S2 and S4 as if they were reports from any client connected to its ports, and generates a consolidated proxy report to the designated router. In this way, the router receives a single consolidated report from that entire subnet.

After the switches learn which ports are requesting access to the IP Multicast stream, all other ports not responding to the queries are blocked from receiving the IP Multicast (Figure 42).

Figure 42 Ethernet Routing Switch 3510-24T filtering IP multicast streams (2 of 2)

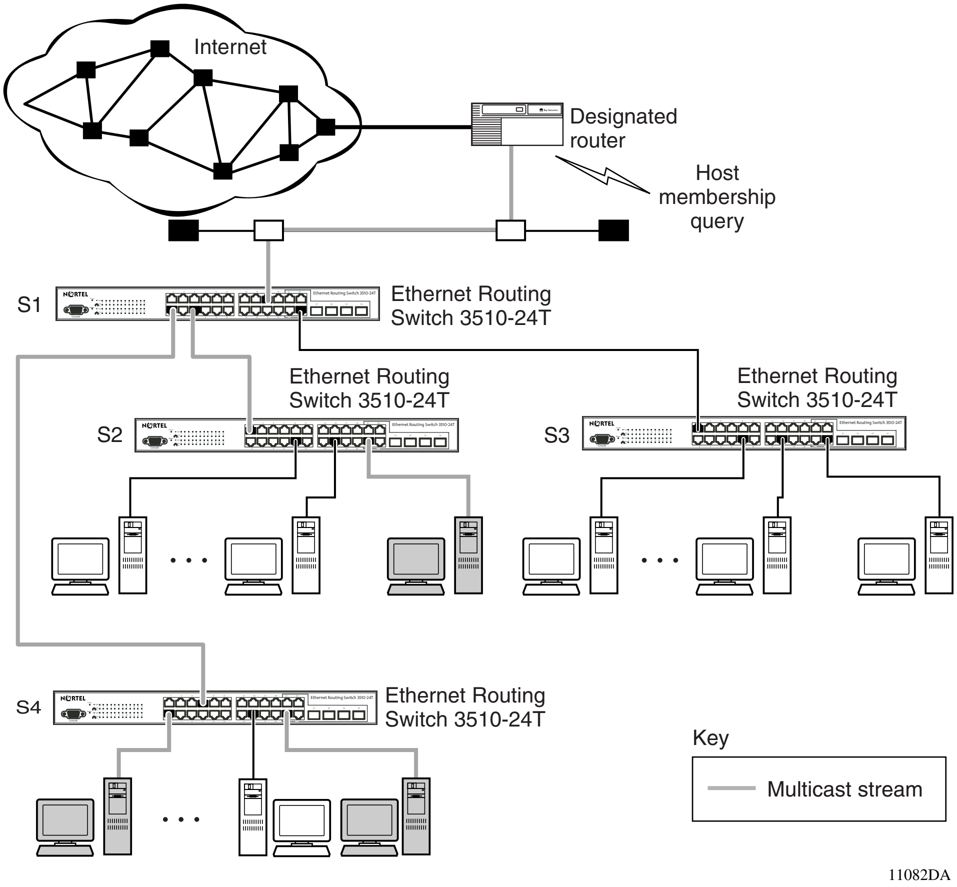
Internet
Designated router
Host membership query
S1
Ethernet Routing Switch 3510-24T
Ethernet Routing Switch 3510-24T
S2
Ethernet Routing Switch 3510-24T
S3
S4
Ethernet Routing Switch 3510-24T
Key
Multicast stream
11082DA

The consolidated proxy report generated by the switch remains transparent to layer 3 of the International Standardization Organization, Open Systems Interconnection (ISO/OSI) model. (The switch IP address and MAC address are not part of proxy report generation.) The last reporting IGMP group member in each VLAN represents all of the hosts in that VLAN and IGMP group.

IGMP Snooping Configuration Rules

The IGMP snooping feature operates according to specific configuration rules. When configuring your switch for IGMP snooping, consider the following rules that determine how the configuration reacts in any network topology:

- You have a maximum of 256 groups on the Ethernet Routing Switch 3510-24T.
- A port that is configured for port mirroring cannot be configured as a static router port.
- If a MultiLink Trunk member is configured as a static router port, all of the MultiLink trunk members are configured as static router ports. Also, if a static router port is removed, and it is a MultiLink Trunk member, all MultiLink trunk members are removed as static router port members, automatically.
- Static router ports must be port members of at least one VLAN.
- If a port is configured as a static router port, it is configured as a static router port for all VLANs on that port. The IGMP configuration is propagated through all VLANs of that port.
- If a static router port is removed, the membership for that port is removed from all VLANs of that port.
- The IGMP snooping feature is not STP-dependent.
- The IGMP snooping feature is not Rate Limiting-dependent.
- The snooping field must be enabled for the proxy field to have any valid meaning.
- Static router ports are configured per VLAN and per IGMP Version.

Note: Because IGMP snooping is set up per VLAN, all IGMP changes are implemented according to the VLAN configuration for the specified ports.

See Appendix B for configuration flowcharts that can help you use this feature.

IEEE 802.1p Prioritizing

For more information on prioritizing traffic, refer to Chapter 4, “Policy-enabled networks.”

You can use the VLAN Configuration screens to prioritize the order in which the switch forwards packets, on a per-port basis. For example, if messages from a specific segment are crucial to your operation, you can set the switch port connected to that segment to a higher priority level (by default, all switch ports are set to low priority). Untagged packets received by the switch on that port are tagged according to the priority level you assign to the port (see Figure 43).

Figure 43 Prioritizing packets

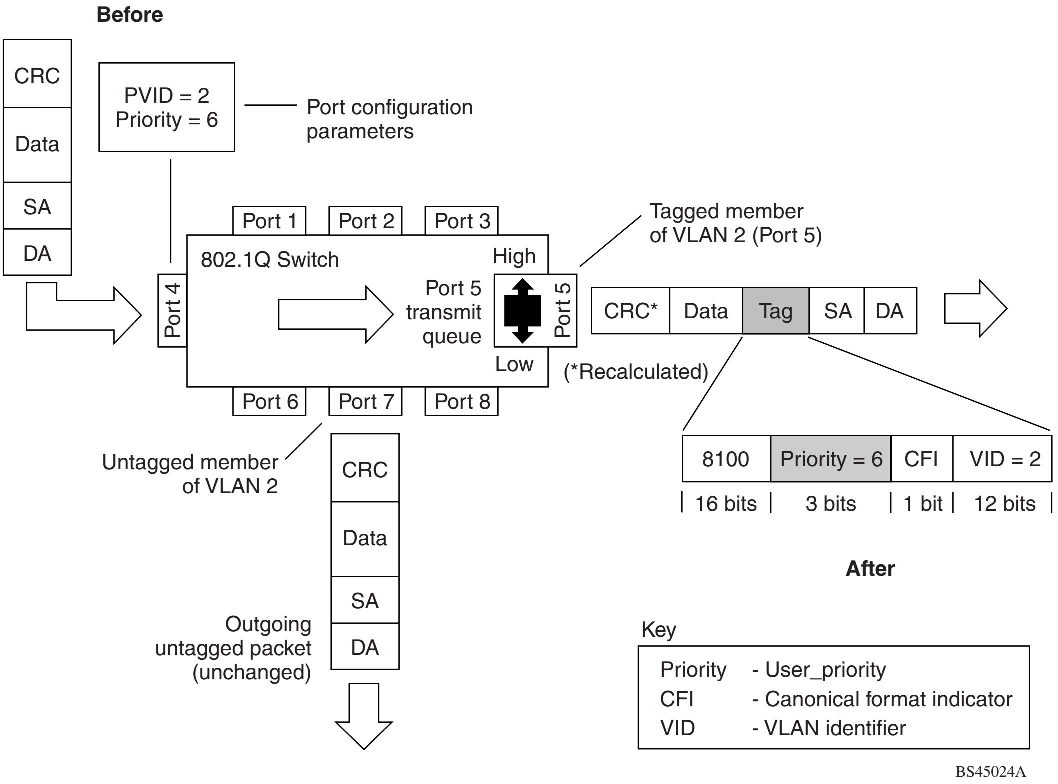

The newly tagged frame is read within the switch and sent to the port’s high or low transmit queue for disposition.

MultiLink Trunks

MultiLink Trunks allow you to group up to four switch ports together to form a link to another switch or server, thus increasing aggregate throughput of the interconnection between the devices (up to 8 GB in full-duplex mode). You can configure up to six MultiLink Trunks. MultiLink Trunking software detects misconfigured (or broken) trunk links and redirects traffic on the misconfigured or broken trunk link to other trunk members within that trunk.

You can use the Trunk Configuration screen with the CI menus, the Web-based management system, the CLI, or DM to create switch-to-switch and switch-to-server MultiLink Trunk links.

Note: Guidelines on configuring VLANs, STGs, and MLT are found throughout this chapter.

Client/Server Configuration Using MultiLink Trunks

Clients accessing data from the servers (FS1 and FS2) are provided with maximized bandwidth through trunks T1, T2, T3, T4, and T5. Trunk members (the ports making up each trunk) do not have to be consecutive switch ports; you can select ports randomly, as shown by T5.

With spanning tree *enabled*, one of the trunks (T2 or T3) acts as a redundant (backup) trunk to Switch S2. With spanning tree *disabled*, you must configure trunks T2 and T3 into separate VLANs for this configuration to function efficiently.

Figure 44 shows an example of how MultiLink Trunking can be used in a client/ server configuration. In this example, both servers connect directly to Switch S1. FS2 is connected through a trunk configuration (T1). The switch-to-switch connections are through trunks (T2, T3, T4, and T5).

Figure 44 Client/server configuration example

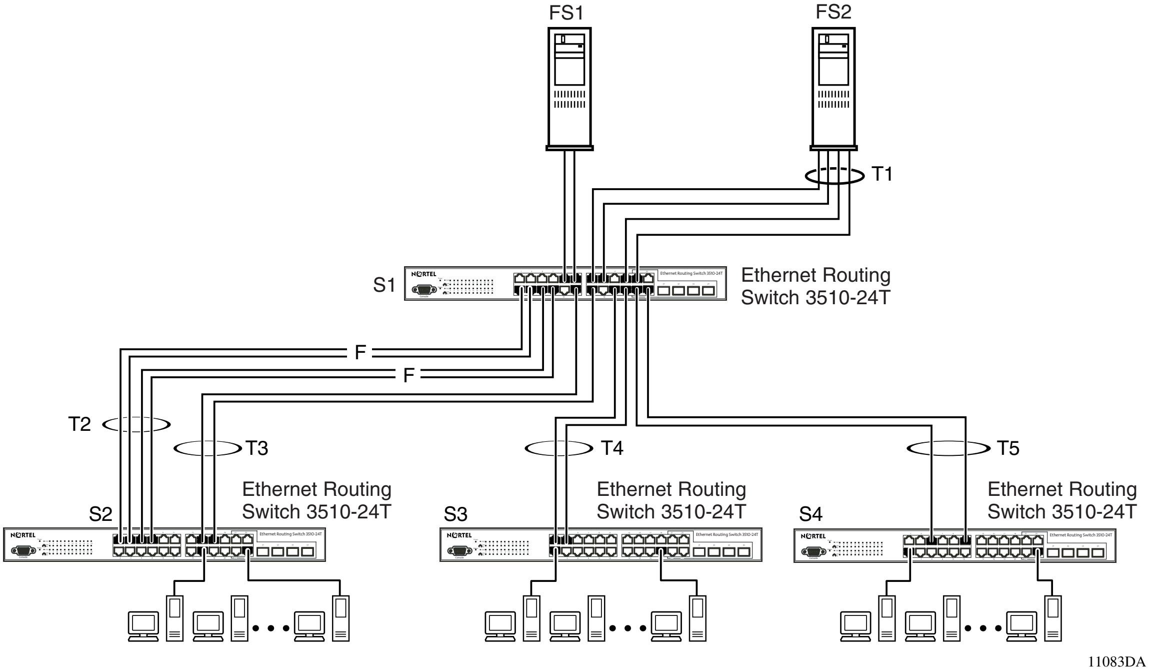
FS1
FS2
T1
S1
Ethernet Routing Switch 3510-24T
F
F
T2
T3
T4
T5
S2
Ethernet Routing Switch 3510-24T
S3
Ethernet Routing Switch 3510-24T
S4
Ethernet Routing Switch 3510-24T
11083DA

For detailed information about configuring trunks, see Chapter 3.

Before You Configure Trunks

When you create and enable a trunk, the trunk members (switch ports) take on certain settings necessary for correct operation of the MultiLink Trunking feature.

Before you configure your MultiLink Trunk, you must consider these settings, along with specific configuration rules, as follows:

1 Read the configuration rules provided in the next section, "MultiLink Trunking Configuration Rules".

2 Determine which switch ports (up to four) are to become *trunk members* (the specific ports making up the trunk). A minimum of two ports are required for each trunk.

Ensure that the chosen switch ports are set to Enabled.

Trunk member ports must have the same VLAN configuration.

3. All network cabling should be complete and stable before configuring any trunks, to avoid configuration errors.
4. Consider how the existing spanning tree will react to the new trunk configuration.
5. Consider how existing VLANs will be affected by the addition of a trunk.

MultiLink Trunking Configuration Rules

The MultiLink Trunking feature is deterministic, it operates according to specific configuration rules. Consider the following rules that determine how the MultiLink Trunk reacts in any network topology when creating trunks:

- Any port that participates in MultiLink Trunking must be an active port (set to Enabled via the Port Configuration screen or through network management).
- All trunk members must have the same VLAN configuration before the Trunk Configuration screen's Trunk Status field can be set to Enabled using CI menus.
- When an active port is configured in a trunk, the port becomes a *trunk member* when you set the Trunk Status field to Enabled. The spanning tree parameters for the port then change to reflect the new trunk settings.
- All trunk members must be in the same spanning tree group and can belong to only one spanning tree group.
- If you change the spanning tree participation of any trunk member to Enabled or Disabled, the spanning tree participation of all members of that trunk changes similarly.
- If you change the VLAN settings of any trunk member, the VLAN settings of all members of that trunk change similarly.
- When you set any trunk member to Disabled (not active) through the Port Configuration screen or through network management, the trunk member is removed from the trunk. The trunk member has to be reconfigured to rejoin the trunk through the Trunk Configuration screen on the CI menus, or another management system. A screen prompt precedes this action when you are using CI menus. A trunk member cannot be disabled if there are only two trunk members on the trunk.

Note: Do not disable the lowest numbered port in a trunk if Spanning Tree is enabled.

- You cannot configure a trunk member as a monitor port.
- Entire trunks cannot be monitored by a monitor port; however, trunk members can be monitored (see "Port-based Mirroring Configuration"").
- All trunk members must have identical IGMP configurations.
- If you change the IGMP snooping configuration for any trunk member, the IGMP snooping settings for all trunk members change.
- Nortel recommends that you do not enable MAC Address Security (or BaySecure) on trunk ports.
- The order of the port numbers is used in the packet distribution in MLT on the Ethernet Routing Switch 3510-24T.

Spanning Tree Considerations for MultiLink Trunks

The spanning tree Path Cost parameter is recalculated based on the aggregate bandwidth of the trunk.

For example, Figure 45 shows a two-port trunk (T1) with two port members operating at an aggregate bandwidth of 2 GB, with a comparable Path Cost of 1. Trunk 2 has two ports at 100 Mbps with a Path Cost of 5.

When the Path Cost calculations for both trunks are equal, the software chooses the trunk with the larger aggregate bandwidth (T1) to determine the most efficient path. Also, the trunk cannot span multiple spanning tree groups.

Note: The default spanning tree Path Cost for all gigabit ports (GB) is always equal to 1.
Be careful when configuring trunks so that you do not add one GB link physically in front of another trunk; the trunk will be blocked because they both have a Path Cost of 1.

Figure 45 Path Cost arbitration example

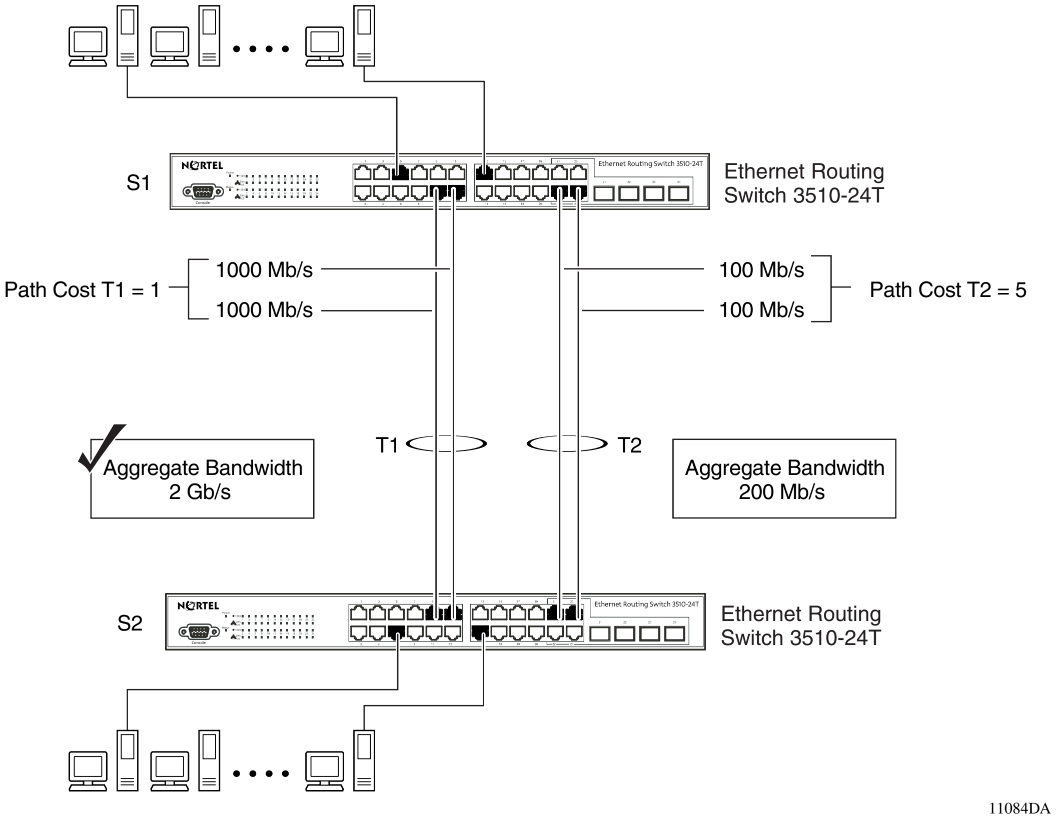
S1
Ethernet Routing Switch 3510-24T
Path Cost T1 = 1
1000 Mb/s
1000 Mb/s
100 Mb/s
100 Mb/s
Path Cost T2 = 5
T1
T2
Aggregate Bandwidth 2 Gb/s
Aggregate Bandwidth 200 Mb/s
S2
Ethernet Routing Switch 3510-24T
11084DA

The switch can also detect trunk member ports that are physically misconfigured. For example, in Figure 46, trunk member ports 2, 4, and 6 of Switch S1 are configured *correctly* to trunk member ports 7, 9, and 11 of Switch S2. The Spanning Tree Port Configuration screen for each switch shows the port state field for each port in the Forwarding state.

Figure 46 Example 1: correctly configured trunk

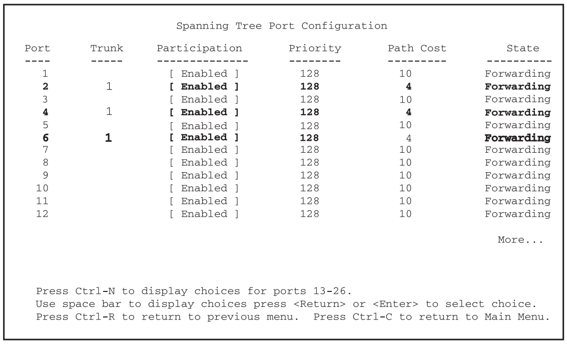

```
                       Spanning Tree Port Configuration

Port      Trunk     Participation      Priority      Path Cost        State
----      -----     --------------     --------      ---------      ----------
  1                  [ Enabled ]         128            10          Forwarding
  2         1        [ Enabled ]         128             4          Forwarding
  3                  [ Enabled ]         128            10          Forwarding
  4         1        [ Enabled ]         128             4          Forwarding
  5                  [ Enabled ]         128            10          Forwarding
  6         1        [ Enabled ]         128             4          Forwarding
  7                  [ Enabled ]         128            10          Forwarding
  8                  [ Enabled ]         128            10          Forwarding
  9                  [ Enabled ]         128            10          Forwarding
 10                  [ Enabled ]         128            10          Forwarding
 11                  [ Enabled ]         128            10          Forwarding
 12                  [ Enabled ]         128            10          Forwarding

                                                                     More...



  Press Ctrl-N to display choices for ports 13-26.
  Use space bar to display choices press <Return> or <Enter> to select choice.
  Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

S1 Port Configuration screen

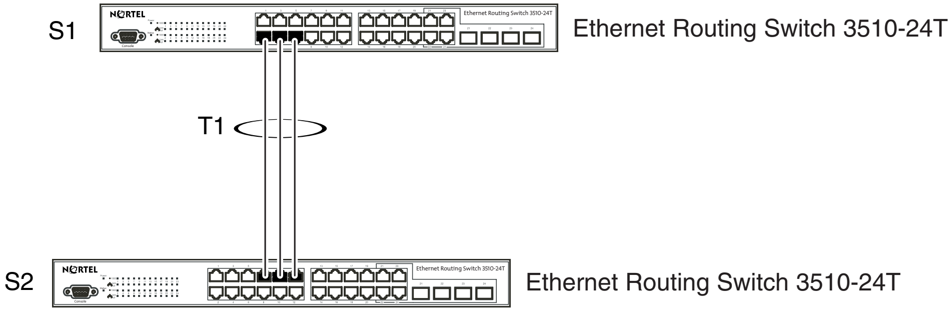
S1
Ethernet Routing Switch 3510-24T
T1
S2
Ethernet Routing Switch 3510-24T

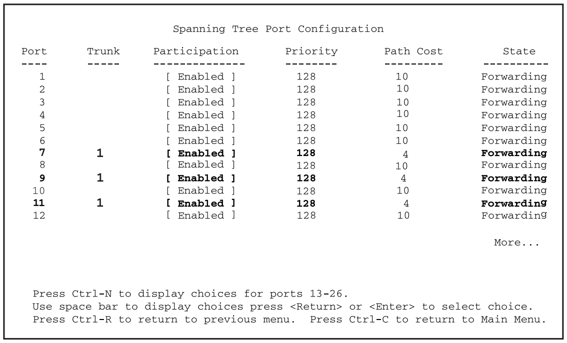

```
                       Spanning Tree Port Configuration

Port      Trunk     Participation      Priority      Path Cost        State
----      -----     --------------     --------      ---------      ----------
  1                  [ Enabled ]         128            10          Forwarding
  2                  [ Enabled ]         128            10          Forwarding
  3                  [ Enabled ]         128            10          Forwarding
  4                  [ Enabled ]         128            10          Forwarding
  5                  [ Enabled ]         128            10          Forwarding
  6                  [ Enabled ]         128            10          Forwarding
  7         1        [ Enabled ]         128             4          Forwarding
  8                  [ Enabled ]         128            10          Forwarding
  9         1        [ Enabled ]         128             4          Forwarding
 10                  [ Enabled ]         128            10          Forwarding
 11         1        [ Enabled ]         128             4          Forwarding
 12                  [ Enabled ]         128            10          Forwarding

                                                                     More...



  Press Ctrl-N to display choices for ports 13-26.
  Use space bar to display choices press <Return> or <Enter> to select choice.
  Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

S2 Port Configuration screen

11087DA

If Switch S2's trunk member port 11 is physically disconnected and then reconnected to port 13, the Spanning Tree Port Configuration screen for Switch S1 changes to show port 6 in the Blocking state (Figure 47).

Figure 47 Example 2: detecting a misconfigured port

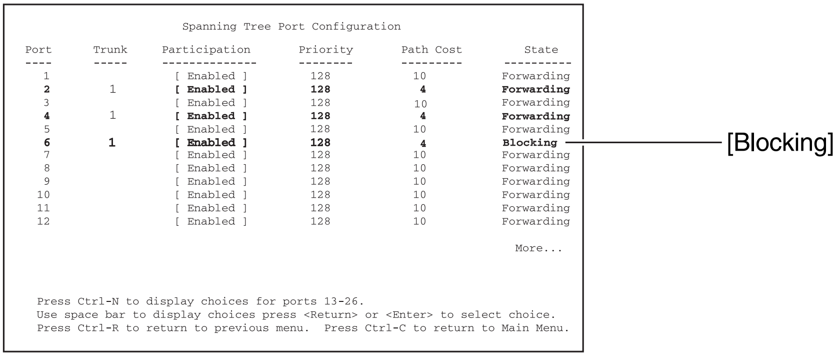

```
                       Spanning Tree Port Configuration

  Port       Trunk      Participation       Priority       Path Cost        State
  ----       -----      --------------      --------       ---------      ----------
    1                    [ Enabled ]          128             10          Forwarding
    2          1         [ Enabled ]          128              4          Forwarding
    3                    [ Enabled ]          128             10          Forwarding
    4          1         [ Enabled ]          128              4          Forwarding
    5                    [ Enabled ]          128             10          Forwarding
    6          1         [ Enabled ]          128              4          Blocking
    7                    [ Enabled ]          128             10          Forwarding
    8                    [ Enabled ]          128             10          Forwarding
    9                    [ Enabled ]          128             10          Forwarding
   10                    [ Enabled ]          128             10          Forwarding
   11                    [ Enabled ]          128             10          Forwarding
   12                    [ Enabled ]          128             10          Forwarding

                                                                            More...


    Press Ctrl-N to display choices for ports 13-26.
    Use space bar to display choices press <Return> or <Enter> to select choice.
    Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

S1 Port Configuration screen

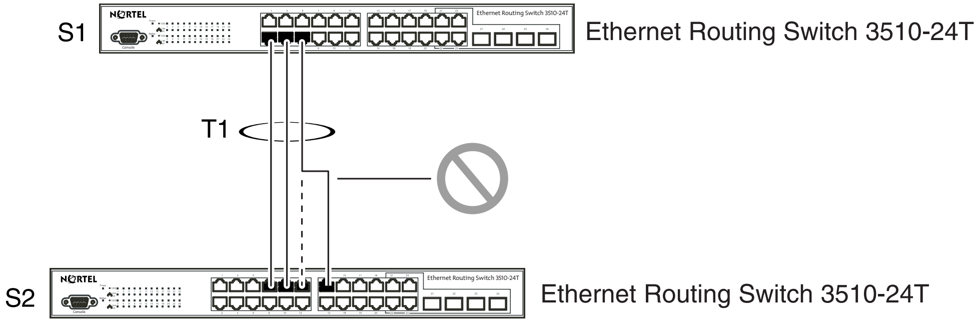

S1
Ethernet Routing Switch 3510-24T
T1
S2
Ethernet Routing Switch 3510-24T

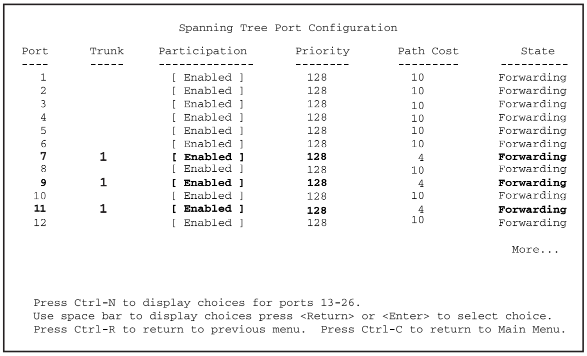

```
                       Spanning Tree Port Configuration

  Port       Trunk      Participation       Priority       Path Cost        State
  ----       -----      --------------      --------       ---------      ----------
    1                    [ Enabled ]          128             10          Forwarding
    2                    [ Enabled ]          128             10          Forwarding
    3                    [ Enabled ]          128             10          Forwarding
    4                    [ Enabled ]          128             10          Forwarding
    5                    [ Enabled ]          128             10          Forwarding
    6                    [ Enabled ]          128             10          Forwarding
    7          1         [ Enabled ]          128              4          Forwarding
    8                    [ Enabled ]          128             10          Forwarding
    9          1         [ Enabled ]          128              4          Forwarding
   10                    [ Enabled ]          128             10          Forwarding
   11          1         [ Enabled ]          128              4          Forwarding
   12                    [ Enabled ]          128             10          Forwarding

                                                                            More...


    Press Ctrl-N to display choices for ports 13-26.
    Use space bar to display choices press <Return> or <Enter> to select choice.
    Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

S2 Port Configuration screen

11088DA

Additional Tips About the MultiLink Trunking Feature

When you create a MultiLink Trunk, the individual trunk members (the specific ports that make up the trunk) logically connect and react as a single entity. For example, if you change spanning tree parameters for *any* trunk member, the spanning tree parameters for *all* trunk members change.

To change port membership in MultiLink Trunking, you must:

- Disable the trunk
- Make the change
- Re-enable the trunk

All configured trunks are indicated in the Spanning Tree Configuration screen. The Trunk field lists the active trunks, adjacent to the port numbers that correspond to the specific trunk member for that trunk.

The Spanning Tree parameters are display-only on the MultiLink Trunk screens.

To change any Spanning Tree parameters for the trunk members, you must reconfigure the Spanning Tree parameters.

The trunk is also viewed by management stations as a single spanning tree port. The spanning tree port is represented by the trunk member with the lowest port number. For example, if ports 13, 14, 15, and 16 are trunk members of trunk T1, the management station views trunk T1 as spanning tree port 13.

Note: Do not disable the lowest numbered port in a trunk if Spanning Tree is enabled.

See also Appendix B for configuration flowcharts that can help you use this feature.

Port Mirroring

You can designate one of your switch ports to monitor traffic on any two specified switch ports (port-based) or to monitor traffic to or from any two specified addresses that the switch has learned (address-based).

The following sections provide sample configurations for both monitoring modes available with the Port Mirroring feature:

- Port-based mirroring
- Address-based mirroring

A sample Port Mirroring Configuration screen accompanies each network configuration example. Note that the displayed screens do not show all of the screen prompts that precede some actions.

Note: Use the CI menus, the CLI, or the Web-based management system to configure port mirroring.

For example, when you configure a switch for port mirroring or when you modify an existing port mirroring configuration, the new configuration does not take effect until you respond [Yes] to the following screen prompt:

```
Is your port mirroring configuration complete?     [ Yes ]
```

Port-based Mirroring Configuration

Figure 48 shows an example of a port-based mirroring configuration where port 20 is designated as the monitor port for ports 21 and 22 of Switch S1. Although this example shows ports 21 and 22 monitored by the monitor port (port 20), any of the trunk members of T1 and T2 can also be monitored.

In this example, Figure 48 shows port X and port Y as members of Trunk T1 and Trunk T2. Port X and port Y are not required to always be members of Trunk T1 and Trunk T2.

Note: Trunks cannot be monitored and trunk members cannot be configured as monitor ports (see "MultiLink Trunking Configuration Rules").

Figure 48 shows the Port Mirroring Configuration screen setup for this example.

Figure 48 Port-based mirroring configuration example

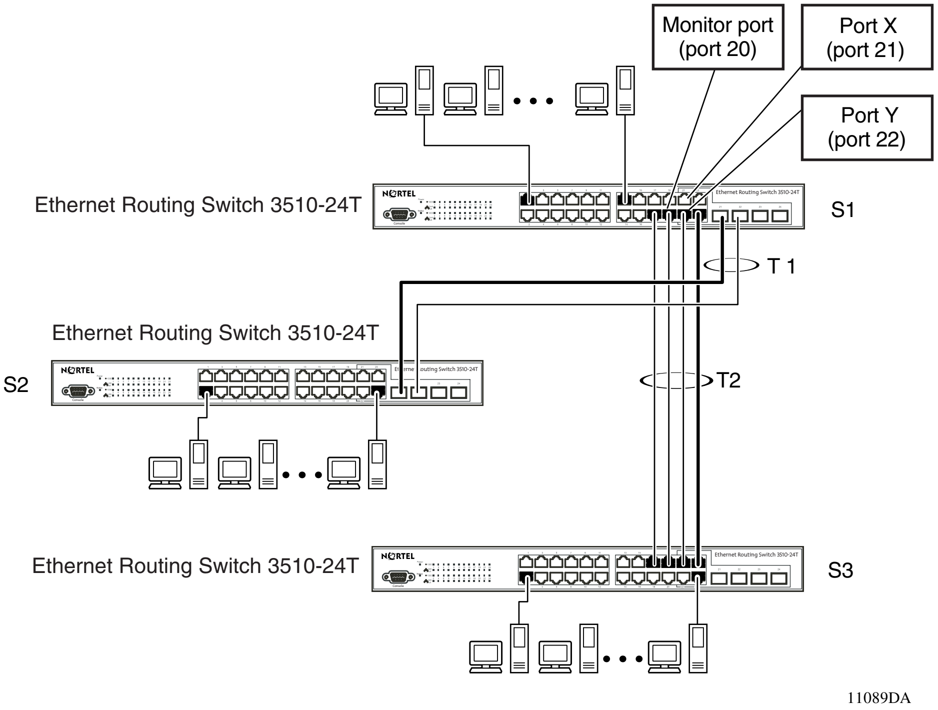
Monitor port
(port 20)
Port X
(port 21)
Port Y
(port 22)
Ethernet Routing Switch 3510-24T
S1
T 1
Ethernet Routing Switch 3510-24T
S2
T2
Ethernet Routing Switch 3510-24T
S3
11089DA

In the configuration example shown in Figure 48, the designated monitor port (port 20) can be set to monitor traffic in any of the following modes:

- Monitor all traffic received by port X.
- Monitor all traffic transmitted by port X.
- Monitor all traffic received and transmitted by port X.

- Monitor all traffic received by port X or transmitted by port Y.
- Monitor all traffic received by port X (destined to port Y) and then transmitted by port Y.
- Monitor all traffic received/transmitted by port X and transmitted/received by port Y (conversations between port X and port Y).

As shown in the Port Mirroring Configuration screen example (Figure 49), port 20 is designated as the Monitor Port for ports 21 and 22 in Switch S1.

The screen data displayed at the bottom of the screen shows the currently active port mirroring configuration.

Figure 49 Port Mirroring Configuration port-based screen example

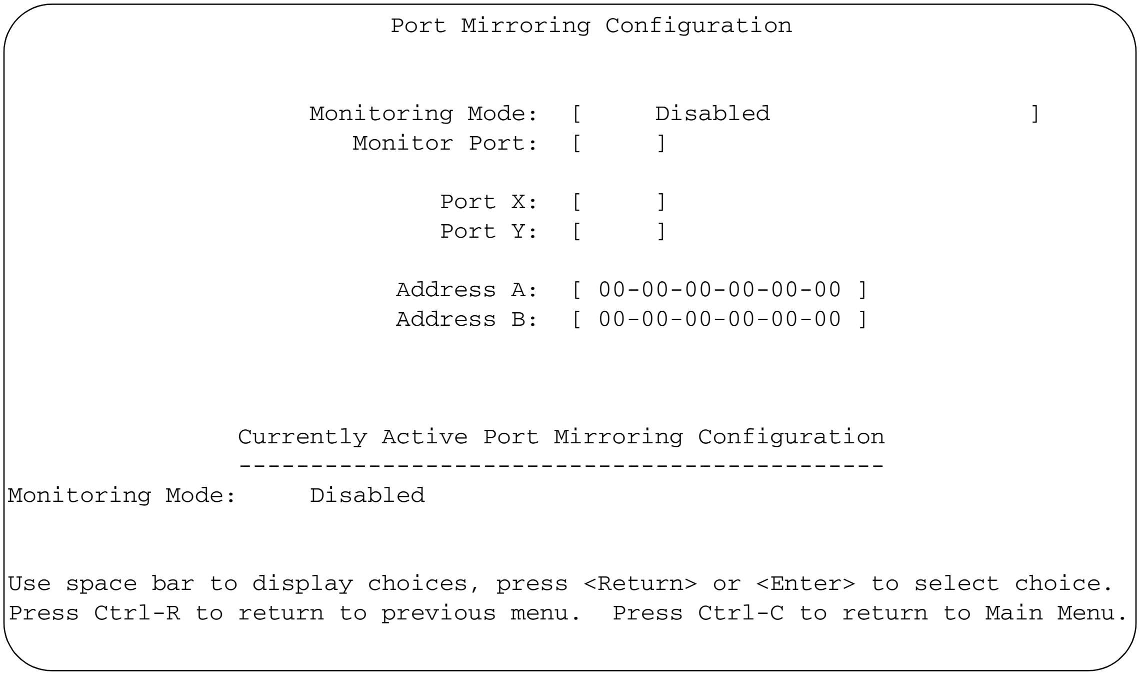

```
                           Port Mirroring Configuration


                     Monitoring Mode:  [     Disabled                   ]
                        Monitor Port:  [     ]

                              Port X:  [     ]
                              Port Y:  [     ]

                           Address A:  [ 00-00-00-00-00-00 ]
                           Address B:  [ 00-00-00-00-00-00 ]



                Currently Active Port Mirroring Configuration
                ---------------------------------------------
Monitoring Mode:     Disabled


Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Address-based Mirroring Configuration

Figure 50 shows an example of an address-based mirroring configuration where port 20, the designated monitor port for Switch S1, is monitoring traffic occurring between address A and address B.

In this configuration, the designated monitor port (port 20) can be set to monitor traffic in any of the following modes:

- Monitor all traffic transmitted from address A to any address.
- Monitor all traffic received by address A from any address.
- Monitor all traffic received by or transmitted by address A.
- Monitor all traffic transmitted by address A to address B.
- Monitor all traffic between address A and address B (conversation between the two stations).

Figure 50 Address-based mirroring configuration example

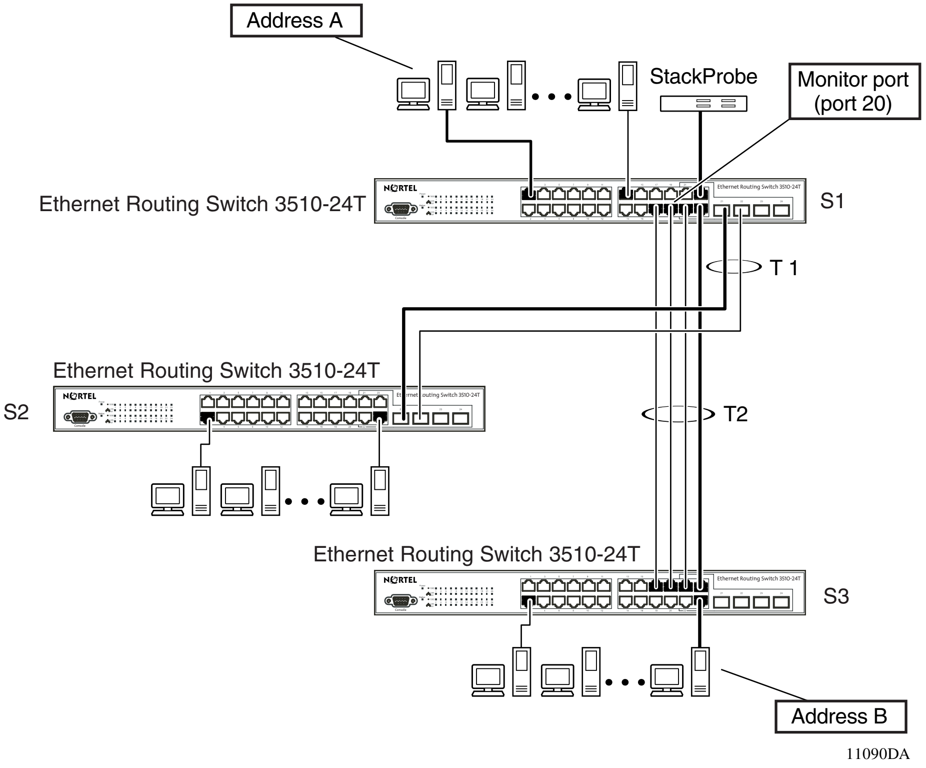
Address A
StackProbe
Monitor port
(port 20)
Ethernet Routing Switch 3510-24T
S1
T 1
Ethernet Routing Switch 3510-24T
S2
T2
Ethernet Routing Switch 3510-24T
S3
Address B
11090DA

In this example, port 20 becomes the designated Monitor Port for Switch S1 when you press Enter in response to the [Yes] screen prompt.

➡ **Note:** The screen data displayed at the bottom of the screen changes to show the *new* currently active port mirroring configuration *after* you press Enter.

➡ **Note:** When you enter MAC addresses in this screen, they are also displayed in the MAC Address Table screen (see Chapter 3).

Port Mirroring Configuration Rules

The following configuration rules apply to any port mirroring configuration:

- You cannot configure a monitor port as a trunk member or IGMP member.
- A monitor port cannot be used for normal switch functions.
- When you configure a port as a monitor port, the port is automatically disabled from participating in the spanning tree. When you reconfigure the port as a standard switch port (no longer a monitor port), the port is enabled for spanning tree participation.
- When you create a *port-based* port mirroring configuration, be sure that the monitor port and both of the mirrored ports, port X and port Y, have the same configuration. Use the VLAN Configuration screen to configure the VLAN (see Chapter 3).
- VLAN configuration settings for any ports configured for port-based mirroring cannot be changed. Use the Port Mirroring Configuration screen to disable port mirroring (or reconfigure the port mirroring ports), then change the VLAN configuration settings.
- For port-based monitoring of traffic, use one of the following modes for monitoring broadcast, IP Multicast, or unknown DA frames:
  - — Monitor all traffic received by port X.
  - — Monitor all traffic transmitted by port X.
  - — Monitor all traffic received and transmitted by port X.

For more information about using the Port Mirroring feature, see Chapter 3.

See also Appendix B for configuration flowcharts that can help you use this feature.

Figure 51 shows the Port Mirroring Configuration screen setup for this example.

Figure 51 Port Mirroring Configuration address-based screen example

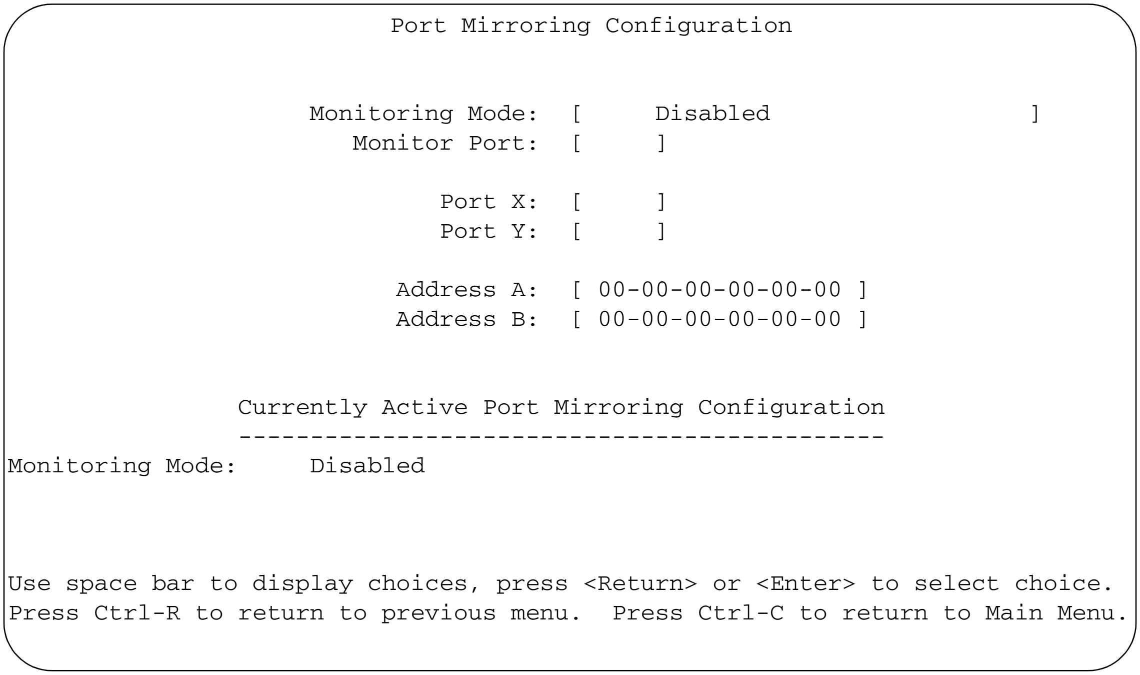

```
                           Port Mirroring Configuration


                     Monitoring Mode:  [     Disabled                  ]
                        Monitor Port:  [     ]

                              Port X:  [     ]
                              Port Y:  [     ]

                           Address A:  [ 00-00-00-00-00-00 ]
                           Address B:  [ 00-00-00-00-00-00 ]


                Currently Active Port Mirroring Configuration
                ---------------------------------------------
Monitoring Mode:     Disabled



Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Chapter 3
Using the Console Interface

This chapter describes how to configure and manage the Ethernet Routing Switch 3510-24T using the menu-driven console interface (CI).

This chapter covers the following topics:

- "Accessing the CI Menus and Screens", next
- "Using the CI Menus and Screens"
- "Main Menu"

Accessing the CI Menus and Screens

You can access the console interface (CI) menus and screens locally through a console terminal attached to your Ethernet Routing Switch 3510-24T, remotely through a dial-up modem connection, or in-band through a Telnet session (see Chapter 1).

Note: If you have a properly configured BootP server in your network, it detects the IP address; you will not need to configure the IP address.

For information about SNMP, see your network management documentation. You can also manage the Ethernet Routing Switch 3510-24T using the Nortel Command Line Interface (NCLI), the Web-based management system, or Device Manager (DM). For more information on using these management systems, consult the "Related Publications" in the Preface.

Using the CI Menus and Screens

The CI menus and screens provide options that allow you to configure and manage Ethernet Routing Switch 3510-24T. Help prompts at the bottom of each menu and screen explain how to enter data in the highlighted field and how to navigate the menus and screens.

The console port default settings are:

- 9600 baud
- 8 data bits
- 1 stop bit
- no parity as the communications format
- flow control set to enabled (Xon/Xoff supported)

Some CI screen options allow you to toggle among several possible values; other options allow you to set or modify a parameter.

Using Telnet to Access the CI Menus and Screens

When you use Telnet to access the CI menus and screens, set the terminal Preferences to VT100 Arrows and VT-100/ANSI and as shown in Figure 52.

Figure 52 Terminal preference settings

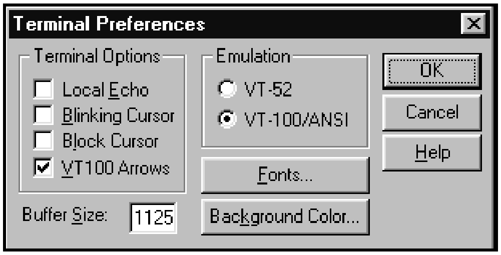

Terminal Preferences

Terminal Options

☐ Local Echo
☐ Blinking Cursor
☐ Block Cursor
☑ VT100 Arrows

Emulation

○ VT-52
● VT-100/ANSI

Buffer Size: 1125

Navigating the CI Menus and Screens

Use the following methods to navigate the CI menus and screens.

To select a menu option:

1. Use the arrow keys to highlight the option name.
2. Press [Enter].

   The option takes effect immediately after you press [Enter].

   Alternatively, you can press the key corresponding to the underlined letter in the option name. For example, to select the Switch Configuration option in the main menu, press the w key. Note that the text characters are not case-sensitive.

Following are the additional navigation aids:

- To toggle between values in a form:
  - Use the spacebar to highlight the value.
  - Press [Enter].
- To clear a string field:
  - Position the cursor in the string field.
  - Press [Ctrl]-K.
- To return to the previous menu, press [Ctrl]-R.
- To go to the next screen in a series, press [Ctrl]-N.
- To return to the main menu at any time, press [Ctrl]-C.
- Press [Ctrl + Backspace] or [Del] to delete entered text.
- Options that appear in brackets (for example, [Enabled]) are user-settable options.

Screen Fields and Descriptions

Figure 53 shows a map of the CI screens. The remainder of this chapter describes the CI screens and their fields, beginning with the main menu.

Figure 53 Map of console interface screens

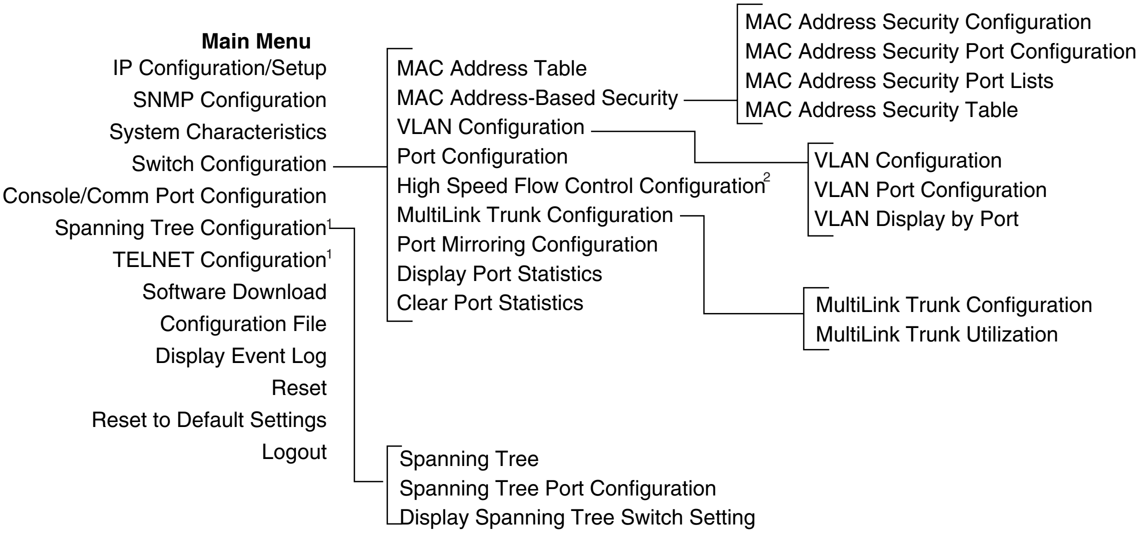
Main Menu
IP Configuration/Setup
SNMP Configuration
System Characteristics
Switch Configuration
Console/Comm Port Configuration
Spanning Tree Configuration[1]
TELNET Configuration[1]
Software Download
Configuration File
Display Event Log
Reset
Reset to Default Settings
Logout
MAC Address Table
MAC Address-Based Security
VLAN Configuration
Port Configuration
High Speed Flow Control Configuration[2]
MultiLink Trunk Configuration
Port Mirroring Configuration
Display Port Statistics
Clear Port Statistics
MAC Address Security Configuration
MAC Address Security Port Configuration
MAC Address Security Port Lists
MAC Address Security Table
VLAN Configuration
VLAN Port Configuration
VLAN Display by Port
MultiLink Trunk Configuration
MultiLink Trunk Utilization
Spanning Tree
Spanning Tree Port Configuration
Display Spanning Tree Switch Setting

BS45041F

The CI screens for your specific switch model show the correct model name in the main menu screen title and the correct number of ports and port types in the Port Configuration screen.

Note: The field values shown in the CI screens in this section are provided as examples only.

Main Menu

This section describes the options available from the CI main menu (Figure 54). The CI screens and submenus for these options are described in the following sections.

Note: Some menu options shown in this main menu example and in other screen examples in this chapter may not appear on your screen, depending on the switch options installed.

Figure 54 Console interface main menu

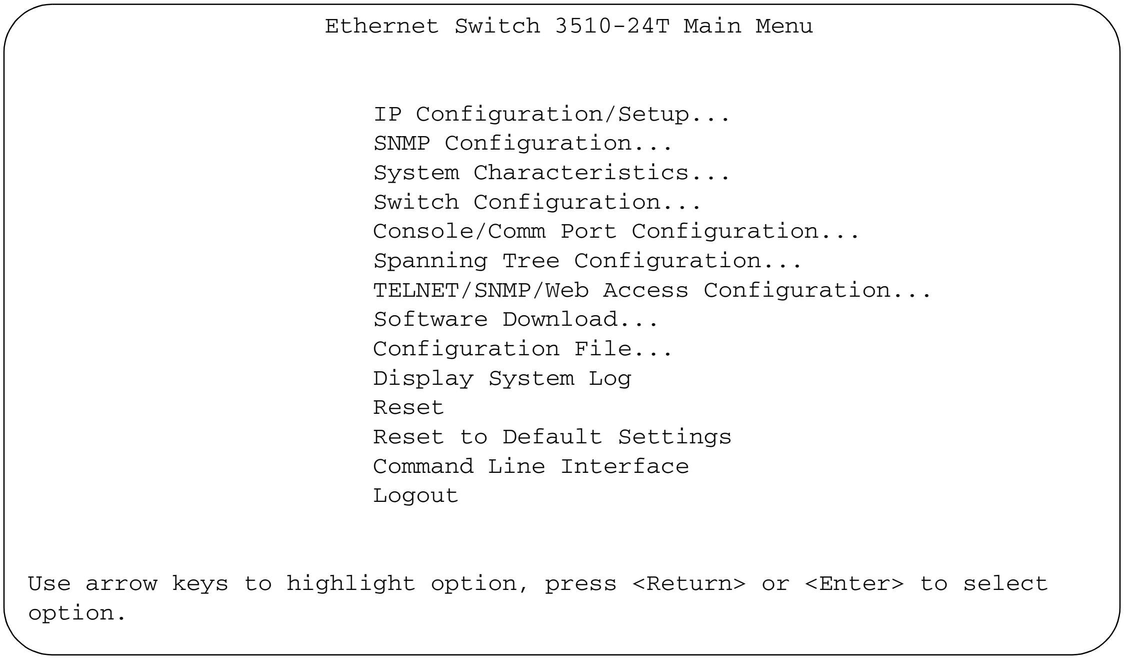

```
                     Ethernet Switch 3510-24T Main Menu


                          IP Configuration/Setup...
                          SNMP Configuration...
                          System Characteristics...
                          Switch Configuration...
                          Console/Comm Port Configuration...
                          Spanning Tree Configuration...
                          TELNET/SNMP/Web Access Configuration...
                          Software Download...
                          Configuration File...
                          Display System Log
                          Reset
                          Reset to Default Settings
                          Command Line Interface
                          Logout


Use arrow keys to highlight option, press <Return> or <Enter> to select
option.
```

Table 23 describes the CI main menu options

Table 23 Console interface Main Menu options

Option	Description
IP Configuration/Setup...	Displays the IP Configuration/Setup screen (see "IP Configuration/Setup Screen"" on page 148). This screen allows you to set or modify IP configuration parameters and to ping other network devices.
SNMP Configuration...	Displays the SNMP Configuration screen (see "SNMP Configuration Screen"" on page 154). This screen allows you to set or modify the SNMP read-only community and read-write community strings, enable or disable the authentication trap and the link Up/down trap, set the IP address of trap receivers, and set the trap community strings.
System Characteristics...	Displays the System Characteristics screen (see "System Characteristics Screen"" on page 156). This screen allows you to view switch characteristics, including number of resets, power status, hardware and software version, and MAC address. This screen also contains three user-configurable fields: sysContact, sysName, and sysLocation.
Switch Configuration...	Displays the Switch Configuration Menu screen (see "Switch Configuration Menu Screen"" on page 158). This menu provides the following configuration options: MAC Address Table, MAC Address-Based Security, EAPOL Security Configuration, VLAN Configuration, Port Configuration, High Speed Flow Control, MultiLink Trunk Configuration, Port Mirroring Configuration, Rate Limiting Configuration, IGMP Configuration, Display Port Statistics, and Clear All Port Statistics.
Console/Comm Port Configuration...	Displays the Console/Comm Port Configuration screen (see "Console/Comm Port Configuration Screen"" on page 212). This screen allows you to configure and modify the console/Comm port parameters, including the console port speed and password settings for the switch.
Spanning Tree Configuration...	Displays the Spanning Tree Configuration Menu (see "Spanning Tree Configuration Menu Screen"" on page 219). This menu provides the following options: Spanning Tree Group Configuration, Spanning Tree Port Configuration, Display Spanning Tree Switch Settings, and Display Spanning Tree VLAN Membership.
TELNET/SNMP/Web Access Configuration...	Displays the TELNET/SNMP/Web Access Configuration screen (see "TELNET/SNMP/Web Access Configuration Screen"" on page 232). This screen allows you to set your switch to enable a user at a remote console terminal to communicate with the Ethernet Routing Switch 3510-24T as if the console terminal were directly connected to it. You can have up to four active Telnet sessions running at one time in either a standalone switch. You can use the Command Line Interface (CLI), DM, or Web-based management system or these menus with a Telnet session. This screen also allows you to set the switch to allow up to 10 IP addresses to access the switch using either these management systems or SNMP access
Software Download...	Displays the Software Download screen (see "Software Download Screen"" on page 235). This screen allows you to revise the Ethernet Routing Switch 3510-24T software image that is located in nonvolatile flash memory (NVRAM).

Table 23 Console interface Main Menu options (continued)

Option	Description
Configuration File...	Displays the Configuration File Menu screen (see “Configuration File Menu Screen”” on page 239). This menu provides the following options: Configuration File Download/Upload and ASCII Configuration File Download.
Display System Log	Displays the System Log screen (see “System Log Screen”).
Reset	Resets the switch with the current configuration settings. This option is followed by a screen prompt that precedes the action. Enter Yes to reset the switch; enter No to abort the option: • When you select this option, the switch resets, runs a self-test, then displays the Nortel logo screen. Press [Ctrl]-Y to access the Ethernet Routing Switch 3510-24T main menu.
Reset to Default Settings	Resets the switch to the factory default configuration settings. This option is followed by a screen prompt that precedes the action. Enter Yes to reset the switch to the factory default configuration settings; enter No to abort the option: • When you select this option, the switch resets, runs a self-test, then displays the Nortel logo screen. Press [Ctrl]-Y to access the Ethernet Routing Switch 3510-24T main menu.
	⛔ **Caution:** If you choose the Reset to Default Settings option, all of your configured settings will be replaced with factory default settings when you press [Enter]
	⛔ **Achtung:** Bei Auswahl des Befehls zur Rücksetzung auf die Standardeinstellungen werden alle von Ihnen konfigurierten Einstellungen durch die werkseitigen Standardeinstellungen ersetzt, wenn Sie die Eingabetaste drücken.
	⛔ **Attention:** Si vous restaurez la configuration usine, votre configuration courante sera remplacée par la configuration usine dès que vous appuierez sur [Entrée].
	⛔ **Precaución:** Si selecciona el comando Restaurar valores predeterminados, todos los valores de configuración se sustituirán por las valores predeterminados en fábrica al pulsar [Intro].
	⛔ **Attenzione:** Nel caso in cui si selezioni la reimpostazione dei valori di default, tutte le impostazioni configurate verranno sostituite dai default di fabbrica premendo il tasto [Invio].
	⛔ 注意：「デフォルトの設定にリセット」コマンドを選択すると、現在のコンフィグレーションされた設定は、[Enter]を押したとき、工場出荷時の設定に変更されます。

Table 23 Console interface Main Menu options (continued)

Option	Description
Command Line Interface	Allows a properly authorized user to initiate a CLI management session. Refer to NCLI Configuration Guide for Ethernet Routing Switch 3510-24T, Software Release 4.0.3 for information on using the CLI.
Logout	Allows a user in a Telnet session or a user working at a password-protected console terminal to terminate the session.

IP Configuration/Setup Screen

The IP Configuration/Setup screen (Figure 55) allows you to set or modify the Ethernet Routing Switch 3510-24T IP configuration parameters. Data that you enter in the user-configurable fields takes effect as soon as you press [Enter].

To open the IP Configuration/Setup screen:

➨ Choose IP Configuration/Setup (or press i) from the main menu.

Figure 55 IP Configuration/Setup screen

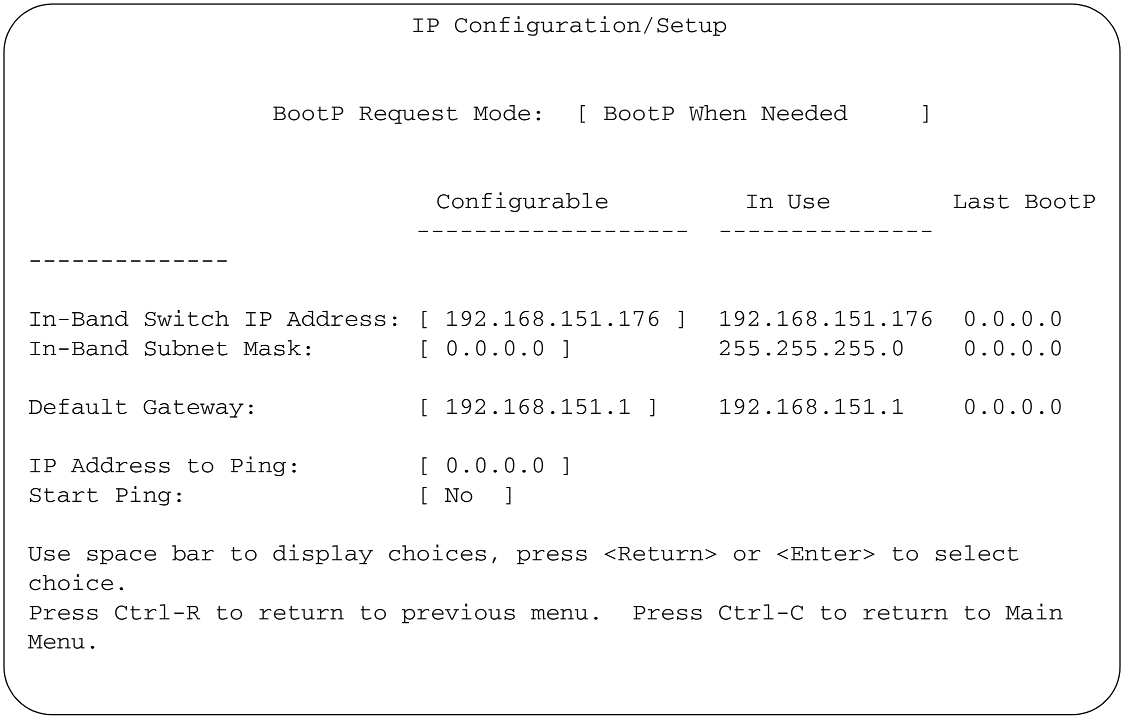

```
                          IP Configuration/Setup


                BootP Request Mode:  [ BootP When Needed     ]


                           Configurable         In Use         Last BootP
                          -------------------  ---------------
--------------

In-Band Switch IP Address: [ 192.168.151.176 ]  192.168.151.176  0.0.0.0
In-Band Subnet Mask:       [ 0.0.0.0 ]          255.255.255.0    0.0.0.0

Default Gateway:           [ 192.168.151.1 ]    192.168.151.1    0.0.0.0

IP Address to Ping:        [ 0.0.0.0 ]
Start Ping:                [ No  ]

Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 24 describes the IP Configuration/Setup screen fields.

Table 24 IP Configuration/Setup screen fields

Field	Description	
BootP Request Mode	One of four modes of operation for BootP. (See "Choosing a BootP Request Mode"" on page 152 for details about the four modes.)	
	Default Value	BootP When Needed
	Range	BootP Disabled, BootP When Needed, BootP Always, BootP or Last Address
Configurable	Column header for the user-configurable IP configuration fields in this screen.	
In Use	Column header for the read-only fields in this screen. The read-only data displayed in this column represents IP configuration that is currently in use.	
Last BootP	Column header for the read-only fields in this screen. The read-only data displayed in this column represents IP configuration obtained from the last BootP reply received.	
	Default Value	0.0.0.0 (no IP address assigned)
	Range	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, between 0 and 255, separated by a decimal point
In-Band Switch IP Address	The in-band IP address of the switch.	
	Default Value	0.0.0.0 (no IP address assigned)
	Range	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, between 0 and 255, separated by a decimal point
	Note: When the IP address is entered in the In-Band IP Address field, and the In-Band Subnet Mask field value is not present, the software provides an *in-use* default value for the In-Band Subnet Mask field that is based on the class of the IP address entered in the In-Band IP Address field.	

Table 24 IP Configuration/Setup screen fields (continued)

Field	Description	
In-Band Subnet Mask	The subnet address mask associated with the in-band IP address shown on the screen (see In-Band Switch IP Address field). Network routers use the subnet mask to determine the network or subnet address portion of a host's IP address. The bits in the IP address that contain the network address (including the subnet) are set to 1 in the address mask, and the bits that contain the host identifier are set to 0.	
	Default Value	0.0.0.0 (no subnet mask assigned)
	Range	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, between 0 and 255, separated by a decimal point
Default Gateway	The IP address of the default gateway.	
	Default Value	0.0.0.0 (no IP address assigned)
	Range	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, between 0 and 255, separated by a decimal point
IP Address to Ping	The IP address of the network device you want to ping.	
	Default Value	0.0.0.0 (no IP address assigned)
	Range	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, between 0 and 255, separated by a decimal point
Start Ping	Pings the selected network device when you choose Yes.	
	Default Value	No
	Range	No, Yes

➡ **Note:** The read-only fields in this screen are updated based on the BootP mode specified in the BootP Request Mode field. (See "Choosing a BootP Request Mode"" on page 152 for more information.)

Choosing a BootP Request Mode

The BootP Request Mode field in the IP Configuration screen allows you to choose a method the switch uses to broadcast BootP requests.

Following are the list of BootP Modes:

- BootP When Needed
- BootP Always
- BootP Disabled
- BootP or Last Address

Note: Whenever the switch is broadcasting BootP requests, the BootP process will eventually time out if a reply is not received. When the process times out, the BootP request mode automatically changes to BootP Disabled mode. To restart the BootP process, change the BootP request mode to any of the three following modes:

- BootP When Needed
- BootP Always
- BootP or Last Address.

BootP When Needed

Allows the switch to request an IP address if one has not already been set from the console terminal. When selected, this mode operates as follows:

- When the IP data is entered from the console terminal, the data becomes the in-use address of the switch and BootP requests are not broadcast. The switch can be managed using this in-band IP address.
- When the in-band IP address is not set from the console terminal, the switch broadcasts BootP requests until it receives a BootP reply containing an IP address. If the switch does not receive a BootP reply that contains an IP address, the switch cannot be managed in-band.

If an IP address is *not* currently in use, these actions take effect immediately. If an IP address *is* currently in use, these actions take effect only after the switch is reset or power cycled.

BootP When Needed is the default setting for the Ethernet Routing Switch 3510-24T.

BootP Always

Allows the switch to be managed only when configured with the IP address obtained from the BootP server. When selected, this mode operates as follows:

- The switch continues to broadcast BootP requests, regardless of whether an in-band IP address is set from the console terminal.
- If the switch receives a BootP reply that contains an in-band IP address, the switch uses this new in-band IP address.
- If the switch does not receive a BootP reply, the switch cannot be managed using the in-band IP address set from the console terminal.

If an IP address is *not* currently in use, these actions take effect immediately. If an IP address *is* currently in use, these actions take effect only after the switch is reset or power cycled.

BootP Disabled

Allows the switch to be managed only by using the IP address set from the console terminal. When selected, this mode operates as follows:

- The switch does not broadcast BootP requests, regardless of whether an IP address is set from the console terminal.
- The switch can be managed only by using the in-band switch IP address set from the console terminal.

These actions take effect after the switch is reset or power cycled, even if an IP address is not currently in use.

BootP or Last Address

Allows the switch to be managed even if a BootP server is not reachable. When selected, this mode operates as follows:

- When the IP data is entered from the console terminal, the data becomes the in-band address of the switch and BootP requests are not broadcast. The switch can be managed using this in-band IP address.
- When the in-band IP address is not set from the console terminal, the switch broadcasts BootP requests until it receives a BootP reply containing an in-band IP address. If the switch does not receive a BootP reply that contains an in-band IP address within 10 minutes, the switch uses the last in-band IP address it received from a BootP server. This IP information is displayed in the Last BootP column.

If an IP address is *not* currently in use, these actions take effect immediately. If an IP address *is* currently in use, these actions take effect only after the switch is reset or power cycled.

SNMP Configuration Screen

The SNMP Configuration Screen (Figure 56) allows you to set or modify the SNMP configuration parameters (for SNMP v1 and 2c; not for v3).

To open the SNMP Configuration screen:

➨ Choose SNMP Configuration (or press m) from the main menu.

Figure 56 SNMP Configuration screen

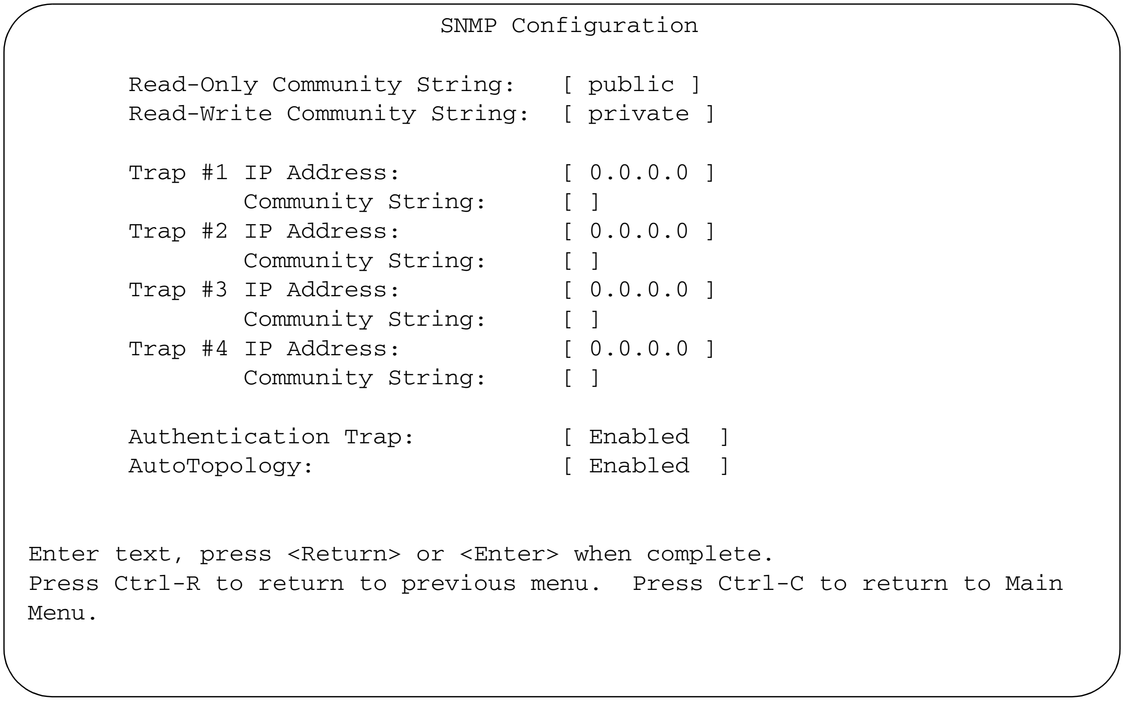

```
                             SNMP Configuration

       Read-Only Community String:   [ public ]
       Read-Write Community String:  [ private ]

       Trap #1 IP Address:           [ 0.0.0.0 ]
               Community String:     [ ]
       Trap #2 IP Address:           [ 0.0.0.0 ]
               Community String:     [ ]
       Trap #3 IP Address:           [ 0.0.0.0 ]
               Community String:     [ ]
       Trap #4 IP Address:           [ 0.0.0.0 ]
               Community String:     [ ]

       Authentication Trap:          [ Enabled  ]
       AutoTopology:                 [ Enabled  ]


Enter text, press <Return> or <Enter> when complete.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 25 describes the SNMP Configuration screen fields.

Table 25 SNMP Configuration screen fields

Field	Description	
Read-Only Community String	The community string used for in-band read-only SNMP operations.	
	Default Value	public
	Range	Any ASCII string of up to 32 printable characters
Read-Write Community String	The community string used for in-band read-write SNMP operations.	
	Default Value	private
	Range	Any ASCII string of up to 32 printable characters

Table 25 SNMP Configuration screen fields (continued)

Field	Description	
Trap #1 IP Address*	Number one of four trap IP addresses. Successive trap IP address fields are numbered 2, 3, and 4. Each trap address has an associated community string (see Community String).	
	Default Value	0.0.0.0 (no IP address assigned)
	Range	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point
Community String	The community string associated with one of the four trap IP addresses (see Trap #1 IP Address).	
	Default Value	Zero-length string
	Range	Any ASCII string of up to 32 printable characters
Authentication Trap	Determines whether a trap will be sent when there is an SNMP authentication failure.	
	Default Value	Enabled
	Range	Enabled, Disabled
Autotopology	Allows you to enable or disable the switch participation in Autotopology, which allows network topology mapping of other switches in your network.	
	Default Value	Enabled
	Range	Disabled, Enabled

* The Trap IP Address and Community String fields can be set using a MIB table (in a Nortel proprietary MIB). The status of the row in the MIB table can be set to Ignore. If the row status is set to Ignore, the fields appear to be set when viewed from the console terminal; however, no traps will be sent to that address until the row status is set to Valid.

System Characteristics Screen

The System Characteristics Screen (Figure 57) allows you to view system characteristics and contains three user-configurable fields: sysContact, sysName, and sysLocation.

To open the System Characteristics screen:

➨ Choose System Characteristics (or press s) from the main menu.

Figure 57 System Characteristics screen

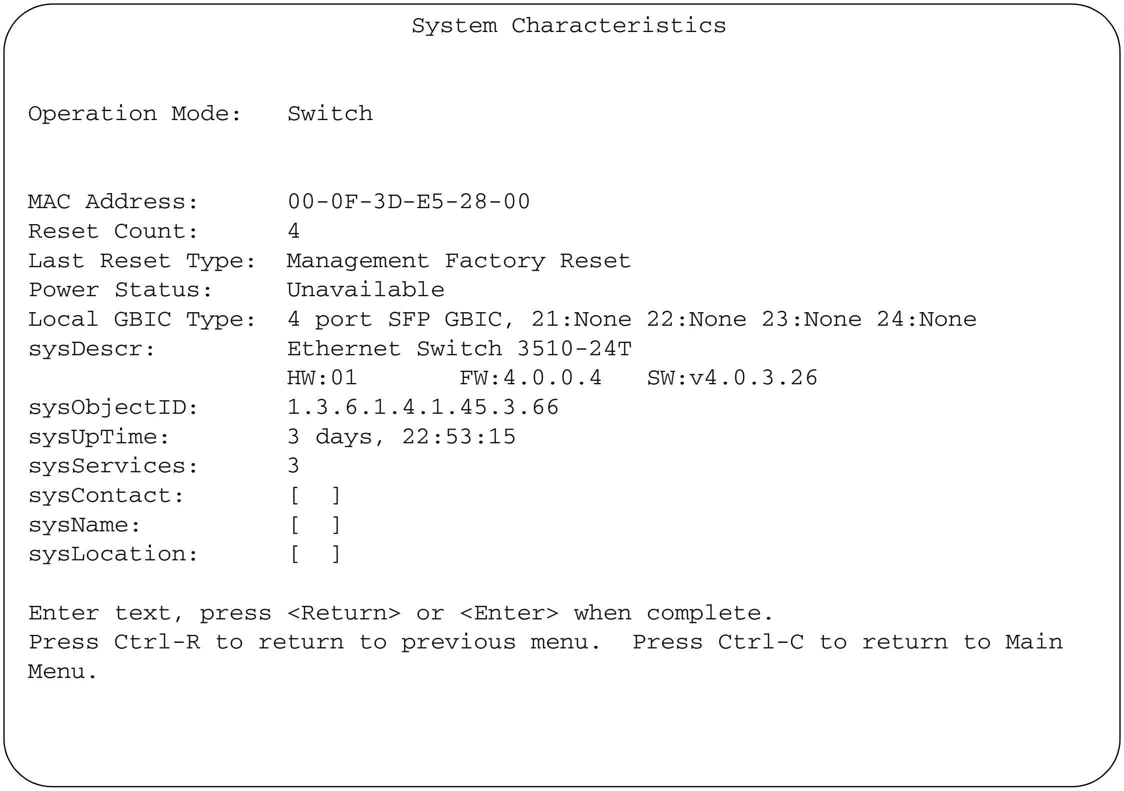

```
                         System Characteristics


Operation Mode:   Switch


MAC Address:      00-0F-3D-E5-28-00
Reset Count:      4
Last Reset Type:  Management Factory Reset
Power Status:     Unavailable
Local GBIC Type:  4 port SFP GBIC, 21:None 22:None 23:None 24:None
sysDescr:         Ethernet Switch 3510-24T
                  HW:01      FW:4.0.0.4   SW:v4.0.3.26
sysObjectID:      1.3.6.1.4.1.45.3.66
sysUpTime:        3 days, 22:53:15
sysServices:      3
sysContact:       [  ]
sysName:          [  ]
sysLocation:      [  ]

Enter text, press <Return> or <Enter> when complete.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 26 describes the System Characteristics Screen fields.

Table 26 System Characteristics Screen fields

Field	Description
Operation Mode	Read-only field that indicates the operation mode of the Switch. The default value for the Operation mode is Switch only.
MAC Address	The MAC address of the switch.
Reset Count	A read-only field that indicates the number of resets since the operational firmware was first loaded on the switch. Default Value 1 Range 0 to 232 -1 (4,294,967,295)

Table 26 System Characteristics Screen fields (continued)

Field	**Description**	
Last Reset Type	A read-only field that indicates the last type of reset.	
	Default Value	Power Cycle
	Range	Power Cycle, Software Download, Management Reset, Management Factory Reset
Power Status	A read-only field that indicates the current power source.	
	Default Value	Primary Power
	Range	Primary Power
Local GBIC Type	A read-only field that indicates the SFP GBIC type that is configured in this unit.	
sysDescr	A read-only field that specifies hardware and software versions.	
sysObjectID	A read-only field that provides a unique identification of the switch, which contains the vendor's private enterprise number.	
sysUpTime	A read-only field that shows the length of time since the last reset. Note that this field is updated when the screen is redisplayed.	
sysServices	A read-only field that indicates the switch's physical and data link layer functionality.	
sysContact	The name and phone number of the person responsible for the switch.	
	Default Value	Zero-length string
	Range	Any ASCII string of up to 56 printable characters*
sysName	A name that uniquely identifies the switch.	
	Default Value	Zero-length string
	Range	Any ASCII string of up to 56 printable characters*
sysLocation	The physical location of the switch.	
	Default Value	Zero-length string
	Range	Any ASCII string of up to 56 printable characters

* Although this field can be set to up to 255 characters from a Network Management Station (NMS), only 56 characters are displayed on the console terminal.

Switch Configuration Menu Screen

The Switch Configuration Menu Screen (Figure 58) allows you to set or modify the switch configuration.

Choose Switch Configuration (or press w) from the main menu to open the Switch Configuration Menu screen (Table 27).

Figure 58 Switch Configuration Menu screen

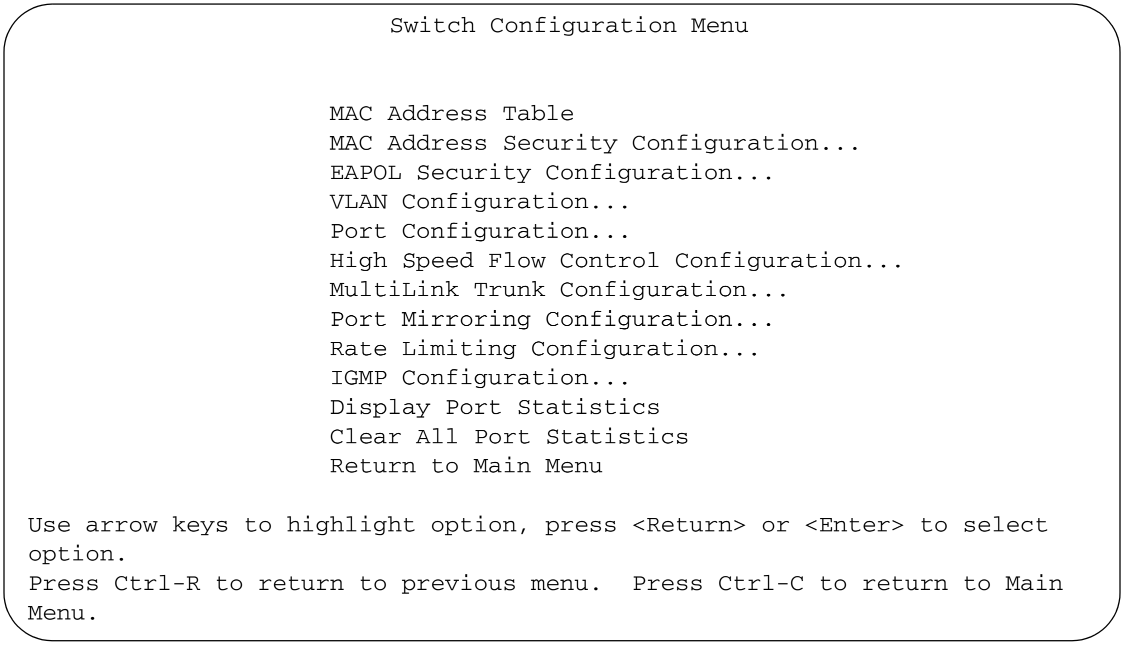

```
                         Switch Configuration Menu


                    MAC Address Table
                    MAC Address Security Configuration...
                    EAPOL Security Configuration...
                    VLAN Configuration...
                    Port Configuration...
                    High Speed Flow Control Configuration...
                    MultiLink Trunk Configuration...
                    Port Mirroring Configuration...
                    Rate Limiting Configuration...
                    IGMP Configuration...
                    Display Port Statistics
                    Clear All Port Statistics
                    Return to Main Menu

Use arrow keys to highlight option, press <Return> or <Enter> to select
option.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 27 describes the Switch Configuration Menu Screen options.

Table 27 Switch Configuration Menu Screen options

Option	Description
MAC Address Table	Displays the MAC Address Table screen (see "MAC Address Table Screen"" on page 161). This screen allows you to view all MAC addresses and their associated port or trunk that the switch has learned, or to search for a particular MAC address (to see if the switch has learned the address).
MAC Address Security Configuration...	Displays the MAC Address Security Configuration menu (see "MAC Address Security Configuration Menu Screen" on page 162). This screen allows you to set up the MAC address security feature and provides the following options: MAC Address Security Configuration, MAC Address Security Port Configuration, MAC Address Security Port Lists, and MAC Address Security Table. This menu allows you to enable and disable security features on the port and trunk levels.
EAPOL Security Configuration...	Displays the EAPOL Security Configuration menu (see "EAPOL Security Configuration Screen"" on page 173). This screen allows you to set up Extensible Authentication Protocol over LAN (EAPOL)-based security.

Table 27 Switch Configuration Menu Screen options (continued)

Option	**Description**
VLAN Configuration...	Displays the Configuration Menu (see "VLAN Configuration Menu Screen"" on page 177). This menu provides the following options: VLAN Configuration, MAC Addresses for MAC-SA Based VLAN, VLAN Port Configuration, and VLAN Display by Port. This menu allows you to create and modify VLANs and to enable the automatic PVID feature.
Port Configuration...	Displays the Port Configuration screen (see "Port Configuration Screen"" on page 188). This screen allows you to configure a specific switch port and all switch ports.
High Speed Flow Control Configuration...	Displays the High Speed Flow Control Configuration screen (see "High Speed Flow Control Configuration Screen"" on page 191).
MultiLink Trunk Configuration...	Displays the MultiLink Trunk Configuration Menu (see "MultiLink Trunk Configuration Menu screen""on page 193). This menu provides the following options: MultiLink Trunk Configuration and MultiLink Trunk Utilization. This menu allows you to create and modify trunks, and to monitor the bandwidth utilization of configured trunks.
Port Mirroring Configuration...	Displays the Port Mirroring Configuration screen (see "Port Mirroring Configuration Screen""on page 199). This screen allows you to designate a single switch port as a traffic monitor for up to two specified ports or addresses.
Rate Limiting Configuration...	Displays the Rate Limiting Configuration screen (see "Rate Limiting Configuration Screen""on page 202). This screen allows you to limit the forwarding rate of broadcast and multicast packets at ingress.
IGMP Configuration...	Displays the IGMP Configuration screen (see "IGMP Configuration Screen"" on page 206). This screen allows you to optimize multicast traffic by setting up IGMP port memberships that filter multicast on a per port basis (see Chapter 1 for more information about this feature).
Display Port Statistics	Displays the Port Statistics screen (see "Port Statistics Screen"" on page 210). This screen allows you to view detailed information about any switch port.
Clear All Port Statistics	Allows you to clear all port statistics. This option is followed by screen prompts that precede a choice of the actions: • If the switch is operating *standalone*, choose one of the following: • Yes, to clear all port statistics for all switch ports • No, to abort the option

MAC Address Table Screen

The MAC Address Table Screen (Figure 59) allows you to view MAC addresses that the switch has discovered or to search for a specific MAC address.

➨ Choose MAC Address Table (or press m) from the Switch Configuration Menu screen to open the MAC Address Table screen (Figure 59).

Figure 59 MAC Address Table Screen

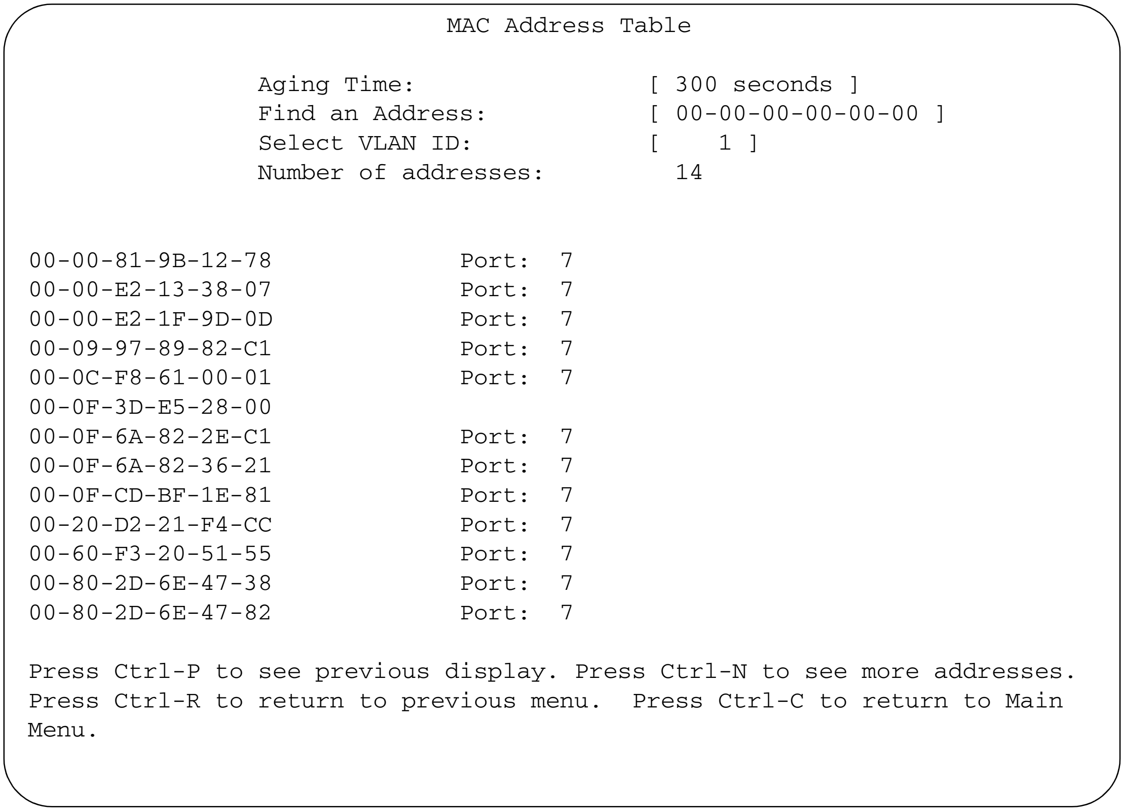

```
                             MAC Address Table

                Aging Time:                [ 300 seconds ]
                Find an Address:           [ 00-00-00-00-00-00 ]
                Select VLAN ID:            [    1 ]
                Number of addresses:         14


00-00-81-9B-12-78             Port:  7
00-00-E2-13-38-07             Port:  7
00-00-E2-1F-9D-0D             Port:  7
00-09-97-89-82-C1             Port:  7
00-0C-F8-61-00-01             Port:  7
00-0F-3D-E5-28-00
00-0F-6A-82-2E-C1             Port:  7
00-0F-6A-82-36-21             Port:  7
00-0F-CD-BF-1E-81             Port:  7
00-20-D2-21-F4-CC             Port:  7
00-60-F3-20-51-55             Port:  7
00-80-2D-6E-47-38             Port:  7
00-80-2D-6E-47-82             Port:  7

Press Ctrl-P to see previous display. Press Ctrl-N to see more addresses.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 28 describes the MAC Address Table screen fields.

Table 28 MAC Address Table screen fields

Field	Description
Aging Time	Specifies how long a learned MAC address remains in the switch's forwarding database. If an entry is inactive for a period of time that exceeds the specified aging time, the address is removed. Default Value 300 seconds Range 10 to 1,000,000 seconds
Find an Address	Allows the user to search for a specific MAC address. Default Value 00-00-00-00-00-00 (no MAC address assigned) Range 00-00-00-00-00-00 to FF-FF-FF-FF-FF-FF
Select VLAN ID	Enter the VLAN ID number you want to display the MAC addresses for. Default Value 1 Range 1-4094
Number of addresses	Displays the total number of MAC addresses currently learned by the specified VLAN. This number updates dynamically when you press [Ctrl]-P or [Ctrl]-N to scroll through the list.

MAC Address Security Configuration Menu Screen

The MAC Address Security Configuration Menu screen (Figure 60) allows you to specify a range of system responses to unauthorized network access to your switch. The system response can range from sending a trap to disabling the port. The network access control is based on the MAC sources addresses of the authorized stations. You can specify a list of up to 448 MAC source addresses that are authorized to access the switch. You can also specify the ports that each MAC source address is allowed to access.

The port access options include:

- NONE
- ALL
- Single or multiple ports that are specified in a list, for example, 1-4, 6,9, etc. You must also include the MAC address of any router connected to any secure ports.

You can configure the switch to drop all packets with specified MAC destination addresses (DA). You can enter up to 10 specific MAC DAs you want filtered. The packet with the specified MAC DA will be dropped regardless of the ingress port, source address (SA) intrusion, or VLAN membership.

➡ **Note:** You must use either the Web-based management system or the CLI to configure MAC DA filtering.

When the switch software detects a security violation on the specified MAC SAs, the response can be to send a trap, turn on the Destination Address (DA) filtering that is based on SA filtering, disable the specific port, or any combination of these three options. The packet will be blocked on an invalid MAC source address.

To open the MAC Address Security Configuration Screen:

➨ Choose MAC Address Security Configuration from the Switch Configuration Menu.

Figure 60 MAC Address Security Configuration Menu screen

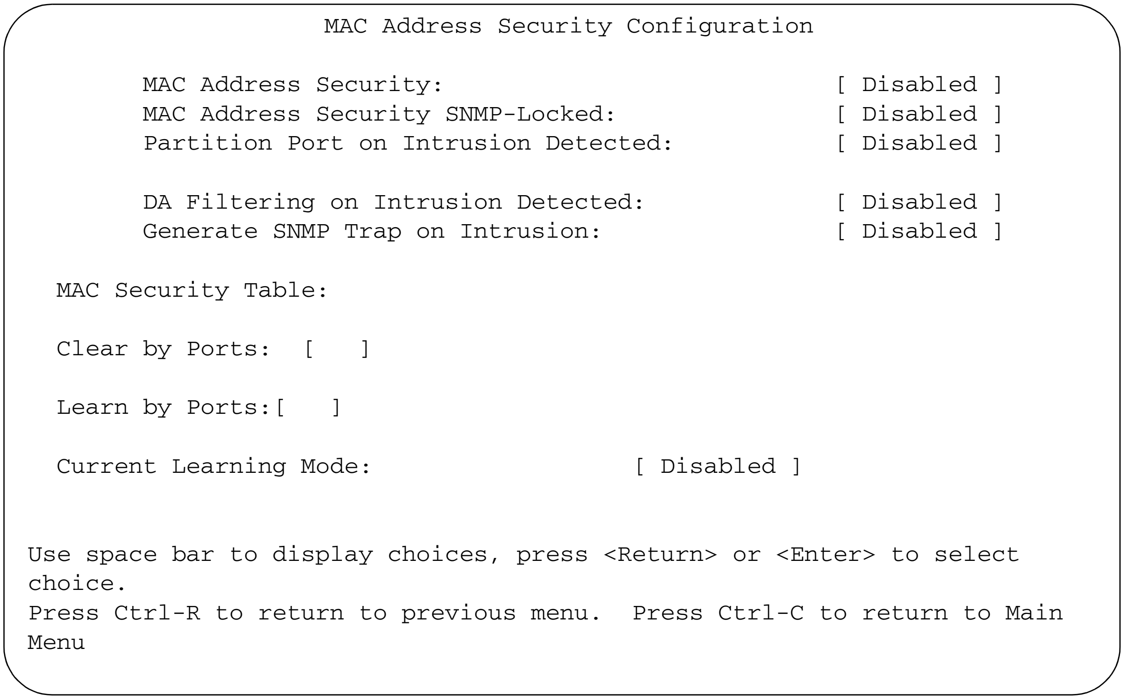

```
                     MAC Address Security Configuration

        MAC Address Security:                           [ Disabled ]
        MAC Address Security SNMP-Locked:               [ Disabled ]
        Partition Port on Intrusion Detected:           [ Disabled ]

        DA Filtering on Intrusion Detected:             [ Disabled ]
        Generate SNMP Trap on Intrusion:                [ Disabled ]

  MAC Security Table:

  Clear by Ports:  [   ]

  Learn by Ports:[   ]

  Current Learning Mode:                  [ Disabled ]


Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu
```

Table 29 describes the MAC Address Security Configuration Menu options.

Table 29 MAC Address Security Configuration Menu Options

Option	Description
MAC Address Security Configuration...	Displays the MAC Address Security Configuration screen (see "MAC Address Security Configuration Menu Screen"). This screen allows you to Enable or Disable the MAC Address Security feature.
MAC Address Security Port Configuration...	Displays the MAC Address Security Port Configuration screen (see "MAC Address Security Port Configuration Screen""). This screen allows you to Enable or Disable MAC Security for each port.
MAC Address Security Port Lists...	Displays the MAC Address Security Port Lists screen (see "MAC Address Security Port Lists Screens"). This screen allows you to create port lists that can be used as an *allowed source port list* for a MAC address in the MAC Address Security Table screen.
MAC Address Security Table...	Displays the MAC Address Security Table screen (see "MAC Address Security Table Screens"). This screen allows you to specify the MAC addresses that are allowed to access the switch.

MAC Address Security Configuration Screen

The MAC Address Security Configuration screen (Figure 61) allows you to enable or disable the MAC address security feature and to specify the appropriate system responses to any unauthorized network access to your switch.

➨ Choose MAC Address Security Configuration from the MAC Address Security Configuration Menu to open the MAC Address Security Configuration screen.

Figure 61 MAC Address Security Configuration screen

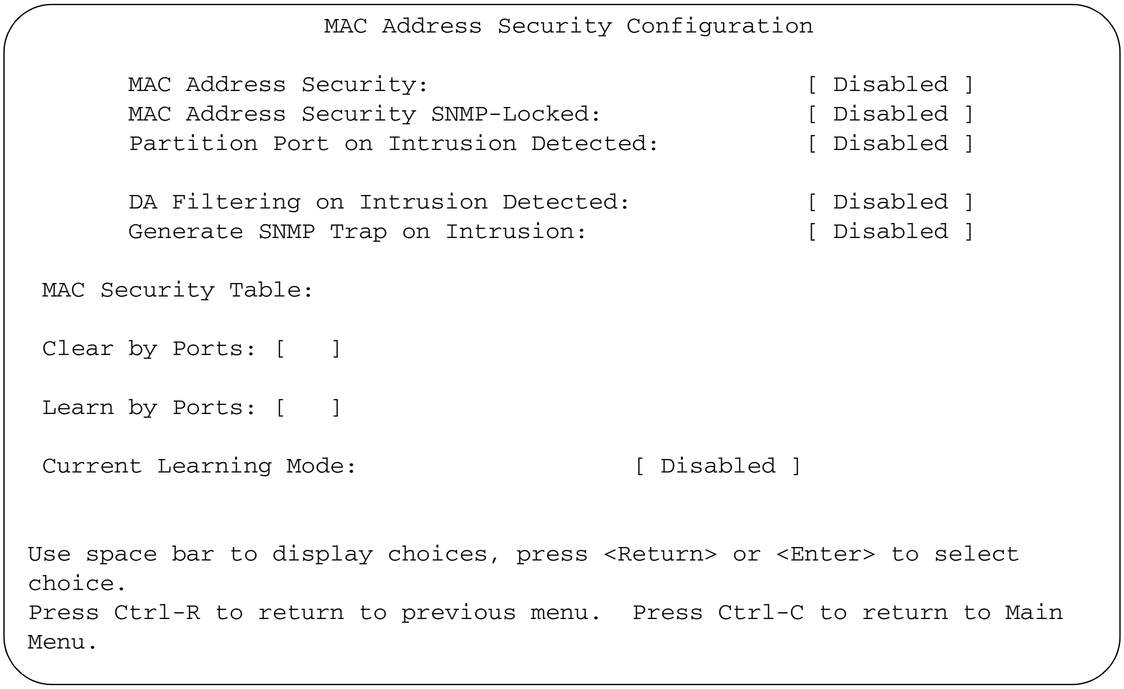

```
                     MAC Address Security Configuration

       MAC Address Security:                                 [ Disabled ]
       MAC Address Security SNMP-Locked:                     [ Disabled ]
       Partition Port on Intrusion Detected:                 [ Disabled ]

       DA Filtering on Intrusion Detected:                   [ Disabled ]
       Generate SNMP Trap on Intrusion:                      [ Disabled ]

 MAC Security Table:

 Clear by Ports: [   ]

 Learn by Ports: [   ]

 Current Learning Mode:                    [ Disabled ]


Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 30 describes the MAC Address Security Configuration screen fields.

Table 30 MAC Address Security Configuration fields

Field	Description
MAC Address Security	When this field is set to enabled, the software checks source MAC addresses of packets that arrive on secure ports against MAC addresses listed in the MAC Address Security Table for allowed membership. If the software detects a source MAC address that is not an allowed member, the software registers a MAC intrusion event. Default Disabled Range Disabled, Enabled
MAC Address Security SNMP-Locked	When this field is set to enabled, the MAC address security screens cannot be modified using SNMP (SNMP includes the DM management system). Default Disabled Range Disabled, Enabled
Partition Port on Intrusion Detected	This field value determines how the switch reacts to an intrusion event. When an intrusion even is detected (see MAC Address Security field description) the specified switch port is set to Disabled (partitioned from other switch ports). When the field is set to: • Disabled - the port remains enabled, even if an intrusion event is detected. • Enabled - the port becomes disabled, then automatically resets to enabled depending on the value set in the Partition Time field. • Forever - the port becomes disabled, and remains disabled (partitioned). The Partition Time field cannot be used to automatically to reset the port to Enabled if you set this field to Forever. You can always manually set the port’s status field to enabled using the Port Configuration screen (see “Port Configuration Screen”). Default Disabled Range Disabled, Enabled, Forever
Partition Time	This field appears only when the Partition Port on Intrusion Detected field is set to enabled. This field determines the length of time a partitioned port remains disabled. This field is not operational when the Partition Port on Intrusion Detected field is set to Forever. Default 1 second (the value 0 indicates forever) Range 0-65535 seconds

Table 30 MAC Address Security Configuration fields (continued)

Field	Description
DA Filtering on Intrusion Detected	When set to enabled, this field isolates the intruding node by filtering (discarding) packets sent to that MAC address. Default Disabled Range Disabled, Enabled
Generate SNMP Trap on Intrusion	When set to enabled and a MAC intrusion event is detected, the software issues an SNMP trap message to all registered SNMP trap addresses (see "SNMP Configuration Screen"). Default Disabled Range Disabled, Enabled
Clear by Ports	This field clears the specified port (or ports) that are listed in the Allowed Source Port(s) field of the MAC Address Security Table screen (see "MAC Address Security Table Screens"). When you specify a port (or ports) to be cleared using this field, the specific port (or ports) will be cleared for each of the entries listed in the MAC Address Security Table. If you totally clear the allowed Source Port(s) field (leaving a blank field) for an entry, the associated MAC address for that entry is also cleared. Default NONE Range NONE, ALL, a port number list (for example,1,6 etc)
Learn by Ports	All source MAC addresses of any packets received on the specified port (or ports) are added to the MAC Security Table when the Current Learning Mode field is set to Enabled. You cannot include any of the port values you have chosen for the secure ports field. Default NONE Range NONE, ALL, a port number list (for example, 1,6 etc)
Current Learning Mode	Indicates the current learning mode for the switch ports. When this field is set to Learning in Progress, all source MAC addresses of any packets received on the specified port (or ports) are added to the MAC Security Table (maximum of 448 MAC address entries allowed). If you exceed the limit of 448 entries, the system prompts you with an alert message. Default Disabled Range Enabled, Disabled

MAC Address Security Port Configuration Screen

The MAC Address Security Port Configuration Screens (Figure 62 and Figure 63) allow you to set or modify your MAC address port security configuration on a per port basis.

To open the MAC Address Security Port Configuration Screen:

➨ Choose MAC Address Security Port Configuration from the MAC Address Security Configuration Menu.

Figure 62 MAC Security Port Configuration screen (1 of 2)

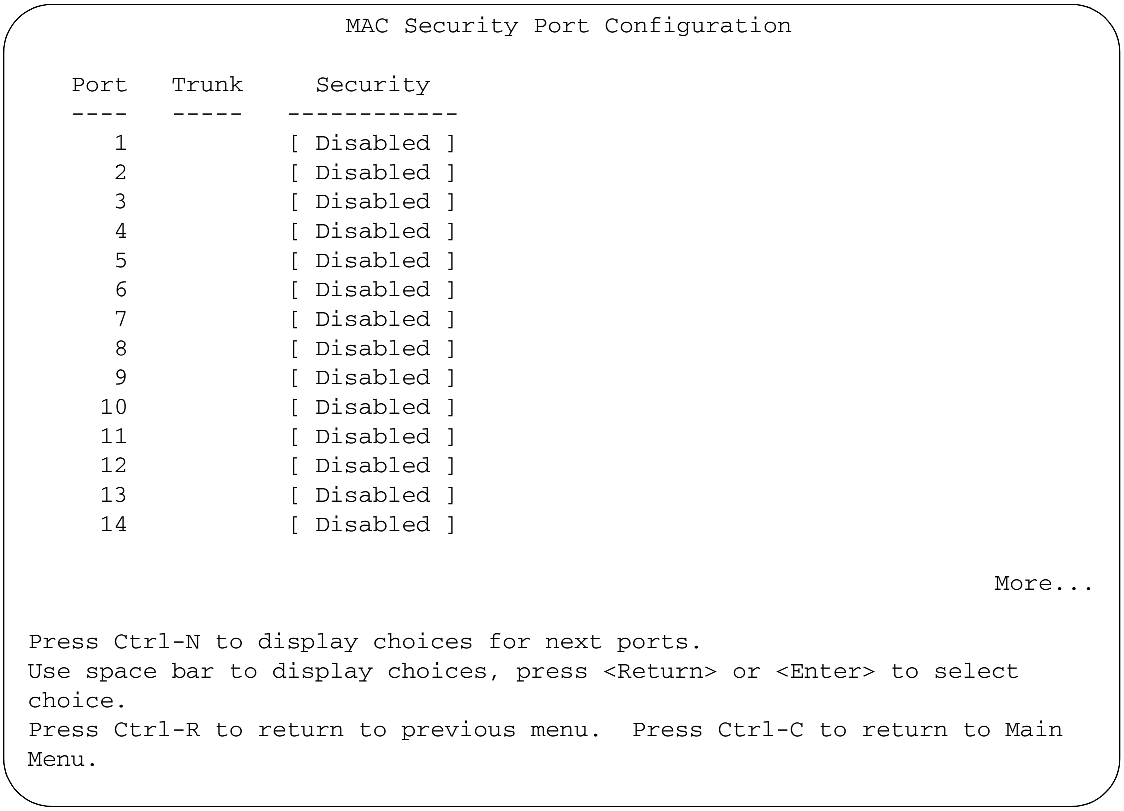

```
                        MAC Security Port Configuration

    Port    Trunk      Security
    ----    -----    ------------
       1             [ Disabled ]
       2             [ Disabled ]
       3             [ Disabled ]
       4             [ Disabled ]
       5             [ Disabled ]
       6             [ Disabled ]
       7             [ Disabled ]
       8             [ Disabled ]
       9             [ Disabled ]
      10             [ Disabled ]
      11             [ Disabled ]
      12             [ Disabled ]
      13             [ Disabled ]
      14             [ Disabled ]

                                                                   More...

Press Ctrl-N to display choices for next ports.
Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Figure 63 MAC Security Port Configuration screen (2 of 2)

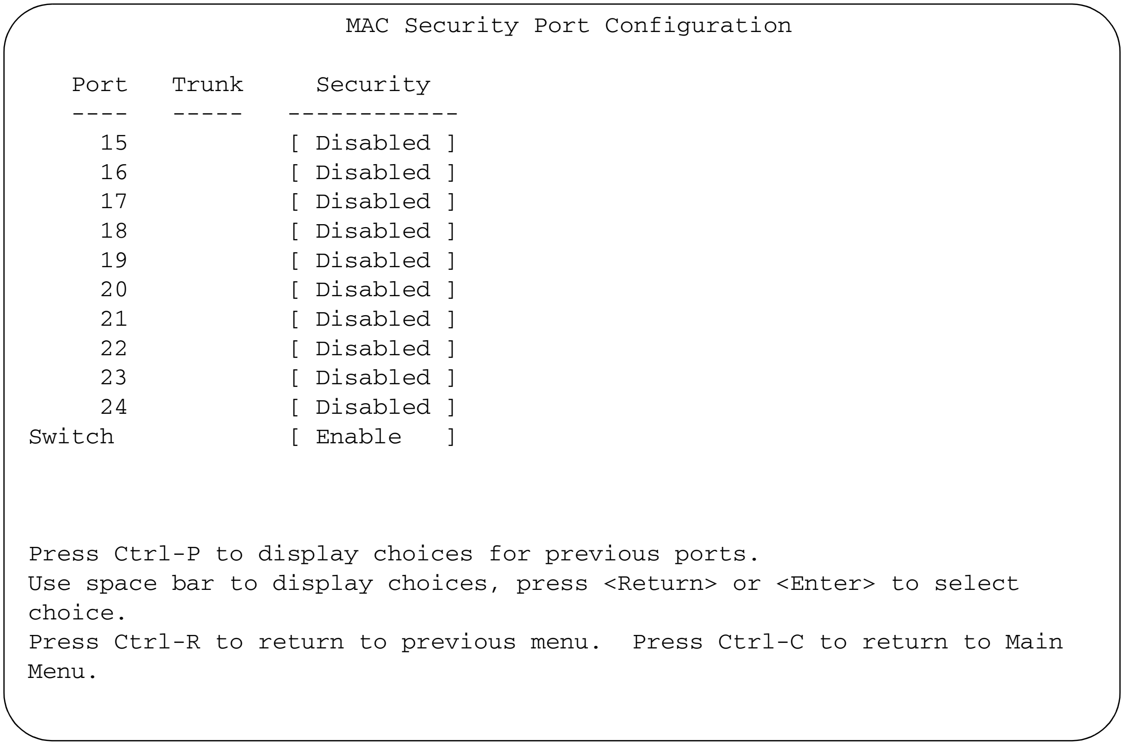

```
                          MAC Security Port Configuration

    Port   Trunk     Security
    ----   -----   ------------
      15           [ Disabled ]
      16           [ Disabled ]
      17           [ Disabled ]
      18           [ Disabled ]
      19           [ Disabled ]
      20           [ Disabled ]
      21           [ Disabled ]
      22           [ Disabled ]
      23           [ Disabled ]
      24           [ Disabled ]
 Switch            [ Enable   ]



 Press Ctrl-P to display choices for previous ports.
 Use space bar to display choices, press <Return> or <Enter> to select
 choice.
 Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
 Menu.
```

Table 31 describes the MAC Security Port Configuration screen fields.

Table 31 MAC Security Port Configuration screen fields

Field	Description
Port	Displays a numbered port list.
Trunk	Displays the trunk number if the port is a member of that trunk. Default blank field
Security	This field value determines whether or not security is enabled or disabled on the port level or switch level. Default Disabled Range Disabled, Enabled

MAC Address Security Port Lists Screens

The MAC Address Security Port Lists screens allow you to create port lists that can be used as *allowed source port lists* for a specified MAC address in the MAC Address Security Table screen. You can create as many as 32 port lists, using up to five MAC Address Security Port Lists screens (see Figure 64).

Figure 64 MAC Address Security Port Lists screens

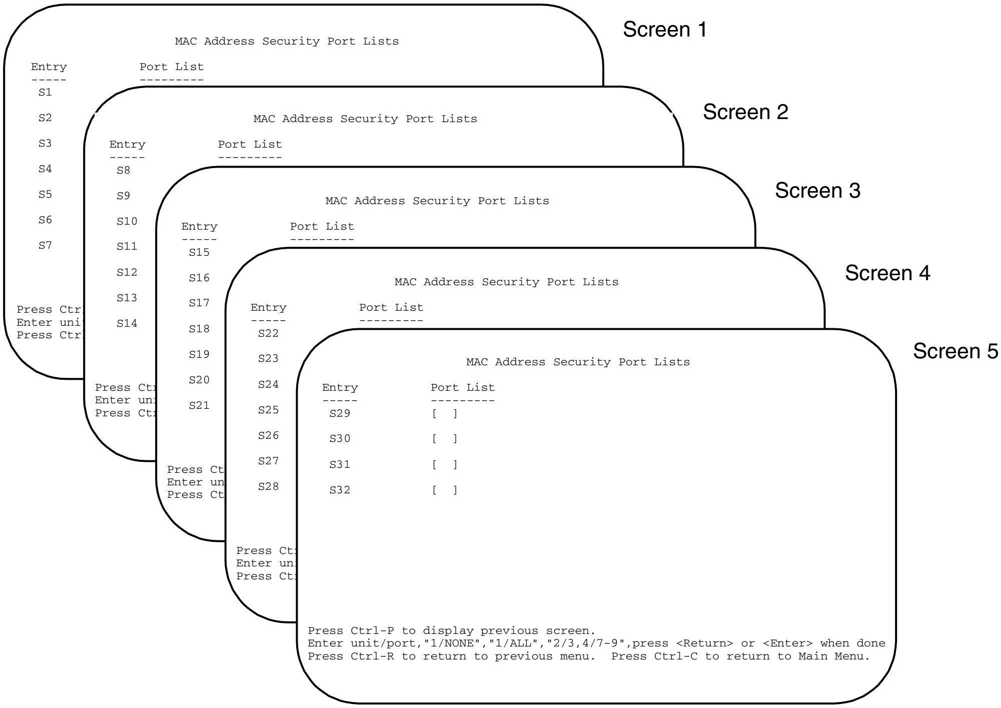

```
                    MAC Address Security Port Lists

 Entry          Port List
 -----          ---------
  S1
  S2
  S3
  S4
  S5
  S6
  S7

Press Ctr
Enter uni
Press Ctr
```

```
                    MAC Address Security Port Lists

 Entry          Port List
 -----          ---------
  S8
  S9
  S10
  S11
  S12
  S13
  S14

Press Ct
Enter un
Press Ct
```

```
                    MAC Address Security Port Lists

 Entry          Port List
 -----          ---------
  S15
  S16
  S17
  S18
  S19
  S20
  S21

Press Ct
Enter un
Press Ct
```

```
                    MAC Address Security Port Lists

 Entry          Port List
 -----          ---------
  S22
  S23
  S24
  S25
  S26
  S27
  S28

Press Ct
Enter un
Press Ct
```

```
                    MAC Address Security Port Lists

 Entry          Port List
 -----          ---------
  S29           [  ]
  S30           [  ]
  S31           [  ]
  S32           [  ]

Press Ctrl-P to display previous screen.
Enter unit/port,"1/NONE","1/ALL","2/3,4/7-9",press <Return> or <Enter> when done
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

To open the MAC Address Security Lists screen:

➨ Choose MAC Address Security Lists from the MAC Address Security Configuration Menu.

The options for allowed port access include: NONE, ALL, and ports that are specified in a list (for example, 1,6 etc.). Refer to Port List syntax for more information.

MAC Address Security Table Screens

The MAC Address Security Table screens allow you specify the ports that each MAC address is allowed to access. You must also include the MAC addresses of any routers that are connected to any secure ports.

There are 16 available MAC Address Security Table screens that you can use to create up to 448 MAC address entries (28 per screen).

➨ Choose MAC Address Security Table from the MAC Address Security Configuration Menu to open the MAC Address Security Table screen (Figure 65).

Figure 65 MAC Address Security Table screen

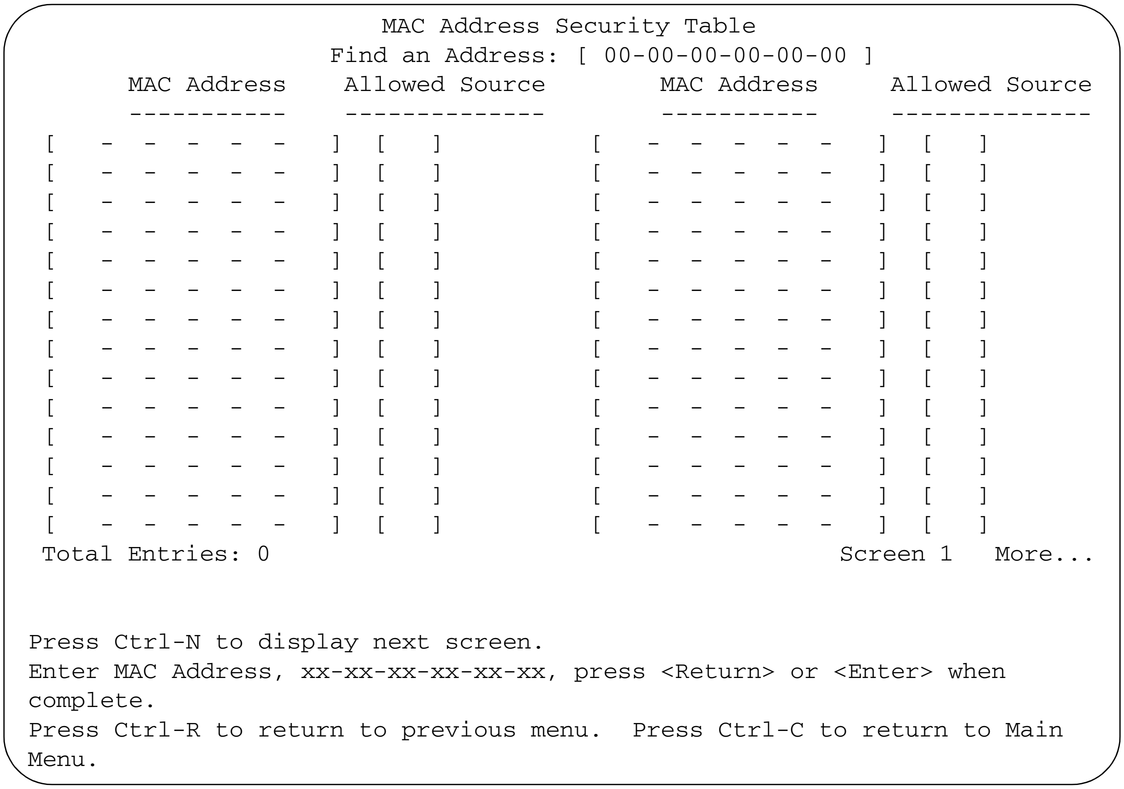

```
                          MAC Address Security Table
                     Find an Address: [ 00-00-00-00-00-00 ]
        MAC Address      Allowed Source           MAC Address      Allowed Source
        -----------      --------------           -----------      --------------
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  [    -  -  -  -  -   ]  [   ]          [    -  -  -  -  -   ]  [   ]
  Total Entries: 0                                         Screen 1   More...


 Press Ctrl-N to display next screen.
 Enter MAC Address, xx-xx-xx-xx-xx-xx, press <Return> or <Enter> when
 complete.
 Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
 Menu.
```

Table 32 describes the MAC Address Security Table screen fields.

Table 32 MAC Address Security Table Screen Fields

Field	Description
Find an Address	Allows you to search for a specific MAC address that is used in any of the MAC Address Security Table screens.
MAC Address	Allows you to specify up to 448 MAC addresses that are authorized to access the switch. You can specify the ports that each MAC address is allowed to access using the Allowed Source field (see next field description). The specified MAC address does not take effect until the Allowed Source field is set to some value (a single port number or a port list value that you previously configured in the MAC Address Security Port Lists screen). You can clear an existing MAC address field by entering zero (0) in the field and pressing [Enter]. Default - - - - - (no address assigned) Range A range of 6 Hex Octets, separated by dashes (multicast* and broadcast addresses are not allowed).
Allowed Source	Allows you to specify the ports that each MAC address is allowed to access. The options for the Allowed Source field include a single port number or a list value that you have previously configured in the MAC Address Security Port Lists screen (for example 3,4,6, S1, S5 ect). Default - (Blank field) Range A single port or a port list value (for example, 1/3, 1/6, 3/4, S1, S5, etc.).
Total Entries	Displays the total number of entries on the entire table. Note: When the table is full, the system returns the following message: `Table full`.

* Multicast address -- Note that the first octet of any multicast address will always be an odd number.

EAPOL Security Configuration Screen

The EAPOL Security Configuration screen (Figure 66) allows you to selectively limit access to the switch based on an authentication mechanism that uses Extensible Authentication Protocol (EAP) to exchange authentication information between the switch and an authentication server.

Note: Before you use the EAPOL Security Configuration screen, you must configure your Primary RADIUS Server and RADIUS Shared Secret.

You will also need to set up specific user accounts on your RADIUS server:

- User names
- Passwords
- VLAN IDs
- Port priority

You can set up these parameters directly on your RADIUS server. For detailed instructions about configuring your RADIUS server, refer to your RADIUS server documentation.

Note: Do not enable EAPOL security on the switch port that is connected to the RADIUS server.

To open the EAPOL Security Configuration screen:

➨ Choose EAPOL Security Configuration (or press e) from the Switch Configuration Menu.

Figure 66 EAPOL Security Configuration screen

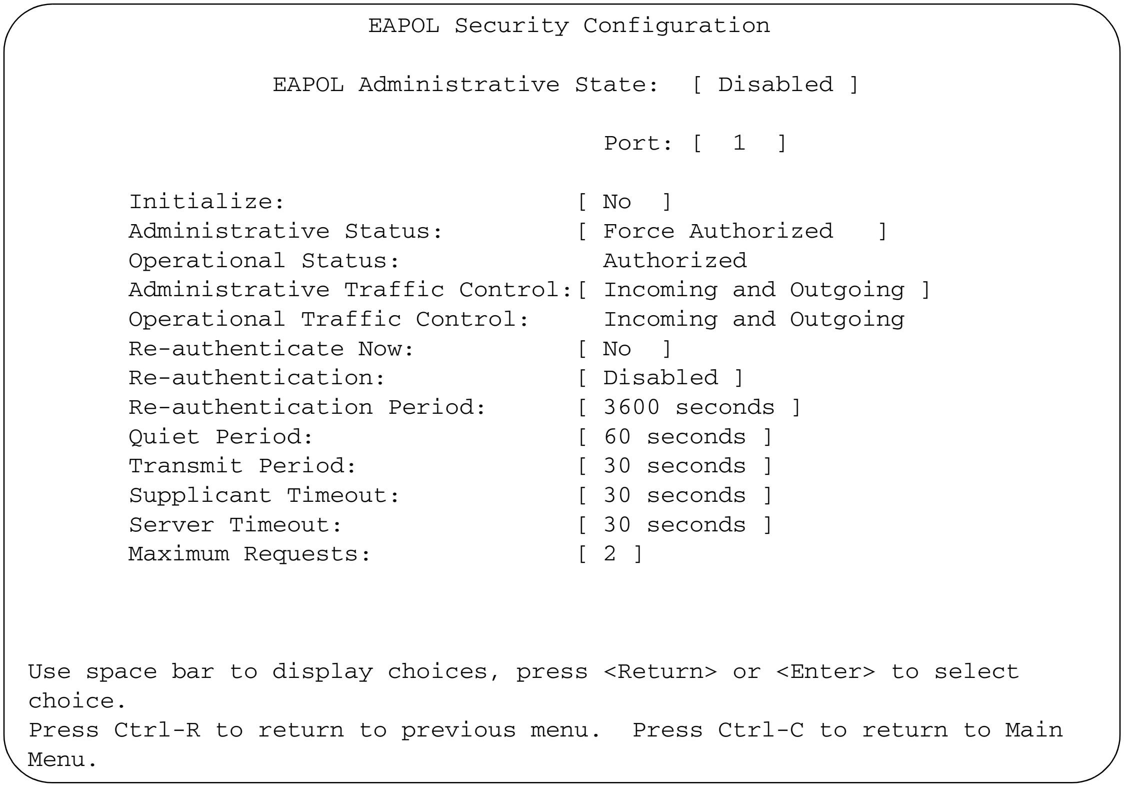

```
                        EAPOL Security Configuration

                 EAPOL Administrative State:  [ Disabled ]

                                        Port: [  1  ]

       Initialize:                    [ No  ]
       Administrative Status:         [ Force Authorized   ]
       Operational Status:              Authorized
       Administrative Traffic Control:[ Incoming and Outgoing ]
       Operational Traffic Control:     Incoming and Outgoing
       Re-authenticate Now:           [ No  ]
       Re-authentication:             [ Disabled ]
       Re-authentication Period:      [ 3600 seconds ]
       Quiet Period:                  [ 60 seconds ]
       Transmit Period:               [ 30 seconds ]
       Supplicant Timeout:            [ 30 seconds ]
       Server Timeout:                [ 30 seconds ]
       Maximum Requests:              [ 2 ]



Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 33 describes the EAPOL Security Configuration screen options.

Table 33 EAPOL security configuration screen options

Option	Description
EAPOL Administrative State	Allows you to enable or disable EAPOL for your switch. When this field is set to disabled (the default state), the Operational Status for all of the switch ports is set to Authorized (no security restriction). Default Disabled Range Disabled, Enabled

Table 33 EAPOL security configuration screen options (continued)

Option	**Description**
Port	Allows you to select a specified unit's (see preceding Unit field) port number to view or configure. To view or configure another port, type its port number and press [Enter], or press the spacebar to toggle the port numbers. If you set this field value to All, other screen field values you modify apply to *all* ports for the specified unit.The All value is also useful when you want to apply modified field values to most of, but not all of, your switch's ports. Default 1
	Range 1 to 24,ALL; 1 to 48, ALL
Initialize	Allows you to activate EAPOL authentication for the specified unit/port. Default No Range No,Yes
Administrative Status	Allows you to set the EAPOL authorization status for the specified unit/port. Default Force Authorized Range Force Authorized,Force Unauthorized,Auto • Force Authorized means the specified unit/port authorization status is *always* authorized. • Force Unauthorized means the specified unit/port authorization status is *always* Unauthorized. • Auto means the specified unit/port authorization status depends on the EAP authentication results.
Operational Status	A read-only field that shows the current authorization status for the specified unit/port. This read-only field does not appear when the Unit/Port field value is set to All. Default Authorized Range Authorized,Unauthorized
Administrative Traffic Control	Allows you to choose whether EAPOL authentication is set for incoming and outgoing traffic or for incoming traffic only. For example, if you set specified unit/port field value to Incoming and Outgoing, and EAPOL authentication fails,then both incoming and outgoing traffic on the specified unit/port is blocked. Default Incoming and Outgoing Range Incoming and Outgoing,Incoming Only
Operational Traffic Control	A read-only field that indicates the current administrative traffic control configuration for the specified unit/port (see preceding field description). This read-only field does not appear when the Port field value is set to All. Default Incoming and Outgoing Range Incoming and Outgoing,Incoming Only

Table 33 EAPOL security configuration screen options (continued)

Option	Description	
Re-authenticate Now	Allows you to activate EAPOL authentication for the specified unit/port immediately, without waiting for the Re-Authentication Period to expire.	
	Default	No
	Range	No,Yes
Re-authentication	Allows you to repeat EAPOL authentication for the specified unit/port according to the time interval value configured in the Re-Authentication Period field (see next field description).	
	Default	Enabled
	Range	Enabled,Disabled
Re-authentication Period	When the Re-Authentication field value (see preceding field) is set to enabled, this field allows you to specify the time period between successive EAPOL authentications for the specified unit/port.	
	Default	3600 seconds
	Range	1 to 604800 seconds
Quiet Period	Allows you to specify the time period between any single EAPOL authentication failure and the start of a new EAPOL authentication attempt.	
	Default	60 seconds
	Range	0 to 65535 seconds
Transmit Period	Allows you to specify how long the switch waits for the supplicant to respond to EAP Request/Identity packets.	
	Default	30 seconds
	Range	1 to 65535 seconds
Supplicant Timeout	Allows you to specify how long the switch waits for the supplicant to respond to all EAP packets, except EAP Request/Identity packets.	
	Default	30 seconds
	Range	1 to 65535 seconds
Server Timeout	Allows you to specify how long the switch waits for the RADIUS server to respond to all EAP packets.	
	Default	30 seconds
	Range	1 to 65535 seconds
Maximum Requests	Allows you to specify the number of times the switch attempts to resend EAP packets to a supplicant.	
	Default	2 attempts
	Range	1 to 10 attempts

VLAN Configuration Menu Screen

The VLAN Configuration Menu screen (Figure 67) allows you to select the appropriate screen to configure up to 256 VLANs. VLAN 1 is port-based by default. You can configure the remaining 255 VLANs to be of any appropriate combination of types.

You can configure as many as 255 protocol-based VLANs, with up to 7 different protocols**.** The number of different protocols you can configure depends on the number of hexadecimal values (PID values) associated with the protocol type. Some protocol types use more than one PID value. Refer to "Predefined Protocol Identifier (PID) description". A port may not be a member of more than one protocol-based VLAN with the same PID. (Untagged ports cannot belong to different VLANs of the same protocol type; however, tagged ports can.)

When you create VLANs, you can assign various ports (and therefore the devices attached to these ports) to different broadcast domains. Creating VLANs increases network flexibility by allowing you to reassign devices to accommodate network moves, additions, and changes, eliminating the need to change physical cabling.

For detailed information about configuring VLANs, refer to Chapter 2.

To open the VLAN Configuration Menu:

➨ Choose VLAN Configuration (or press v) from the Switch Configuration Menu screen.

Figure 67 VLAN Configuration Menu screen

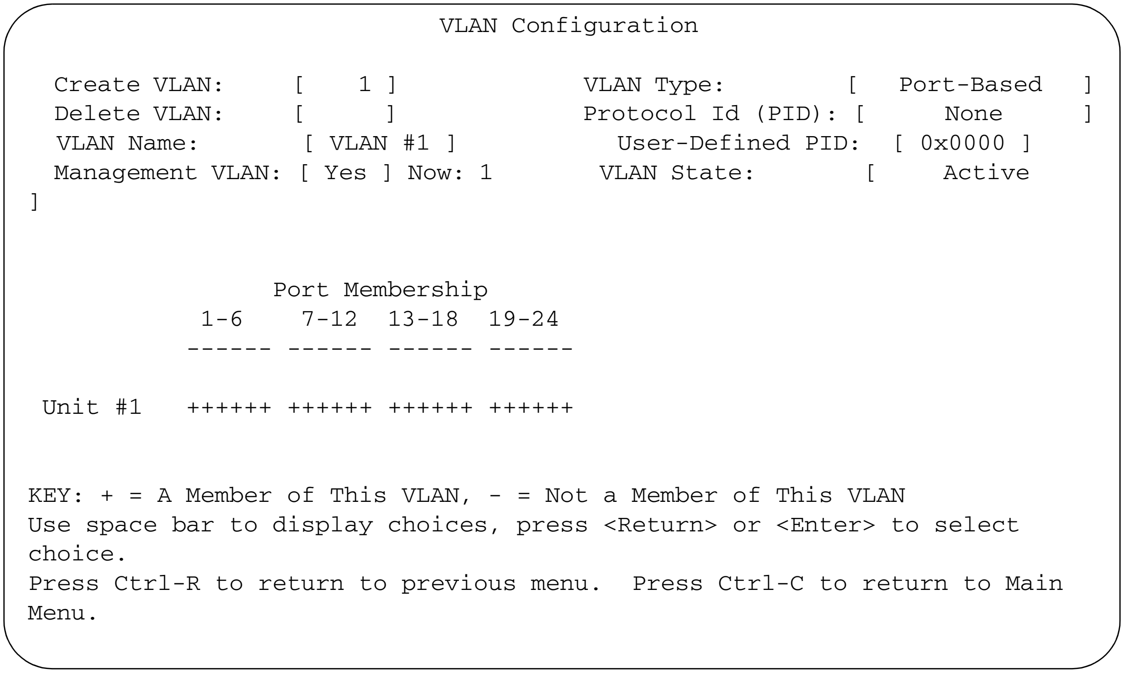

```
                               VLAN Configuration

   Create VLAN:      [    1 ]             VLAN Type:          [   Port-Based    ]
   Delete VLAN:      [      ]             Protocol Id (PID): [      None       ]
   VLAN Name:         [ VLAN #1 ]           User-Defined PID:   [ 0x0000 ]
   Management VLAN: [ Yes ] Now: 1         VLAN State:         [     Active
 ]



                 Port Membership
            1-6    7-12  13-18  19-24
           ------ ------ ------ ------


  Unit #1    ++++++ ++++++ ++++++ ++++++


 KEY: + = A Member of This VLAN, - = Not a Member of This VLAN
 Use space bar to display choices, press <Return> or <Enter> to select
 choice.
 Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
 Menu.
```

Table 34 describes the VLAN Configuration Menu screen options.

Table 34 VLAN Configuration Menu Screen options

Option	Description
VLAN Configuration...	Displays the VLAN Configuration screen (see "VLAN Configuration Screen"" on page 179). This screen allows you to set up VLAN workgroups.
VLAN Port Configuration...	Displays the VLAN Port Configuration screen (see "VLAN Port Configuration Screen"" on page 184). This screen allows you to set up a specific switch port.
VLAN Display by Port...	Displays the VLAN Display by Port screen (see "VLAN Display by Port Screen"" on page 187).

VLAN Configuration Screen

The VLAN Configuration screen (Figure 68) allows you to create and assign VLAN port memberships to standalone unit ports. You can create port-based and policy-based VLANs for the following purposes:

- IEEE 802.1Q port-based VLANs allow you to explicitly configure switch ports as VLAN port members.

  When you create a port-based VLAN, you assign a Port VLAN Identifier (PVID) and specify which ports belong to the VLAN.
- Policy-based VLANs allow you to configure your switch ports as members of a broadcast domain, based on the information within a packet. Policy-based VLANs can localize broadcast traffic and assure that only the policy-based VLAN ports are flooded with the specified packets.

When you configure ports as VLAN port members, they become part of a set of ports that form a broadcast domain for a specific VLAN. You can assign switch ports to a standalone unit ports, as VLAN port members of one or more VLANs.

➡ **Note:** Refer to Chapter 1 and guidelines for configuring spanning tree groups for more information on configuring VLANs.

You can add or remove port members from a VLAN in accordance with the IEEE 802.1Q tagging rules. Refer to Chapter 2 for a description of important terms used with 802.1Q VLANs.

You can also use this screen to create and to delete specific VLANs, to assign VLAN names, and to assign any VLAN as the management VLAN.

To open the VLAN Configuration screen:

➨ Choose VLAN Configuration (or press v) from the VLAN Configuration Menu screen.

Figure 68 VLAN Configuration screen

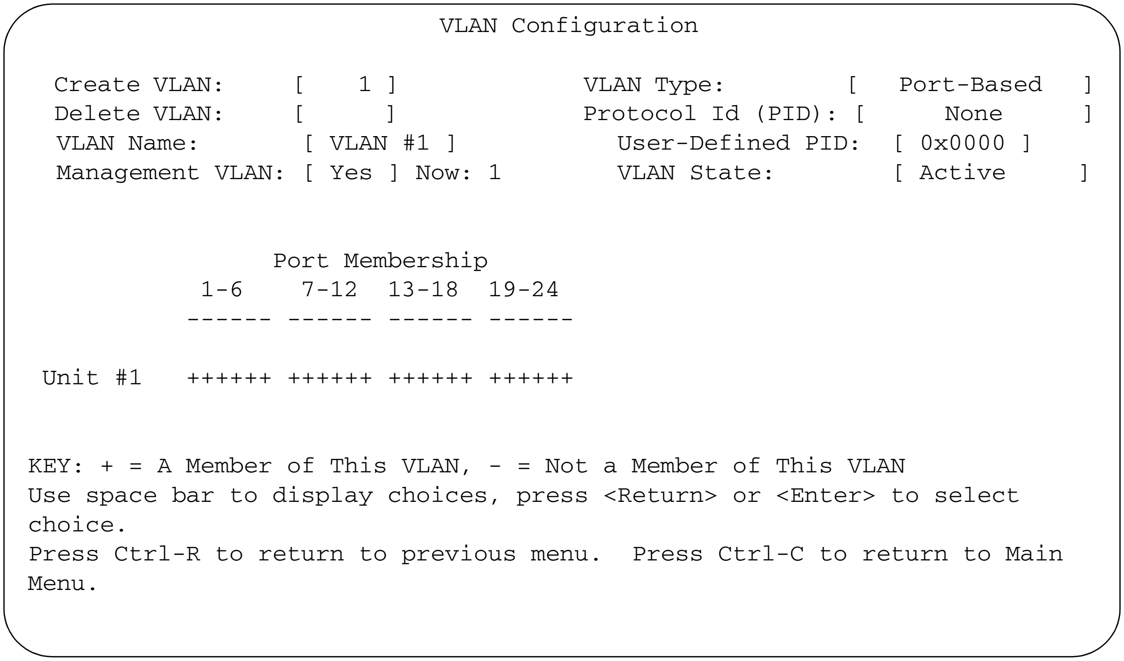

```
                              VLAN Configuration

   Create VLAN:      [    1 ]              VLAN Type:          [   Port-Based   ]
   Delete VLAN:      [      ]              Protocol Id (PID): [      None      ]
   VLAN Name:         [ VLAN #1 ]            User-Defined PID:   [ 0x0000 ]
   Management VLAN: [ Yes ] Now: 1          VLAN State:         [ Active      ]


                  Port Membership
             1-6    7-12  13-18  19-24
            ------ ------ ------ ------

  Unit #1   ++++++ ++++++ ++++++ ++++++


 KEY: + = A Member of This VLAN, - = Not a Member of This VLAN
 Use space bar to display choices, press <Return> or <Enter> to select
 choice.
 Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
 Menu.
```

Table 35 describes the VLAN Configuration screen fields.

Table 35 VLAN Configuration screen fields

Field	Description
Create VLAN	Allows you to set up or view configured VLAN workgroups. Enter the number of the new VLAN you want to create or view, then press [Return]. Alternatively, you can use the space bar to toggle through the various configured VLAN workgroups. You can create up to 255 different VLANs (except VLAN #1). By default, all switch ports are assigned as untagged members of VLAN 1 with all ports configured as PVID = 1. See Chapter 1 for more information Default 1 Range 1 to 4094

Table 35 VLAN Configuration screen fields (continued)

Field	Description	
Delete VLAN	Allows you to delete specified VLANs, except the assigned management VLAN (See Management VLAN field). Enter the number of the VLAN you want to delete, then press [Return], or use the space bar to toggle through the selection until you reach the VLAN you want to delete, then press [Return].	
	The specified VLAN is deleted as soon as you press [Return]. The software does not prompt you to reconsider this action. If you delete a VLAN, all configuration parameters that are associated with that VLAN are deleted also. You cannot delete VLAN 1.	
	Default	Blank
	Range	2 to 4094
VLAN Name	Allows you to assign a name field to configured VLANs.	
	Default	VLAN # (*VLAN number*)
	Range	Any ASCII string of up to 16 printable characters
Management VLAN	Allows you to assign any VLAN as the management VLAN. VLAN 1 is the default management VLAN for the switch. To set this field, the VLAN State field value must be Active.	
	Default	No
	Range	Yes, No
VLAN Type	Allows you to select the type of VLAN (port-based or protocol-based) to create. To set this field, the VLAN State field value must be Inactive.	
	Default	Port-based
	Range	Port-based, Protocol-based
Protocol ID (PID)	Allows you to set the protocol type of your VLAN (to set this field, the VLAN State field value must be Inactive). You can choose from any of predefined supported protocols (see "Predefined Protocol Identifier (PID) description"), or you can create your own user-defined protocol-based VLAN (see the User-defined PID field description for more information). **NOTE**: You can choose up to 7 PID values (in hex), as shown on Table 36 and Table 37.	
	Default	None
	Range	None, IP Ether2, Ipx 802.3, Ipx 802.2, Ipx Snap, Ipx Ether2, Declat Ether2, Sna 802.2, Sna Ether2, NetBios 802.2, Xns Ether2,Vines Ether2, Ipv6 Ether2, User-Defined, Rarp Ether2
User-Defined PID	Allows you to create your own user-defined VLAN where you specify the Protocol Identifier (PID) for the VLAN. To set this field, the VLAN State field must be set to Inactive. Some restrictions apply. "User-Defined Protocol Identifier Description".	
	Default	0x0000

Table 35 VLAN Configuration screen fields (continued)

Field	Description
	Range Any 16-bit hexadecimal value (for example, 0xABCD)
VLAN State	Allows you to activate your newly created VLAN. The following field values: VLAN Type, Protocol Id (PID), or User-defined PID must be configured appropriately before this field can be set to active. After you set the VLAN State field value to Active, you cannot change the VLAN State, VLAN Type, Protocol Id, or User-defined PID field values, unless you delete the VLAN. If you delete a VLAN, all configuration parameters that are associated with that VLAN are also deleted. Default Inactive Range Inactive, Active
Port Membership	Allows you to assign VLAN port memberships to *standalone unit* ports. The ports can be configured in one or more VLANs. To set this field, you must set the VLAN State field to Active. The Port Membership fields indicate the corresponding VLAN workgroup configuration, if configured. Dashes (-) indicate no VLAN Members are configured.
	The Port Membership fields are displayed in six-port groups (for example, 1-6, 7-12, 13-18). The number of ports displayed depends on the switch model.

Predefined Protocol Identifier (PID) description

Table 36 defines the standard protocol-based VLANs and PID types that are supported by the Ethernet Routing Switch 3510-24T.

Table 36 Predefined Protocol Identifier (PID)

PID Name	Encapsulation	PID Value (hex)	VLAN Type
IP Ether2	Ethernet type 2	0800, 0806	Standard IP on Ethernet Type 2 frames
Ipx 802.3	Ethernet 802.3	FF FF	Novell IPX on Ethernet 802.3 frames
Ipx 802.2	Ethernet 802.2	E0 E0	Novell IPX on Ethernet 802.2 frames
Ipx Snap	Ethernet Snap	8137, 8138	Novell IPX on Ethernet SNAP frames
Declat Ether2	Ethernet type 2	6004	DEC LAT protocol
Sna 802.2	Ethernet 802.2	04**, **04	IBM SNA on IEEE 802.2 frames
Sna Ether2	Ethernet type 2	80D5	IBM SNA on Ethernet Type 2 frames
NetBios 802.2	Ethernet type 2	F0**, **F0	NetBIOS protocol
Xns Ether2	Ethernet type 2	0600, 0807	Xerox XNS
Vines Ether2	Ethernet type 2	0BAD	Banyan VINES

Table 36 Predefined Protocol Identifier (PID) (continued)

PID Name	**Encapsulation**	**PID Value (hex)**	**VLAN Type**
Ipv6 Ether2	Ethernet type 2	86DD	IP version 6
User-Defined	Ethernet type 2, Ethernet 802.2, or Ethernet Snap	User-defined 16 bit value	User-defined protocol-based VLAN (see "Predefined Protocol Identifier (PID) description"" below, for more information).
RARP Ether2	Ethernet type 2	8035	Reverse Address Resolution Protocol (RARP): RARP is a protocol used by some old diskless devices to obtain IP addresses by providing the MAC layer address. When you create a VLAN based on RARP, you can limit the RARP broadcasts to the ports that lead to the RARP server.

User-Defined Protocol Identifier Description

In addition to the standard predefined protocols, user-defined protocol-based VLANs are supported. For user-defined protocol-based VLANs, you specify the protocol identifier (PID) for the VLAN. Any frames that match the specified PID in any of the following ways are assigned to that user-defined VLAN:

- The ethertype for Ethernet type 2 frames
- The PID in Ethernet SNAP frames
- The DSAP or SSAP value in Ethernet 802.2 frames

The following PIDs (Table 37) are reserved and are not available for user-defined PIDs.

Table 37 Reserved PIDs

PID Value (hex)	**Comments**
04**, **04	Sna 802.2
F0**, **F0	NetBIOS 802.2
AAAA	SNAP
0 - 05DC	Overlaps with 802.3 frame length
0600, 0807	Xns Ether2
0BAD	Vines Ether2

Table 37 Reserved PIDs (continued)

PID Value (hex)	Comments
4242	IEEE 802.1D BPDUs
6000 - 6009, 8038	Dec
0800, 0806	Ip Ether2 (including ARP)
8035	RARP Ether2
809B, 80F3	AplTk Ether2Snap
8100	IEEE 802.1Q for tagged frames
8137, 8138	Ipx
80D5	SNA Ether2
86DD	Ipv6 Ether2
8808	Ipx 802.3
E0E0	Ethernet 802.2
F0F0	Ethernet 802.3

VLAN Port Configuration Screen

The VLAN Port Configuration screen (Figure 69) allows you to configure specified switch ports with the appropriate PVID/VLAN association that enables the creation of VLAN broadcast domains (see Chapters 1 and 2 for more information about setting up VLAN broadcast domains).

You can configure specified switch ports to filter (discard) all received untagged frames or unregistered frames (see Chapters 1 and 2). Refer to the guidelines for configuring spanning tree groups in Chapter 1 for more information on configuring ports for tagged or untagged frames.

You can also prioritize the order in which the switch forwards packets, on a per-port basis (see Chapters 1 and 2). Refer to Chapter 4 “Policy-enabled networks,” for more information on prioritizing traffic.

To open the VLAN Port Configuration screen:

➨ Choose VLAN Port Configuration (or press c) from the VLAN Configuration Menu screen.

Figure 69 VLAN Port Configuration screen

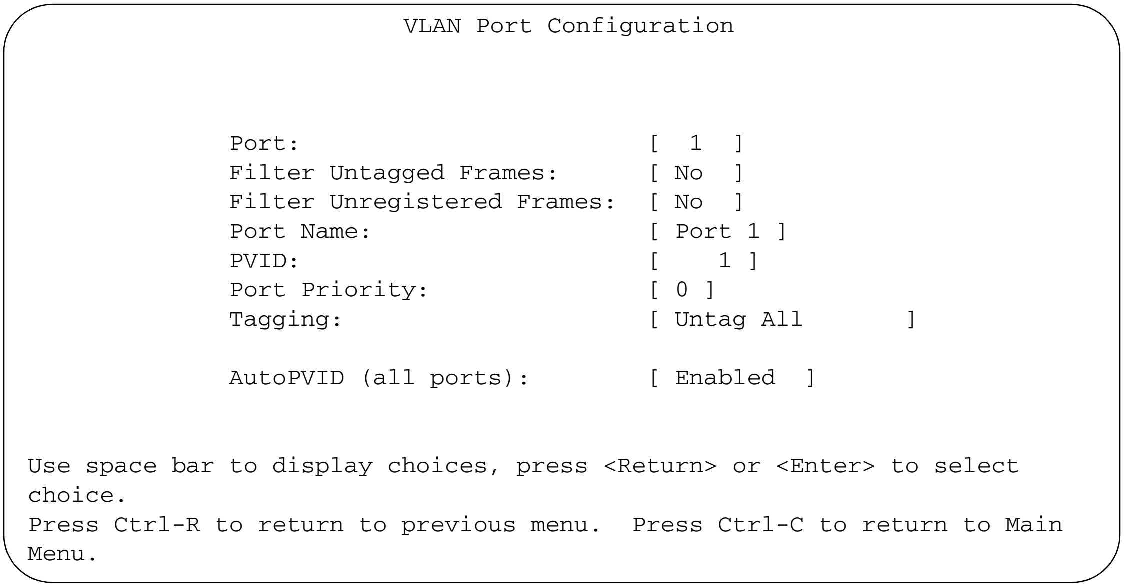

```
                         VLAN Port Configuration



            Port:                          [  1  ]
            Filter Untagged Frames:        [ No  ]
            Filter Unregistered Frames:    [ No  ]
            Port Name:                     [ Port 1 ]
            PVID:                          [    1 ]
            Port Priority:                 [ 0 ]
            Tagging:                       [ Untag All       ]

            AutoPVID (all ports):          [ Enabled  ]


Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 38 describes the VLAN Port Configuration screen fields.

Table 38 VLAN Port Configuration screen fields

Field	Description
Port	Allows you to select the number of the port you want to view or configure. To view another port, type its port number and press [Enter], or press the spacebar to toggle the port numbers.
Filter Untagged Frames	Sets this port to filter (discard) all received untagged frames. This features enables the port to filter frames with a VID that the port does not belong to. Default No Range No, Yes
Filter Unregistered Frames	Sets this port to filter (discard) all received unregistered packets. The Ethernet Routing Switch 3510-24T does not support the Yes option. Default No Range No, Yes

Table 38 VLAN Port Configuration screen fields (continued)

Field	**Description**
Port Name	The default port name assigned to this port. You can change this field to any name that is up to 16 characters long. Default Port *x* Range Any ASCII string of up to 16 printable characters
PVID	Associates this port with a specific VLAN. For example, a port with a PVID of 3 assigns all untagged frames received on this port to VLAN 3. Default 1 Range 1 to 4094
Port Priority	Prioritizes the order in which the switch forwards packets received on specified ports. Default 0 Range 0 to 7
Tagging	Allows you to assign VLAN Port Membership tagging options to this port, as follows: • Untag All: Any VLAN that this port is a member of *will not* be 802.1Q tagged. • Tag All: Any VLAN that this port is a member of will be 802.1Q tagged.
Tagging (continued)	• Tag PVID Only: Only frames whose VLAN ID match the PVID value assigned to the egress port will be tagged. • Untag PVID Only: All frames are tagged except those whose VLAN ID matches the PVID value assigned to the egress port. Setting this field value on any port to Tag All causes incoming untagged packets to be assigned to the PVID VLAN. They will no longer be classified based on the information within the packet, even if they are members of a policy-based VLAN. Default Untag All
	Range Untag All, Tag All, Tag PVID Only, Untag PVID Only
AutoPVID	When enabled, AutoPVID sets the PVID of a port when the port is added to a VLAN, unless the port is configured as Tagg All or Untag PVID Only. Default Enabled Range Enabled, Disabled

VLAN Display by Port Screen

The VLAN Display by Port screen (Figure 70) allows you to view VLAN characteristics associated with a specified switch port.

Choose VLAN Display by Port (or press d) from the VLAN Configuration Menu screen to open the VLAN Display by Port screen.

Figure 70 VLAN Display by Port screen

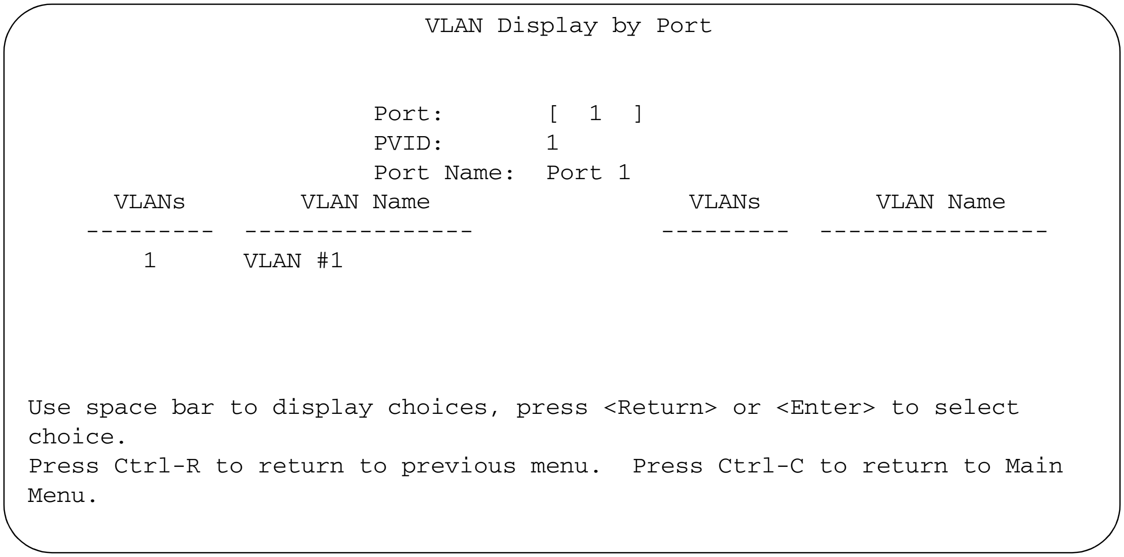

```
                         VLAN Display by Port


                    Port:        [  1  ]
                    PVID:        1
                    Port Name:   Port 1
    VLANs         VLAN Name                   VLANs         VLAN Name
  ---------   ----------------              ---------   ----------------
      1       VLAN #1





Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 39 describes the VLAN Display by Port screen fields.

Table 39 VLAN Display by Port screen fields

Field	Description
Port	Allows you to select the number of the port you want to view. To view another port, type its port number and press [Enter], or press the spacebar on your keyboard to toggle the port numbers.
PVID	Read-only field that indicates the PVID setting for the specified port.
Port Name	Read-only field that indicates the port name assigned to the specified port.
VLANs	Column header for the read-only fields listing the VLANs associated with the specified port.
VLAN Name	Column header for the read-only fields listing the VLAN Names associated with the specified port.

Port Configuration Screen

The Port Configuration screen (Figures 71 and 72) allows you to configure specific switch ports or all switch ports. You can enable or disable the port status of specified switch ports, set the switch ports to autonegotiate for the highest available speed of the connected station, or set the speed for selected switch ports (autonegotiation is not supported on fiber optic ports).

You can disable switch ports that are trunk members; however, the screen prompts for verification of the request before completing the action. Choosing [Yes] disables the port and removes it from the trunk.

Note: The Autonegotiation fields, the Speed fields, and the Duplex fields are independent of MultiLink Trunking, rate limiting, VLANs, IGMP Snooping, and the STP.

To open the Port Configuration screen:

➨ Choose Port Configuration (or press p) from the Switch Configuration Menu screen.

Figure 71 Port Configuration screen 1

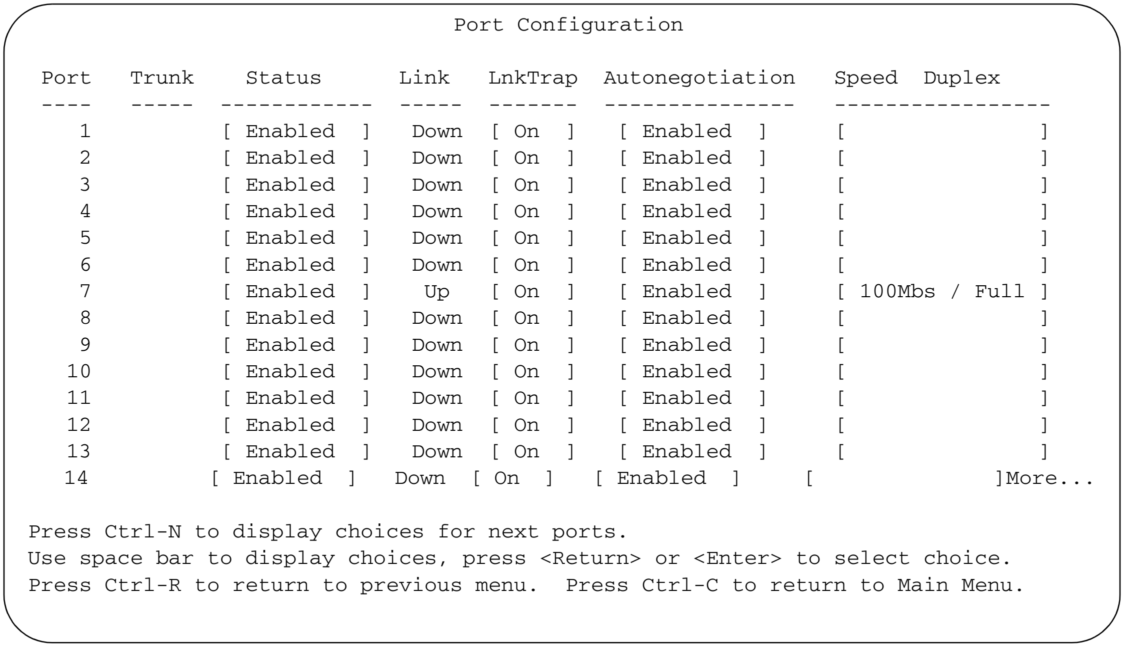

```
                                 Port Configuration

 Port   Trunk    Status      Link   LnkTrap  Autonegotiation   Speed  Duplex
 ----   -----  ------------  -----  -------  ---------------   -----------------
    1          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
    2          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
    3          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
    4          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
    5          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
    6          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
    7          [ Enabled  ]    Up   [ On  ]   [ Enabled  ]     [ 100Mbs / Full ]
    8          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
    9          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   10          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   11          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   12          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   13          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   14         [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [              ]More...

Press Ctrl-N to display choices for next ports.
Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Figure 72 Port Configuration screen 2

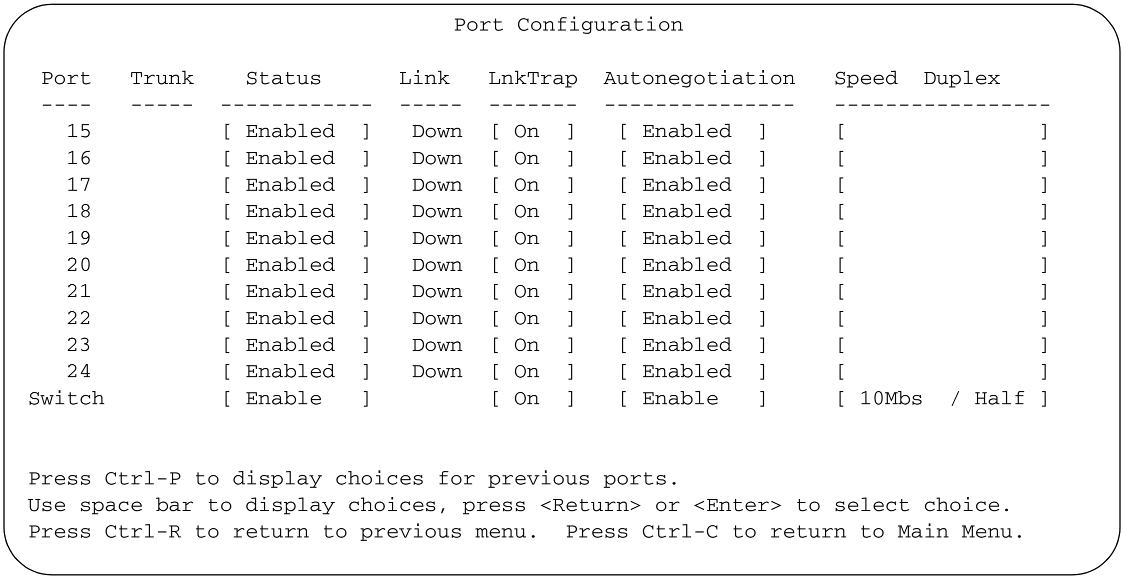

```
                                 Port Configuration

 Port   Trunk    Status      Link   LnkTrap  Autonegotiation   Speed  Duplex
 ----   -----  ------------  -----  -------  ---------------   -----------------
   15          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   16          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   17          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   18          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   19          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   20          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   21          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   22          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   23          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
   24          [ Enabled  ]   Down  [ On  ]   [ Enabled  ]     [               ]
Switch         [ Enable   ]         [ On  ]   [ Enable   ]     [ 10Mbs  / Half ]


Press Ctrl-P to display choices for previous ports.
Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

➡ **Note:** When a SFP GBICs is installed, only the Status field for that SFP GBIC port is configurable. See "High Speed Flow Control Configuration Screen"" on page 191 to set the autonegotiation field for the gigabit MDA port.

Table 40 describes the Port Configuration screen fields.

Table 40 Port Configuration screen fields

Field	**Description**
Port	Indicates the switch port numbers that correspond to the field values in that row of the screen (for example, the field values in row 2 apply to switch port 2). The values that you set in the *Switch* row will affect all switch ports.
Trunk	The read-only data displayed in this column indicates the trunks that correspond to the switch ports specified in the Trunk Members fields of the Trunk Configuration screen (see "MultiLink Trunk Configuration Menu screen" on page 193).
Status	Allows you to disable any of the switch ports. You can also use this field to control access to any switch port. Default Value Enabled Range Enabled, Disabled
Link	A read-only field that indicates the current link state of the corresponding port, as follows: • Up: The port is connected and operational. • Down: The port is not connected or is not operational.
LnkTrap	Allows you to control whether link up/link down traps are sent to the configured trap sink from the switch. Default Value On Range On, Off
Autonegotiation	When enabled, sets the corresponding port speed to match the best service provided by the connected station, up to 1000 Mb/s in full-duplex mode. Autonegotiation can be enabled or disabled for all fiber optic ports. Default Value Enabled Range Enabled, Disabled
Speed/Duplex	Allows you to manually configure any port to support an Ethernet speed of 10 Mb/s or 100 Mb/s, in half- or full-duplex mode, or 1000 Mb/s in full-duplex mode. This field is set (by default) to 1000 Mb/s, full-duplex for gigabit ports only. Default Value 10 Mbps/Half (when Autonegotiation is Disabled) Range 10Mbps/Half, 10Mbs/Full, 100Mbs/Half,100Mbs/Full,1000Mbs/Full

High Speed Flow Control Configuration Screen

The High Speed Flow Control Configuration screen (Figure 73) allows you to set the flow control parameters all gigabit ports.

➨ Choose High Speed Flow Control Configuration (or press h) from the Switch Configuration Menu screen to open the High Speed Flow Control Configuration screen.

Figure 73 High Speed Flow Control Configuration

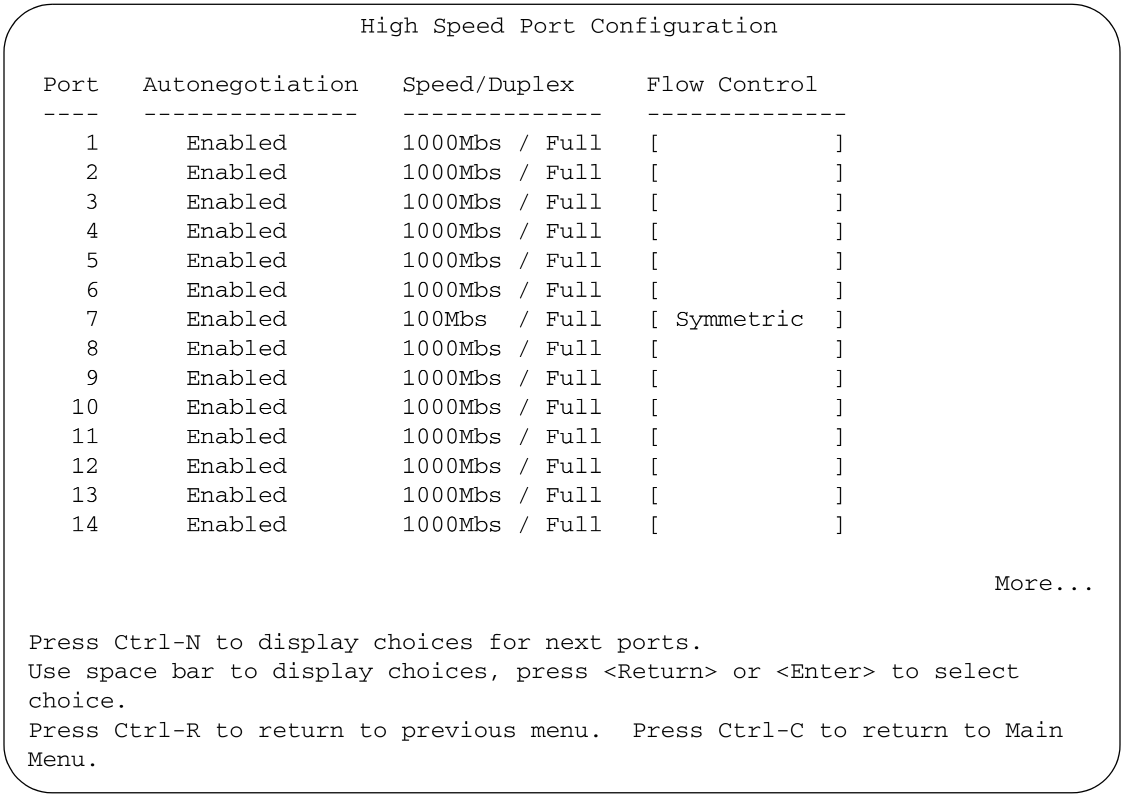

```
                          High Speed Port Configuration

 Port    Autonegotiation    Speed/Duplex     Flow Control
 ----    ---------------    --------------   --------------
    1       Enabled         1000Mbs / Full   [            ]
    2       Enabled         1000Mbs / Full   [            ]
    3       Enabled         1000Mbs / Full   [            ]
    4       Enabled         1000Mbs / Full   [            ]
    5       Enabled         1000Mbs / Full   [            ]
    6       Enabled         1000Mbs / Full   [            ]
    7       Enabled         100Mbs  / Full   [ Symmetric  ]
    8       Enabled         1000Mbs / Full   [            ]
    9       Enabled         1000Mbs / Full   [            ]
   10       Enabled         1000Mbs / Full   [            ]
   11       Enabled         1000Mbs / Full   [            ]
   12       Enabled         1000Mbs / Full   [            ]
   13       Enabled         1000Mbs / Full   [            ]
   14       Enabled         1000Mbs / Full   [            ]

                                                                     More...

Press Ctrl-N to display choices for next ports.
Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 41 describes the High Speed Flow Control Configuration screen fields.

Table 41 High Speed Flow Control Configuration Screen Fields

Field	Description
Port	Displays the port number.
Autonegotiation	Displays autonegotiation setting. **NOTE:** Autonegotiation can be changed in the Port Configuration screen. Default Value Enabled Range Enabled, Disabled
Speed/Duplex	Displays the speed and duplex mode. **NOTE:** The speed/duplex can be changed in the Port Configuration screen.
Flow Control	Allows you to control traffic and avoid congestion on the Gigabit ports. Two modes are available (see "Choosing a High Speed Flow Control Mode"" for details about the two modes). The Flow Control field can be configured only when you set the Autonegotiation field value to Disabled and the speed to 1000Mbs/full duplex. Default Value Disabled (when Autonegotiation is disabled) Range Disabled, Symmetric, Asymmetric

Choosing a High Speed Flow Control Mode

The high speed flow control feature allows you to control traffic and avoid congestion on the Gigabit full-duplex link. If the receive port buffer becomes full, the Ethernet Routing Switch 3510-24T issues a flow-control signal to the device at the other end of the link to suspend transmission. When the receive buffer is no longer full, the switch issues a signal to resume the transmission. You can choose Symmetric or Asymmetric flow control mode.

Note: For high speed flow control, the Ethernet Routing Switch 3510-24T must be connected to a device that is IEEE802.3x-compliant.

Symmetric Mode

This mode allows a port and its link partner to send flow control *pause* frames to each other.

When a pause frame is received (by either the port or its link partner), the port suspends transmission of frames for a number of slot times specified in the control frame or until a pause-release control frame is received. Both devices on the link must support this mode when it is selected.

Asymmetric Mode

This mode allows the link partner to send flow control pause frames to the port. When a pause frame is received, the receiving port suspends transmission of frames for a number of slot times specified in the control frame or until a pause-release control frame is received.

In this mode, the port is disabled from transmitting pause frames to its link partner. Use this mode when the port is connected to a buffered repeater device.

MultiLink Trunk Configuration Menu screen

The MultiLink Trunk Configuration Menu screen (Figure 74) allows you to select the appropriate screen to configure up to six MultiLink Trunks (you can group up to four switch ports together to form each trunk).

You can monitor the bandwidth usage for the trunk member ports within each trunk. For more information about configuring MultiLink Trunks, see Chapters 1 and 2.

Note: When a trunk is not active (Trunk Status field set to Disabled), configuration changes do not take effect until you set the Trunk Status field to Enabled.

To open the MultiLink Trunk Configuration Menu screen:

➨ Choose MultiLink Trunk Configuration (or press t) from the Switch Configuration Menu screen.

Figure 74 MultiLink Trunk Configuration Menu screen

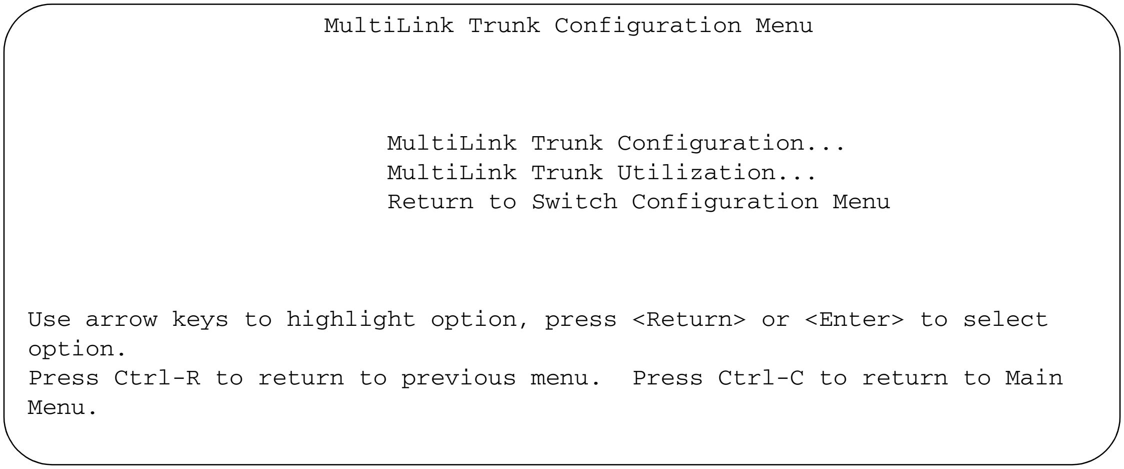

```
                      MultiLink Trunk Configuration Menu



                         MultiLink Trunk Configuration...
                         MultiLink Trunk Utilization...
                         Return to Switch Configuration Menu




Use arrow keys to highlight option, press <Return> or <Enter> to select
option.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 42 describes the MultiLink Trunk Configuration Menu screen options.

Table 42 MultiLink Trunk Configuration Menu screen options

Option	Description
MultiLink Trunk Configuration...	Displays the MultiLink Trunk Configuration screen (Figure 75). This screen allows you to configure up to six MultiLink Trunks within a standalone switch. You can group up to four switch ports together to form each trunk.
MultiLink Trunk Utilization...	Displays the MultiLink Trunk Utilization screen (Figure 76 and Figure 77). This screen allows you to monitor the bandwidth utilization of the configured trunks.

MultiLink Trunk Configuration Screen

The MultiLink Trunk Configuration screen (Figure 75) allows you to configure up to six trunks in a standalone switch.

When the trunks are enabled, the trunk members take on default settings necessary for correct operation of the MultiLink Trunking feature. These default settings can affect the correct operation of your configured network. If you disable a trunk, you may need to reconfigure the specific trunk members switch ports to return to the previous switch configuration. See Chapter 1 for more information.

To open the MultiLink Trunk Configuration screen:

➨ Choose Trunk Configuration (or press t) from the MultiLink Trunk Configuration Menu screen.

Figure 75 MultiLink Trunk Configuration screen

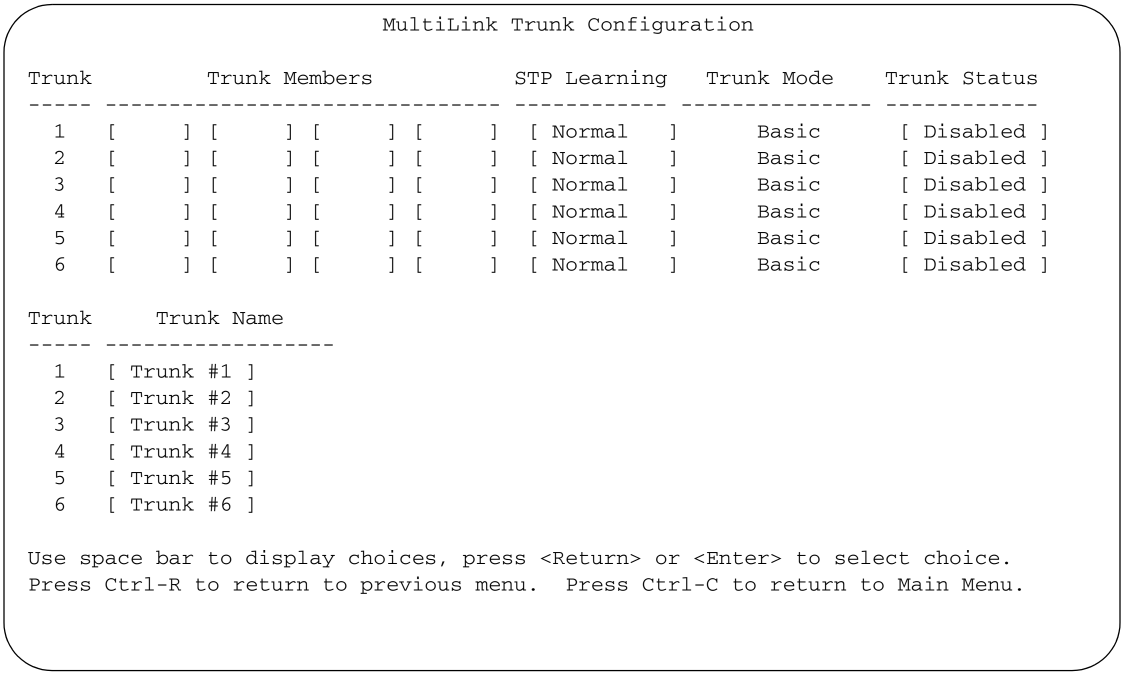

```
                            MultiLink Trunk Configuration

Trunk         Trunk Members           STP Learning   Trunk Mode    Trunk Status
----- ------------------------------- ------------ --------------- ------------
  1   [     ] [     ] [     ] [     ]  [ Normal   ]      Basic      [ Disabled ]
  2   [     ] [     ] [     ] [     ]  [ Normal   ]      Basic      [ Disabled ]
  3   [     ] [     ] [     ] [     ]  [ Normal   ]      Basic      [ Disabled ]
  4   [     ] [     ] [     ] [     ]  [ Normal   ]      Basic      [ Disabled ]
  5   [     ] [     ] [     ] [     ]  [ Normal   ]      Basic      [ Disabled ]
  6   [     ] [     ] [     ] [     ]  [ Normal   ]      Basic      [ Disabled ]

Trunk     Trunk Name
----- ------------------
  1   [ Trunk #1 ]
  2   [ Trunk #2 ]
  3   [ Trunk #3 ]
  4   [ Trunk #4 ]
  5   [ Trunk #5 ]
  6   [ Trunk #6 ]

Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 43 describes the MultiLink Trunk Configuration screen fields.

Table 43 MultiLink Trunk Configuration screen fields

Field	Description
Trunk	Column header for the read-only fields in this screen. The read-only data displayed in the Trunk column indicates the trunk (1 to 6) that corresponds to the switch ports specified in the user-configurable Trunk Members fields.
Trunk Members (Port)	The Trunk Members column contains fields in each row that can be configured to create the corresponding trunk. Each switch port can only be a member of a single trunk. **NOTE**: Trunk members can be removed only if trunk is Disabled. Default Value Blank Range 1 to 24
STP Learning	Specifies the Spanning Tree participation settings for link members in a trunk. Fast is the same as Normal, except that the state transition timer is shortened to two seconds. **NOTE**: The same participation setting is applied to all members in the trunk. Default Value Normal Range Normal, Fast, Disabled
Trunk Mode	Specifies the Trunk Mode for each trunk; always set to Basic. **NOTE:**. When in Basic mode, source MAC addresses are dynamically assigned to specific trunk members for flooding and forwarding, which allows the switch to stabilize and distribute the data streams of source addresses across the trunk members.
Trunk Status	The Trunk Status column contains a single field for each row that allows users to enable or disable any of the trunks. Default Value Disabled Range Enabled, Disabled
Trunk Name	The Trunk Name column contains a single optional field in each row that can be used to assign names to the corresponding configured trunks. The names chosen for this example can provide meaningful information to the user (for example, S1:T1 to FS2 indicates Trunk 1, in switch S1 connects to File Server 2).

MultiLink Trunk Utilization Screen

The MultiLink Trunk Utilization screen (Figure 76 and Figure 77) allows you to monitor the percentage of bandwidth used by configured trunk members. You can choose the type of traffic to monitor.

Figure 76 shows an *example* of bandwidth utilization rates for trunk member ports. Because two screens are necessary to show all of the configured trunks (up to six), the screen prompts you to Press [Ctrl]-N to view trunks five and six.

➨ Choose MultiLink Trunk Utilization (or press u) from the MultiLink Trunk Configuration Menu screen to open the MultiLink Trunk Utilization screen.

Figure 76 MultiLink Trunk Utilization screen (1 of 2)

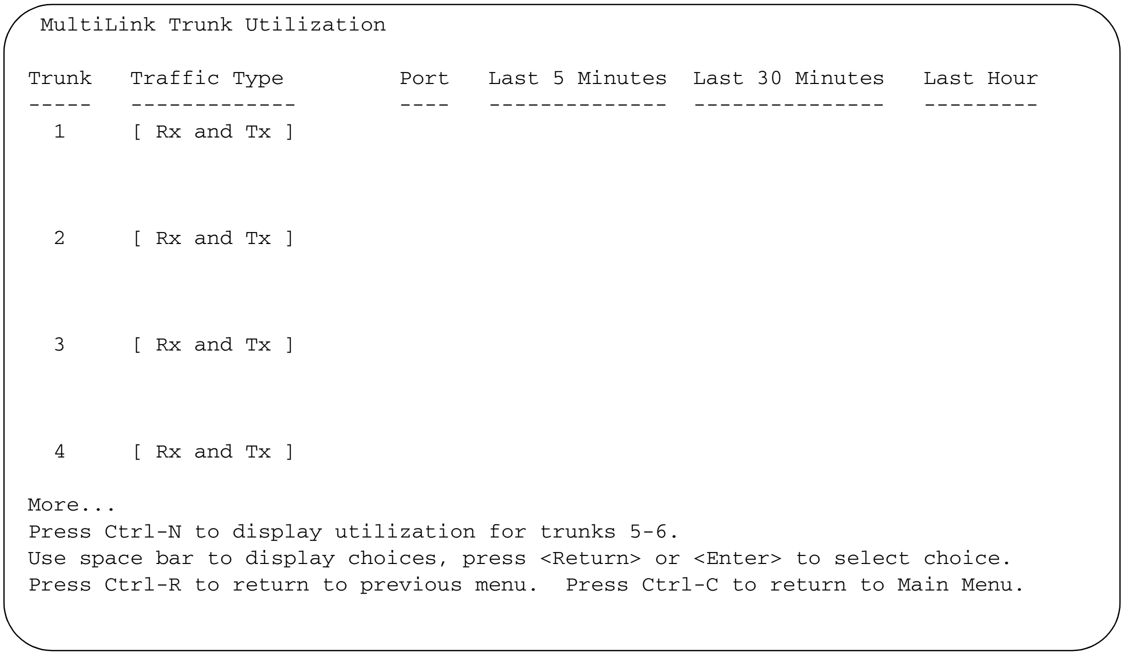

```
 MultiLink Trunk Utilization

Trunk   Traffic Type         Port   Last 5 Minutes  Last 30 Minutes   Last Hour
-----   -------------        ----   --------------  ---------------   ---------
  1     [ Rx and Tx ]



  2     [ Rx and Tx ]



  3     [ Rx and Tx ]



  4     [ Rx and Tx ]

More...
Press Ctrl-N to display utilization for trunks 5-6.
Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Figure 77 MultiLink Trunk Utilization screen (2 of 2)

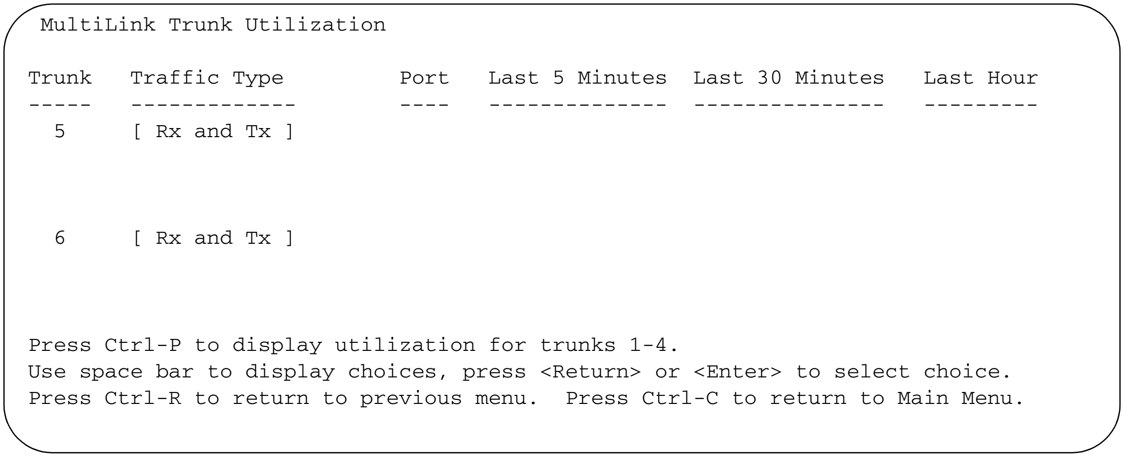

```
 MultiLink Trunk Utilization

Trunk   Traffic Type         Port   Last 5 Minutes  Last 30 Minutes   Last Hour
-----   -------------        ----   --------------  ---------------   ---------
  5     [ Rx and Tx ]




  6     [ Rx and Tx ]




Press Ctrl-P to display utilization for trunks 1-4.
Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 44 describes the MultiLink Trunk Utilization screen fields.

Table 44 MultiLink Trunk Utilization screen fields

Field	Description
Trunk	Column header for the read-only fields in this screen. The read-only data displayed in this column indicates the trunk (1 to 6) that corresponds to the switch ports specified in the Port field.
Traffic Type	Allows you to choose the traffic type to be monitored for percent of bandwidth utilization (see Range). Default Value Rx and Tx Range Rx and Tx, Rx, Tx
Port	Lists the trunk member ports that correspond to the trunk specified in the Trunk column.
Last 5 Minutes	This read-only field indicates the percentage of packets (of the type specified in the Traffic Type field) utilized by the port in the last 5 minutes. This field provides a running average of network activity and is updated every 15 seconds.
Last 30 Minutes	This read-only field indicates the percentage of packets (of the type specified in the Traffic Type field) utilized by the port in the last 30 minutes. This field provides a running average of network activity and is updated every 15 seconds.
Last Hour	This read-only field indicates the percentage of packets (of the type specified in the Traffic Type field) utilized by the port in the last 60 minutes. This field provides a running average of network activity and is updated every 15 seconds.

Port Mirroring Configuration Screen

The Port Mirroring Configuration screen allows you to configure a specific switch port to monitor up to two specified ports or two MAC addresses. You can specify port-based monitoring or address-based monitoring.

For more information about the port mirroring feature, see Chapter 1.

Figure 78 shows an example of a Port Mirroring Configuration screen.

To open the Port Mirroring Configuration screen:

➨ Choose Port Mirroring Configuration (or press i) from the Switch Configuration Menu screen.

Figure 78 Port Mirror Configuration screen

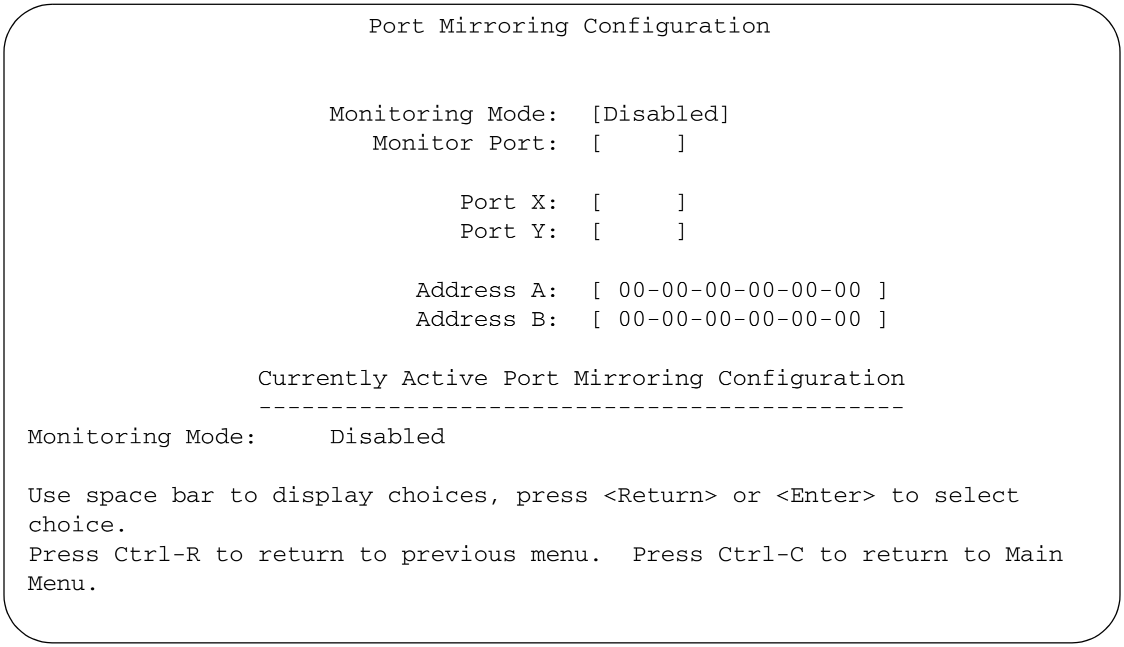

```
                        Port Mirroring Configuration


                     Monitoring Mode:  [Disabled]
                        Monitor Port:  [     ]

                              Port X:  [     ]
                              Port Y:  [     ]

                           Address A:  [ 00-00-00-00-00-00 ]
                           Address B:  [ 00-00-00-00-00-00 ]

                Currently Active Port Mirroring Configuration
                ---------------------------------------------
Monitoring Mode:     Disabled

Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 45 describes the Port Mirroring Configuration screen fields.

Table 45 Port Mirroring Configuration screen fields

Field	Description	
Monitoring Mode	Allows a user to select any one of six port-based monitoring modes or any one of five address-based monitoring modes (see Table 46). Selecting any one of the six *port-based modes* activates the port X and port Y screen fields, where a user can choose up to two ports to monitor. Selecting any one of the five *address-based modes* activates the Address A and Address B screen fields, where a user can specify MAC addresses to monitor.	
	Default Value	Disabled
	Range	See Table 46
Monitor Port	Indicates the port number that is designated as the monitor port.	
	Default Value	Zero-length string
	Range	1 to 8 and 1 to 24 or 1 to 48 (depending on model type)
Port X	Indicates one of the ports that will be monitored by the designated port monitor when one of the port-based monitoring modes is selected. This port will be monitored according to the value of Port X in the Monitoring Mode field (see Table 46).	
	Default Value	Zero-length string
	Range	1 to 24
Port Y	Indicates one of the ports that will be monitored by the designated port monitor when one of the port-based monitoring modes is selected. This port will be monitored according to the value of Port Y in the Monitoring Mode field (see Table 46).	
	Default Value	Zero-length string
	Range	1 to 24
Address A	Indicates the MAC addresses that will be monitored by the designated port monitor when one of the address-based monitoring modes is selected. This port will be monitored according to the value of Address A in the selected Monitoring Mode field (see Table 46).	
	Default Value	00-00-00-00-00-00 (no MAC address assigned)
	Range	00-00-00-00-00-00 to FF-FF-FF-FF-FF-FF
Address B	Indicates the MAC addresses that will be monitored by the designated port monitor when one of the address-based monitoring modes is selected. This port will be monitored according to the value of Address B in the selected Monitoring Mode field (see Table 46).	
	Default Value	00-00-00-00-00-00 (no MAC address assigned)
	Range	00-00-00-00-00-00 to FF-FF-FF-FF-FF-FF

Table 46 describes the various monitoring modes available from the Port Mirroring Configuration screen.

Table 46 Monitoring modes

Field	Description
Port-based:	
Disabled	Default value for this feature.
-> Port X	Monitor all traffic received by Port X.
Port X ->	Monitor all traffic transmitted by Port X.
<-> Port X	Monitor all traffic received and transmitted by Port X.
-> Port X or Port Y ->	Monitor all traffic received by Port X or transmitted by Port Y. Note: Do not use this mode for broadcast or multicast traffic.
-> Port X and Port Y ->	Monitor all traffic received by Port X (destined to Port Y) and then transmitted by Port Y. Note: Do not use this mode for broadcast or multicast traffic
<-> Port X and Port Y <->	Monitor all traffic received/transmitted by Port X and all traffic received/transmitted by Port Y. Note: Do not use this mode for broadcast or multicast traffic
Address-based:	
Disabled	Default value for this feature.
Address A -> any Address	Monitor all traffic transmitted from Address A to any address.
any Address -> Address A	Monitor all traffic received by Address A from any address.
<-> Address A	Monitor all traffic received by or transmitted by Address A.
Address A -> Address B	Monitor all traffic transmitted by Address A to Address B.
Address A <-> Address B	Monitor all traffic between Address A and Address B (conversation between the two stations).

Rate Limiting Configuration Screen

The Rate Limiting Configuration screen allows you to limit the forwarding rate of broadcast and multicast packets at ingress.

Figures 79 and 80 show sample rate limiting values for the two Rate Limiting Configuration screens.

Note: If a port is configured for rate limiting, and it is a MultiLink Trunk member, all trunk member ports implement rate limiting. Also, if a trunk member is implementing rate limiting and the port is disabled from rate limiting, all trunk members are disabled from rate limiting.

To open the Rate Limiting Configuration screen:

➨ Choose Rate Limiting Configuration (or press l) from the Switch Configuration Menu screen.

Figure 79 Rate Limiting Configuration screen (1 of 2)

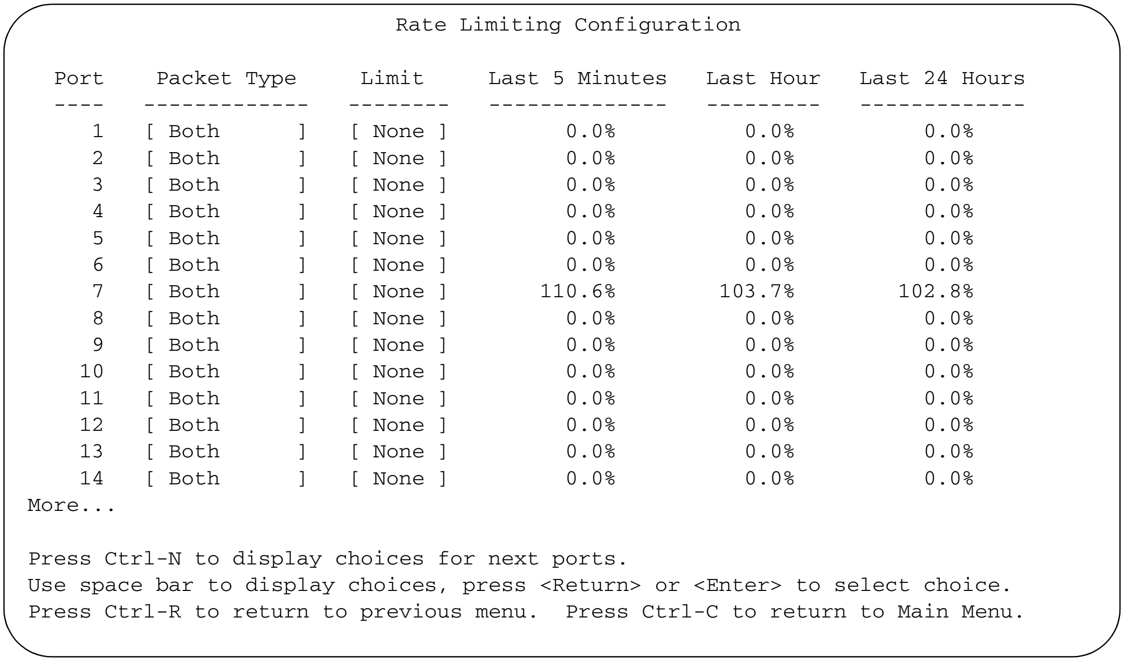

```
                         Rate Limiting Configuration

  Port    Packet Type      Limit     Last 5 Minutes   Last Hour   Last 24 Hours
  ----   -------------   --------   --------------   ---------   -------------
     1   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
     2   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
     3   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
     4   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
     5   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
     6   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
     7   [ Both      ]   [ None ]      110.6%        103.7%        102.8%
     8   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
     9   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
    10   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
    11   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
    12   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
    13   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
    14   [ Both      ]   [ None ]        0.0%          0.0%          0.0%
More...

Press Ctrl-N to display choices for next ports.
Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Figure 80 Rate Limiting Configuration screen (2 of 2)

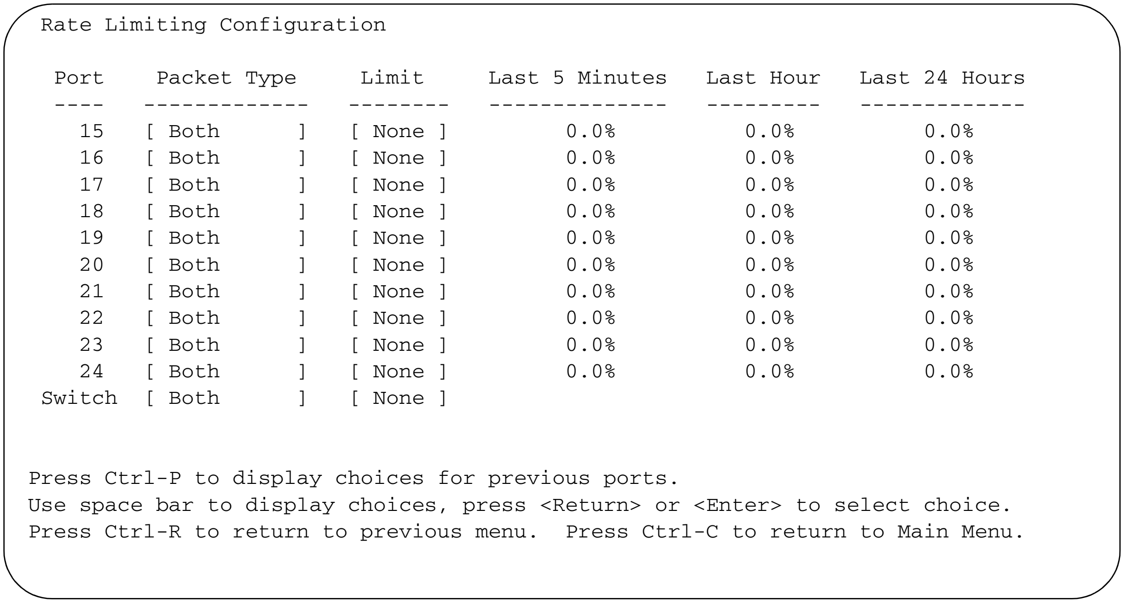

```
 Rate Limiting Configuration

  Port    Packet Type     Limit     Last 5 Minutes   Last Hour   Last 24 Hours
  ----   -------------   --------   --------------   ---------   -------------
    15   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    16   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    17   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    18   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    19   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    20   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    21   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    22   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    23   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
    24   [ Both      ]   [ None ]         0.0%          0.0%          0.0%
 Switch  [ Both      ]   [ None ]


Press Ctrl-P to display choices for previous ports.
Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

You can use this screen to view the percentage of either packet type (or both packet types) *received* on each port.

When the volume of either packet type is high, placing severe strain on the network (often referred to as a “storm”), you can set the forwarding rate of those packet types to *not exceed* a specified percentage of the total available bandwidth. The percentage you set refers to the total available bandwidth, not to a percentage of current traffic.

Table 47 describes the Rate Limiting Configuration screen fields.

Table 47 Rate Limiting Configuration screen fields

Field	Description
Port	Indicates the switch port numbers that correspond to the field values in that row of the screen (for example, the field values in row 2 apply to switch port 2). Note that the values applied in the Switch affect all standalone switch ports.
Packet Type	Allows you to select the packet types for rate-limiting or viewing. Default Value Both Range Both, Multicast, Broadcast
Limit	Sets the percentage of port bandwidth allowed for forwarding the packet types specified in the Packet Type field. When the threshold is exceeded, any additional packets (specified in the Packet Type field) are discarded*. Default Value None Range None, 10%, 9%, 8%, 7%, 6%, 5%, 4%, 3%, 2%, 1%
Last 5 Minutes	This read-only field indicates the percentage of packets (of the type specified in the Packet Type field) received by the port in the last 5 minutes. This field provides a running average of network activity and is updated every 15 seconds. Note that this field indicates the receiving port's view of network activity, regardless of the rate-limiting setting.
Last Hour	This read-only field indicates the percentage of packets (of the type specified in the Packet Type field) received by the port in the last hour. This field provides a running average of network activity and is updated every 5 minutes. Note that this field indicates the receiving port's view of network activity, regardless of the rate-limiting setting.
Last 24 Hours	This read-only field indicates the percentage of packets (of the type specified in the Packet Type field) received by the port in the last 24 hours. This field provides a running average of network activity and is updated every hour. Note that this field indicates the receiving port's view of network activity, regardless of the rate-limiting setting.

* Rate-limiting is disabled if this field is set to None. This allows you to select and view the percentage of specific packet types present in the network, without inadvertently limiting the forwarding rate.

IGMP Configuration Menu Screen

The IGMP Configuration Menu screen (Figure 81) allows you to select the appropriate screen to optimize IP Multicast packets in a bridged Ethernet environment (see Chapter 1).

Note: You can have up to 256 multicast groups with the Ethernet Routing Switch 3510-24T.

To open the IGMP Configuration Menu screen:

➨ Choose IGMP Configuration (or press g) from the Switch Configuration Menu screen.

Figure 81 IGMP Configuration Menu screen

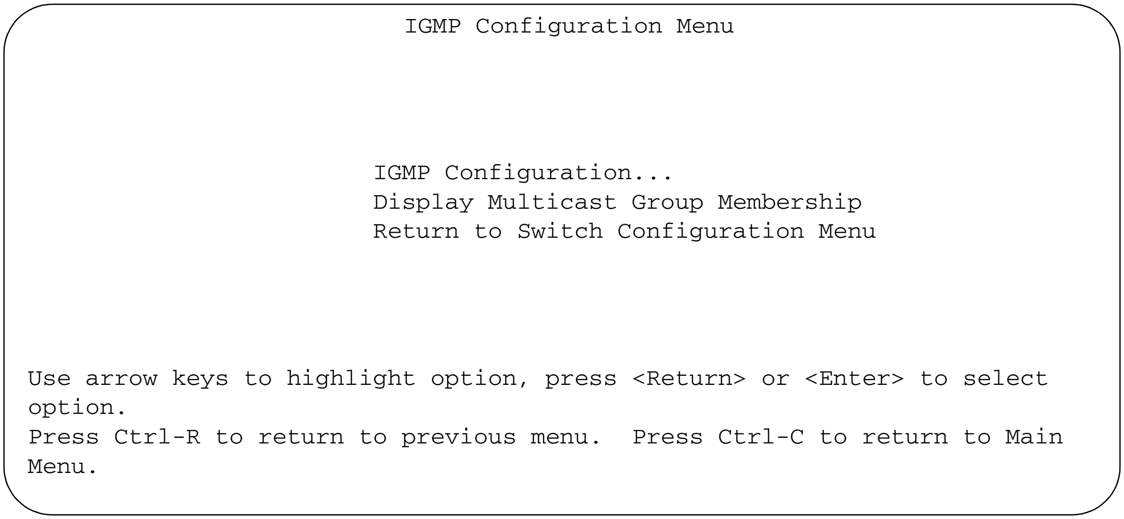

```
                              IGMP Configuration Menu




                            IGMP Configuration...
                            Display Multicast Group Membership
                            Return to Switch Configuration Menu






Use arrow keys to highlight option, press <Return> or <Enter> to select
option.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 48 describes the IGMP Configuration Menu screen options.

Table 48 IGMP Configuration Menu screen options

Option	Description
IGMP Configuration...	Displays the IGMP Configuration screen (see “IGMP Configuration Screen”). This screen allows you to set up IGMP VLAN configurations.

Table 48 IGMP Configuration Menu screen options

Option	Description
Display Multicast Group Membership...	Displays the Multicast Group Membership screen (see“Multicast Group Membership Screen”. This screen allows you to view all IP Multicast addresses that are active in the current LAN.

IGMP Configuration Screen

To open the IGMP Configuration screen:

➨ Choose IGMP Configuration (or press g) from the Switch Configuration Menu screen.

Figure 82 IGMP Configuration screen

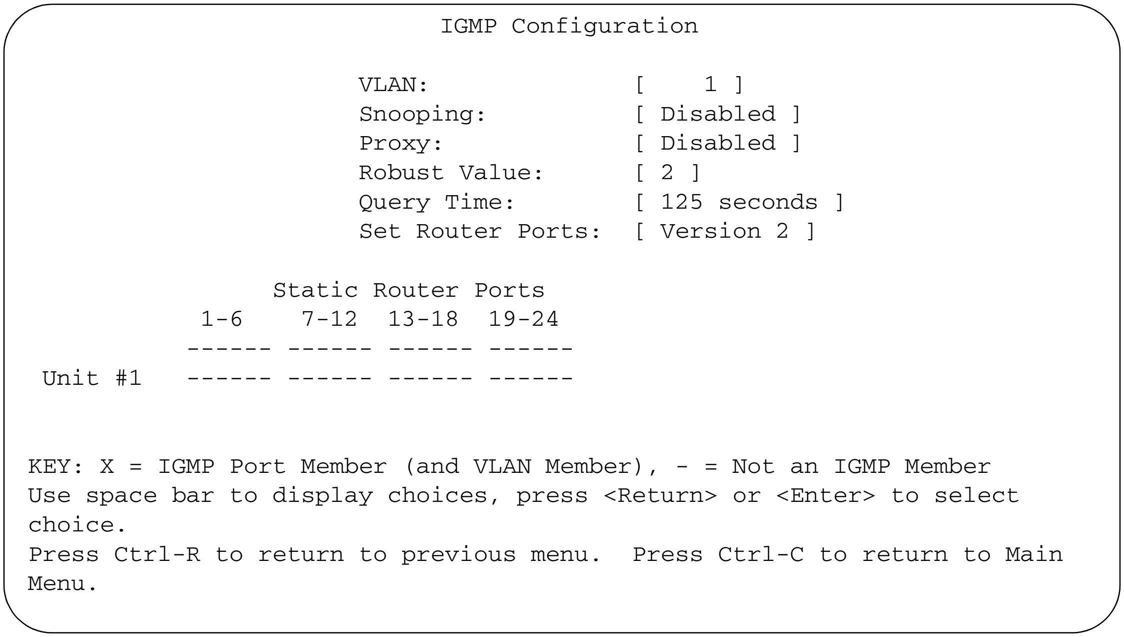

```
                             IGMP Configuration

                       VLAN:              [    1 ]
                       Snooping:          [ Disabled ]
                       Proxy:             [ Disabled ]
                       Robust Value:      [ 2 ]
                       Query Time:        [ 125 seconds ]
                       Set Router Ports:  [ Version 2 ]

                 Static Router Ports
            1-6    7-12  13-18  19-24
           ------ ------ ------ ------
 Unit #1   ------ ------ ------ ------


KEY: X = IGMP Port Member (and VLAN Member), - = Not an IGMP Member
Use space bar to display choices, press <Return> or <Enter> to select
choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 49 describes the IGMP Configuration screen fields.

Table 49 IGMP Configuration screen fields

Field	Description
VLAN	Allows you to set up or view IGMP VLAN configurations on *specified* VLANs. You can use the space bar to toggle to any *existing* IGMP VLAN configurations (the maximum number of VLANs that can be displayed is 256). Default 1 Range 1 to 4094
Snooping	Allows you to enable or disable IGMP Snooping. This field affects *all* VLANs. Default Value Enabled Range Enabled, Disabled
Proxy	Allows the switch to consolidate IGMP Host Membership Reports received on its downstream ports and to generate a consolidated proxy report for forwarding to its upstream neighbor. This field affects specified VLANs. The Proxy field cannot be enabled unless the Snooping field is enabled. Default Value Disabled Range Enabled, Disabled
Robust Value	Allows a user to set the switch to offset expected packet loss on a subnet. The Robust Value is used to tune the subnet for packet losses. If packet losses on a subnet are unacceptably high, you can increase the value in the Robust Value field. This field affects *only* the VLAN specified in the screen's VLAN field (for example, if you change the robust value on the VLAN specified in the screen's VLAN field, other VLANs are not affected). Default Value 2 Range 1 to 255
Query Time	Allows a user to control the number of IGMP messages allowed on the subnet by varying the *Query Interval* (the Query Interval is the interval between general queries sent by the multicast router). This is used to age out the hosts that are learned by the switch. The switch sends only Group-specific queries after the query timeout; that is, for the groups the switch has learned. This field affects *only* the VLAN specified in the screen's VLAN field (for example, if you change the Query Time value field on the VLAN specified in the screen's VLAN field, other VLANs are not affected). Default Value 125 seconds Range 1 to 65535 seconds

Table 49 IGMP Configuration screen fields (continued)

Field	Description
Set Router Ports	Selects the IGMP version according to the IGMPv1 (Version 1) or IGMPv2 (Version 2) standard (see RFC 2236). Use this field in conjunction with the Static Router Ports field (see next field description) to select the IGMP version to set. You can also use this field to view which static router ports are set to Version 1 or to Version 2. Use the space bar to toggle between the two versions and view the static router ports settings. This field affects only selected VLANs. Default Value Version 1 Range Version 1, Version 2
Static Router Ports	Allows a user to assign switch port to be a multicast router. The Ethernet Routing Switch 3510-24T treats such ports as ports connected to a multicast router (regardless of actual connections.). The configured ports do not filter any IP Multicast traffic. The Static Router Ports fields are displayed in six-port groups (for example, 1-6, 7-12, 13-18). This field affects only specified VLANs. Default Value - Range None

Multicast Group Membership Screen

The Multicast Group Membership screen allows you to view configured IP Multicast group addresses for specific VLANs. The screen displays the IP Multicast group addresses associated with ports that are configured within a standalone switch. The displayed addresses are dynamic and can change as clients join (or leave) the various IP Multicast groups.

To open the Multicast Group Membership screen:

➨ Choose Display Multicast Group Membership (or press d) from the IGMP Configuration Menu screen.

Figure 83 Multicast Group Membership screen

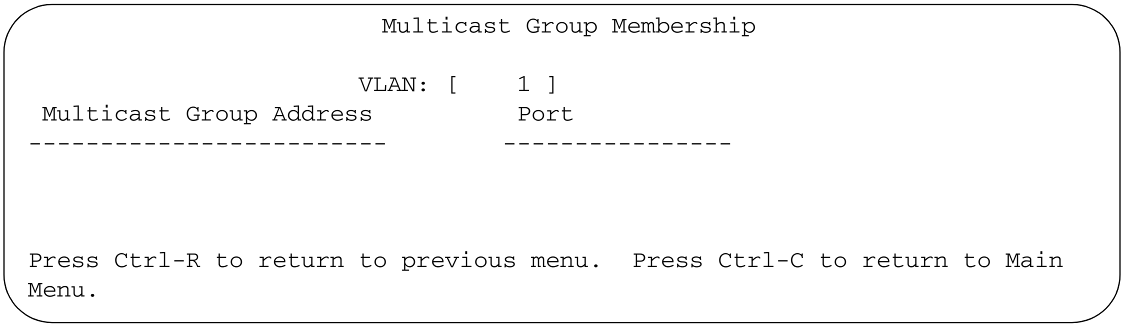

```
                           Multicast Group Membership

                         VLAN: [    1 ]
 Multicast Group Address              Port
-------------------------         ----------------




Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
Menu.
```

Table 50 describes the Multicast Group Membership screen options.

Table 50 Multicast Group Membership screen options

Option	Description
VLAN	Allows you to view multicast group addresses on specified VLANs. You can use the space bar to view group addresses for any existing IGMP VLAN configurations (the maximum number of VLANs that can be displayed is 256).
Multicast Group Address	Displays all of the IP Multicast group addresses that are currently active on the associated port.
Port	Displays the port numbers that are associated with the IP Multicast group addresses displayed in the IP Multicast group address field.

Port Statistics Screen

The Port Statistics screen (Figure 84) allows you to view detailed information about any switch or port in standalone configuration. The screen is divided into two sections (Received and Transmitted) so that you can compare and evaluate throughput or other port parameters. All screen data is updated approximately every 2 seconds.

You can use the Port Statistics screen to clear (reset to zero) port counters for a specific switch or port. Alternatively, you can use the Clear All Port Statistics option to clear port counters for all switches or ports (see “Switch Configuration Menu Screen” on page 158).

To open the Port Statistics screen:

➨ Choose Display Port Statistics (or press d) from the Switch Configuration Menu screen.

Figure 84 Port Statistics screen

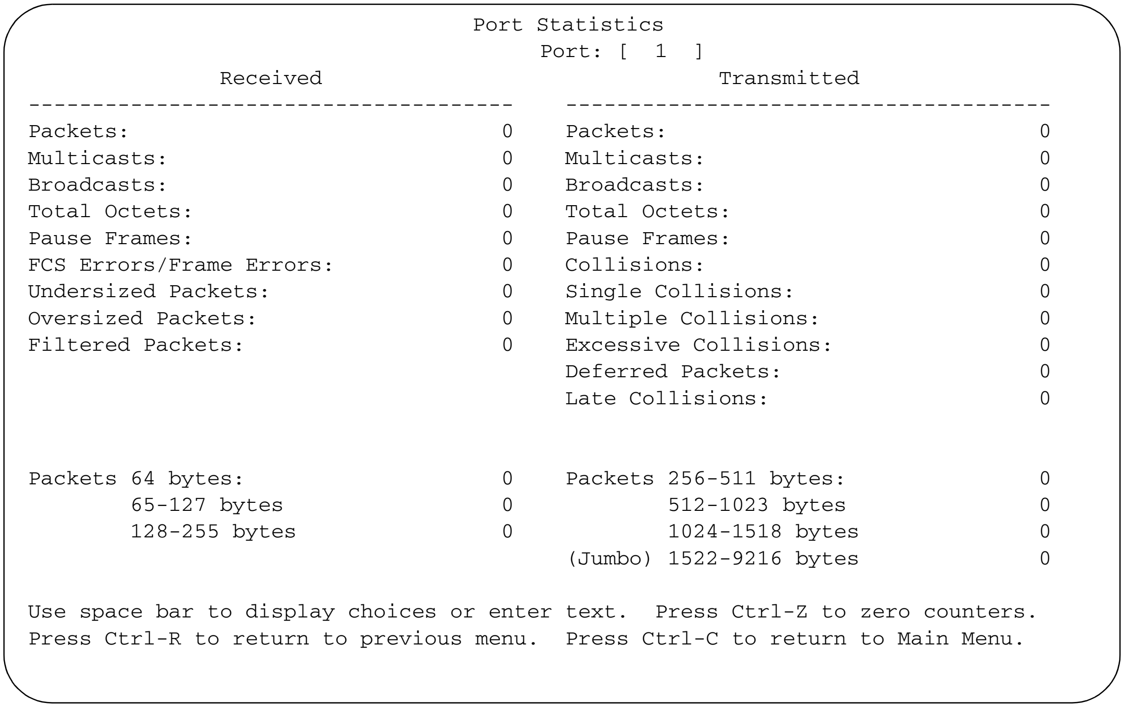

```
                                   Port Statistics
                                        Port: [  1  ]
               Received                               Transmitted
--------------------------------------    --------------------------------------
Packets:                             0    Packets:                             0
Multicasts:                          0    Multicasts:                          0
Broadcasts:                          0    Broadcasts:                          0
Total Octets:                        0    Total Octets:                        0
Pause Frames:                        0    Pause Frames:                        0
FCS Errors/Frame Errors:             0    Collisions:                          0
Undersized Packets:                  0    Single Collisions:                   0
Oversized Packets:                   0    Multiple Collisions:                 0
Filtered Packets:                    0    Excessive Collisions:                0
                                          Deferred Packets:                    0
                                          Late Collisions:                     0


Packets 64 bytes:                    0    Packets 256-511 bytes:               0
        65-127 bytes                 0            512-1023 bytes               0
        128-255 bytes                0            1024-1518 bytes              0
                                          (Jumbo) 1522-9216 bytes              0

Use space bar to display choices or enter text.  Press Ctrl-Z to zero counters.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 51 describes the Port Statistics screen fields.

Table 51 Port Statistics screen fields

Field	Description
Port	Allows you to select the number of the port you want to view or reset to zero. To view another port, type its port number and press [Enter], or press the spacebar on your keyboard to toggle the port numbers.
Packets	Received column: Indicates the total number of packets received on this port, including bad packets, broadcast packets, and multicast packets. Transmitted column: Indicates the total number of packets transmitted successfully on this port, including broadcast packets and multicast packets.
Multicasts	Received column: Indicates the total number of good multicast packets received on this port, excluding broadcast packets. Transmitted column: Indicates the total number of multicast packets transmitted successfully on this port, excluding broadcast packets.
Broadcasts	Received column: Indicates the total number of good broadcast packets received on this port. Transmitted column: Indicates the total number of broadcast packets transmitted successfully on this port.
Total Octets	Received column: Indicates the total number of octets of data (including data in bad packets) received on this port, excluding framing bits but including FCS octets. Transmitted column: Indicates the total number of octets of data transmitted successfully on this port, including FCS octets.
Pause Frames	Transmitted column: Indicates the total number of pause frames transmitted on this port. Pause frames cause the transmitting port to temporarily suspend the transmission of packets when the receiving port's frame buffer is full (Gigabit ports only). Received column: Indicates the total number of pause frames received on this port. Pause frames cause the transmitting port to temporarily suspend the transmission of packets when the receiving port's frame buffer is full (Gigabit ports only)
FCS Errors/Frame Errors	FCS Errors: Indicates the total number of valid-size packets that were received with proper framing but discarded because of cyclic redundancy check (CRC) errors. Frame Errors: Indicates the total number of valid-size packets that were received but discarded because of CRC errors and improper framing
Undersized Packets	Indicates the total number of packets received on this port with fewer than 64 bytes and with proper CRC and framing (also known as short frames or runts).
Oversized Packets	Indicates the total number of packets received on this port with more than 1518 bytes and with proper CRC and framing (also known as oversized frames).
Filtered Packets	Indicates the total number of packets discarded on this port for any reason.
Collisions	Indicates the total number of collisions detected on this port.

Table 51 Port Statistics screen fields (continued)

Field	Description
Single Collisions	Indicates the total number of packets that were transmitted successfully on this port after a single collision.
Multiple Collisions	Indicates the total number of packets that were transmitted successfully on this port after more than one collision.
Excessive Collisions	Indicates the total number of packets lost on this port due to excessive collisions.
Late Collisions	Indicates the total number of packet collisions that occurred after a total length of time that exceeded 512 bit-times of packet transmission.
Packets 64 bytes	Indicates the total number of 64-byte packets received or transmitted on this port.
65-127 bytes	Indicates the total number of 65-byte to 127-byte packets received or transmitted on this port.
128-255 bytes	Indicates the total number of 128-byte to 255-byte packets received or transmitted on this port.
256-511 bytes	Indicates the total number of 256-byte to 511-byte packets received or transmitted on this port.
512-1023 bytes	Indicates the total number of 512-byte to 1023-byte packets received or transmitted on this port.
1024-1518 bytes	Indicates the total number of 1024-byte to 1518-byte packets received or transmitted on this port.
(Jumbo) 1522-9216 bytes	Indicates the total number of jumbo packets (1522 bytes to 9216 bytes received or transmitted on this port.

Console/Comm Port Configuration Screen

The Console/Comm Port Configuration screen (Figure 85) allows you to configure and modify the console/comm port parameters and security features of a standalone switch.

To open the Console/Comm Port Configuration screen:

➨ Choose Console/Comm Port Configuration (or press o) from the main menu.

Figure 85 Console/Comm Port Configuration screen

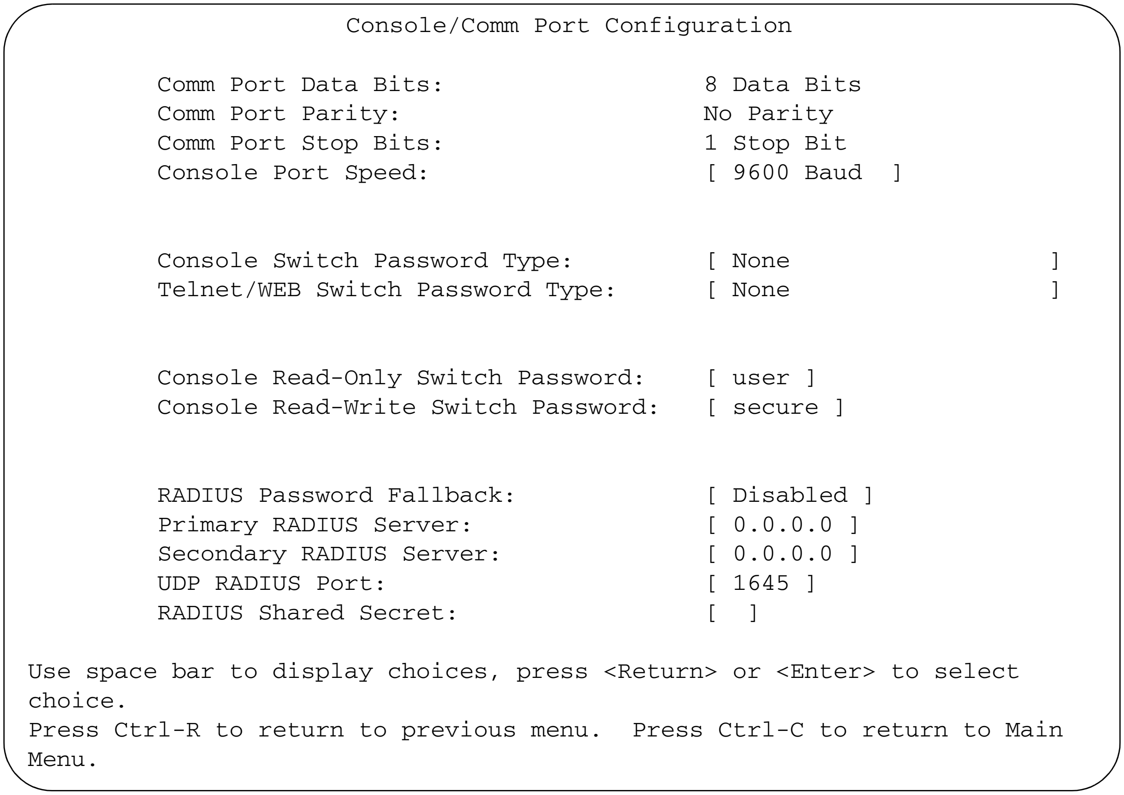

```
                     Console/Comm Port Configuration

          Comm Port Data Bits:                    8 Data Bits
          Comm Port Parity:                       No Parity
          Comm Port Stop Bits:                    1 Stop Bit
          Console Port Speed:                     [ 9600 Baud  ]


          Console Switch Password Type:           [ None                 ]
          Telnet/WEB Switch Password Type:        [ None                 ]


          Console Read-Only Switch Password:      [ user ]
          Console Read-Write Switch Password:     [ secure ]


          RADIUS Password Fallback:               [ Disabled ]
          Primary RADIUS Server:                  [ 0.0.0.0 ]
          Secondary RADIUS Server:                [ 0.0.0.0 ]
          UDP RADIUS Port:                        [ 1645 ]
          RADIUS Shared Secret:                   [  ]

 Use space bar to display choices, press <Return> or <Enter> to select
 choice.
 Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main
 Menu.
```

Table 52 describes the Console/Comm Port Configuration screen fields.

Table 52 Console/Comm Port Configuration screen fields

Field	**Description**
Comm Port Data Bits	A read-only field that indicates the current console/comm port data bit setting.
Comm Port Parity	A read-only field that indicates the current console/comm port parity setting.
Comm Port Stop Bits	A read-only field that indicates the current console/comm port stop bit setting.
Console Port Speed	Allows you to set the console/comm port baud rate to match the baud rate of the console terminal. Default Value: 9600 Baud Range: 2400 Baud, 4800 Baud, 9600 Baud, 19200 Baud, 38400 Baud

Table 52 Console/Comm Port Configuration screen fields (continued)

Field	Description
	Caution: If you choose a baud rate that does not match your console terminal baud rate, you will lose communication with the configuration interface when you press [Enter]. If communication is lost, set your console terminal to match the new service port setting.
	Achtung: Bei Auswahl einer Baud rate, die nicht mit der Baudrate des Konsolenterminals übereinstimmt, geht die Kommunikation mit der Konsolenschnittstelle verloren, wenn Sie die Eingabetaste drücken. Stellen Sie in diesem Fall das Konsolenterminal so ein, daß es mit der neuen Einstellung der Service-Schnittstelle übereinstimmt.
	Attention: Si vous sélectionnez un débit différent de celui de votre terminal, vous perdrez le contact avec l'interface de votre console dès que vous appuierez sur [Entrée]. Pour restaurer la communication, alignez le débit de votre terminal sur le nouveau débit de votre port de service.
	Precaución: Si selecciona una velocidad de transmisión que no coincide con la velocidad de transmisión del terminal de la consola, perderá la comunicación con el interfaz de la consola al pulsar [Intro]. Si se pierde la comunicación, ajuste el terminal de la consola para que coincida con el nuevo valor del puerto de servicio.
	Attenzione: Nel caso in cui si scelga una velocità di trasmissione non corrispondente a quella del terminale della console, la comunicazione con l'interfaccia della console cadrà premendo il tasto [Invio]. Se la comunicazione cade, impostare il terminale della console in modo tale che corrisponda alla nuova impostazione della porta di servizio.
	注意：コンソール・ターミナルのボー・レートに合っていないボー・レートを選択すると、[Enter]を押したときに、コンソール・インタフェイスとの通信が途切れてしまいます。この場合には、新しいサービス・ポート設定に合うようにコンソール・ターミナルを設定してください。
Console Switch Password Type	Enables password protection for accessing the console interface (CI) of a *standalone switch* through a console terminal. If you set this field to require a password, you can use the Logout option to restrict access to the CI. Thereafter, you will need to specify the correct password at the console-terminal prompt. See Console Read-Only Switch Password, Console Read-Write Switch Password, and RADIUS fields for more information. Default Value None Range None, Local Password, RADIUS Authentication

Table 52 Console/Comm Port Configuration screen fields (continued)

Field	**Description**
Telnet/WEB Switch Password Type	Enables password protection for accessing the console interface (CI) of a *standalone switch* through a Telnet session or through a Web interface session. If you set this field to require a password, you can use the Logout option to restrict access to the CI. Thereafter, you will need to specify the correct password at the console-terminal prompt. See Console Read-Only Switch Password, Console Read-Write Switch Password, and RADIUS fields descriptions for more information. Default Value None Range None, Local Password, RADIUS Authentication
Console Read-Only Switch Password	When the Console Switch Password field is set to require a password (for Telnet, for Console, or for Both), this field allows read-only password access to the CI of a *standalone switch*. Users can access the CI using the correct password (see default), but cannot change parameters or use the Reset option or Reset to Default option. Default Value user Range An ASCII string of up to 15 printable characters
Console Read-Write Switch Password	When the Console Switch Password field is set to require a password (for Telnet, for Console, or for Both), this field allows read-write password access to the CI of a *standalone switch*. Users can log in to the CI using the correct password (see default) and can change any parameter. You can change the default passwords for read-only access and read-write access to a private password.
	Default Value: secure Range: Any ASCII string of up to 15 printable characters
	Caution: If you change the system-supplied default passwords, be sure to write the new passwords down and keep them in a safe place. If you forget the new passwords, you cannot access the console interface. In that case, contact Nortel for help.
	Achtung: Wenn Sie die für das System standardmäßig eingestellten Paßwörter ändern, notieren Sie sich die neuen Paßwörter, und bewahren Sie sie an einem sicheren Ort auf. Falls Sie die neuen Paßwörter vergessen, können Sie nicht mehr auf die Konsolenschnittstelle zugreifen. Wenden Sie sich in diesem Fall an Nortel um Unterstützung zu erhalten.

Table 52 Console/Comm Port Configuration screen fields (continued)

Field	Description
	Attention: Si vous changez les mots de passe par défaut du système, assurez-vous de bien noter vos nouveaux mots de passe et de les conserver dans un endroit sûr. Si vous perdez vos nouveaux mots de passe, vous ne pourrez plus accéder à votre interface. Le cas échéant, veuillez contacter Nortel.
	Precaución: Si modifica las contraseñas predeterminadas asignadas por el sistema, asegúrese de anotar las nuevas contraseñas y guárdelas en un lugar seguro. Si olvida las nuevas contraseñas, no podrá acceder al interfaz de la consola. En ese caso, póngase en contacto con Nortel para obtener ayuda al respecto.
	Attenzione: In caso di modifica delle password predefinite nel sistema, assicurarsi di annotare le nuove password e di conservarle in un luogo sicuro. Nel caso in cui le nuove password vengano dimenticate, non sarà possibile accedere all'interfaccia della console. In tal caso, contattare la Nortel per avere assistenza.
	注意：システム装備したデフォルトのパスワードを変更する場合、必ず新しいパスワードを書き留めて安全な場所に保管してください。新しいパスワードを忘れてしまうと、コンソール・インタフェイスにアクセスできません。この場合は、Bay Networksまでご連絡ください。
	Caution: you change the system-supplied default passwords, be sure to write the new passwords down and keep them in a safe place. If you forget the new passwords, you cannot access the console interface. In that case, contact Nortel for help.
	Achtung: Wenn Sie die für das System standardmäßig eingestellten Paßwörter ändern, notieren Sie sich die neuen Paßwörter, und bewahren Sie sie an einem sicheren Ort auf. Falls Sie die neuen Paßwörter vergessen, können Sie nicht mehr auf die Konsolenschnittstelle zugreifen. Wenden Sie sich in diesem Fall an Nortel, um Unterstützung zu erhalten.
	Attention: Si vous changez les mots de passe par défaut du système, assurez-vous de bien noter vos nouveaux mots de passe et de les conserver dans un endroit sûr. Si vous perdez vos nouveaux mots de passe, vous ne pourrez plus accéder à votre interface. Le cas échéant, veuillez contacter Nortel.

Table 52 Console/Comm Port Configuration screen fields (continued)

Field	Description
	Precaución: Si modifica las contraseñas predeterminadas asignadas por el sistema, asegúrese de anotar las nuevas contraseñas y guárdelas en un lugar seguro. Si olvida las nuevas contraseñas, no podrá acceder al interfaz de la consola. En ese caso, póngase en contacto con Nortel para obtener ayuda al respecto.
	Attenzione: In caso di modifica delle password predefinite nel sistema, assicurarsi di annotare le nuove password e di conservarle in un luogo sicuro. Nel caso in cui le nuove password vengano dimenticate, non sarà possibile accedere all'interfaccia della console. In tal caso, contattare la Nortel per avere assistenza.
	注意：システム装備したデフォルトのパスワードを変更する場合、必ず新しいパスワードを書き留めて安全な場所に保管してください。新しいパスワードを忘れてしまうと、コンソール・インタフェイスにアクセスできません。この場合は、Bay Networksまでご連絡ください。
Primary RADIUS Server	The IP address of the Primary RADIUS server. Default 0.0.0.0 (no IP address assigned) Range Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point
Secondary RADIUS Server	The IP address of the Secondary RADIUS server. Default 0.0.0.0 (no IP address assigned) Range Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point
RADIUS UDP Port	The user datagram protocol (UDP) port for the RADIUS server. Default 1645 Range 0 to 65535
RADIUS Shared Secret	Your special switch security code that provides authentication to the RADIUS server. Default Null string (which will not authenticate) Range Any contiguous ASCII string that contains at least 1 printable character, up to a maximum of 16.

Recovering From a Lost Password

To recover a lost password:

1. Attach a console cable to the console port.
2. Start the terminal emulation program.
3. Apply power to the switch.
4. As the switch begins to initialize, press [Ctrl]-C.

   The switch returns the following messages:

```
Ethernet Routing Switch 3510-24T Diagnostics 3.0.0.4
Testing main memory - PASSED
>> Break Recognized - Wait
Resets: 142
Intiializing Flash...
Reading MAC Address...
MAC Address: 00:E0:7B:CC:78:C0
Initializing Switch CBs...
Initializing Switch HW....
Press ‘a’ to run Agent code
Press ‘d’ to Download agent code
Press ‘e’ to display Errors
Press ‘i’ to Initialize config/log/flash
Press ‘p’ to run POST tests.....
```

5. Press `i` to initialize the config.

   This kills the password and the configuration. The system displays the following message:

```
Erase Config/Log Flash Y/N [N]
```

6 Press `Y`.

The system displays the following message:

```
Erasing - Wait ) sec....
Press 'a' to run Agent code
Press 'd' to Download agent code
Press 'e' to display Errors
Press 'i' to Initialize config/log/flash
Press 'p' to run POST tests.....
```

Spanning Tree Configuration Menu Screen

The Spanning Tree Configuration Menu screen (Figure 86) allows you to view spanning tree parameters and configure multiple spanning tree groups (STGs).

➡ **Note:** You must use either the Command Line Interface (CLI) or Device Manager (DM) if you want to configure individual port values for path cost and priority.

➡ **Note:** Before configuring spanning tree groups, refer to Chapter 2 for guidelines and interactions with VLANs and MLT.

To open the Spanning Tree Configuration Menu screen:

➨ Choose Spanning Tree Configuration (or press p) from the main menu.

Figure 86 Spanning Tree Configuration Menu

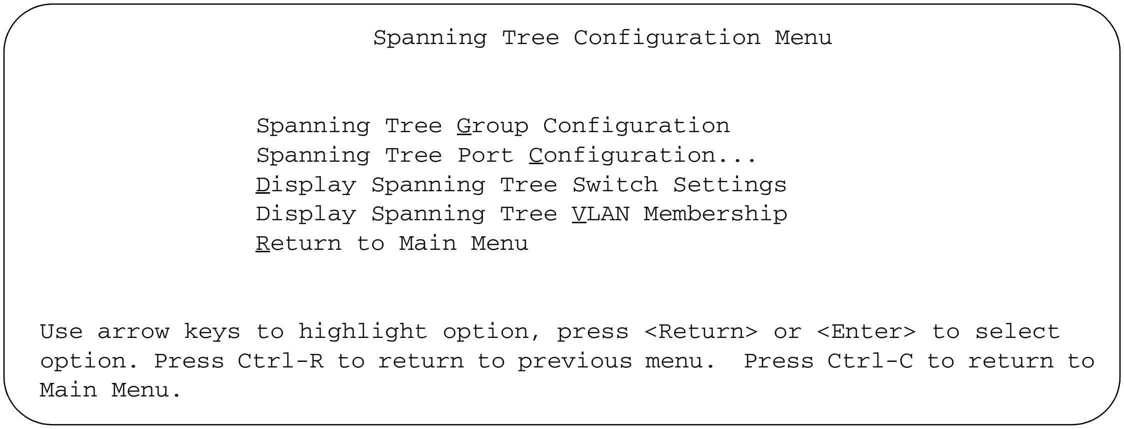

```
                        Spanning Tree Configuration Menu


                  Spanning Tree Group Configuration
                  Spanning Tree Port Configuration...
                  Display Spanning Tree Switch Settings
                  Display Spanning Tree VLAN Membership
                  Return to Main Menu


Use arrow keys to highlight option, press <Return> or <Enter> to select
option. Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to
Main Menu.
```

Table 53 describes the Spanning Tree Configuration Menu screen options.

Table 53 Spanning Tree Configuration Menu screen options

Option	Description
Spanning Tree Group Configuration...	Displays the Spanning Tree Group Configuration screen (see “Spanning Tree Group Configuration Screen”” on page 220).
Spanning Tree Port Configuration...	Displays the Spanning Tree Port Configuration screen (see “Spanning Tree Port Configuration Screen”” on page 224).
Display Spanning Tree Switch Settings	Allows you to display the Spanning Tree Switch Settings screen (see “Spanning Tree Switch Settings Screen”” on page 228).
Display Spanning Tree VLAN Membership	Allows you to display the Spanning Tree VLAN Membership screen (see “Spanning Tree VLAN Membership Screen”” on page 228).

Spanning Tree Group Configuration Screen

The Spanning Tree Group Configuration screen allows you to create and configure spanning tree groups (STGs).

Multiple STGs, up to 8, are available with the Ethernet Routing Switch 3510-24T, and you can configure the VLAN for tagged BPDUs, as well as set the STG multicast MAC address.

To open the Spanning Tree Group Configuration screen:

- ➨ Choose Spanning Tree Group Configuration (or press g) from the Spanning Tree Configuration Menu screen.

Figure 87 shows the Spanning Tree Group Configuration menu.

Figure 87 Spanning Tree Group Configuration menu

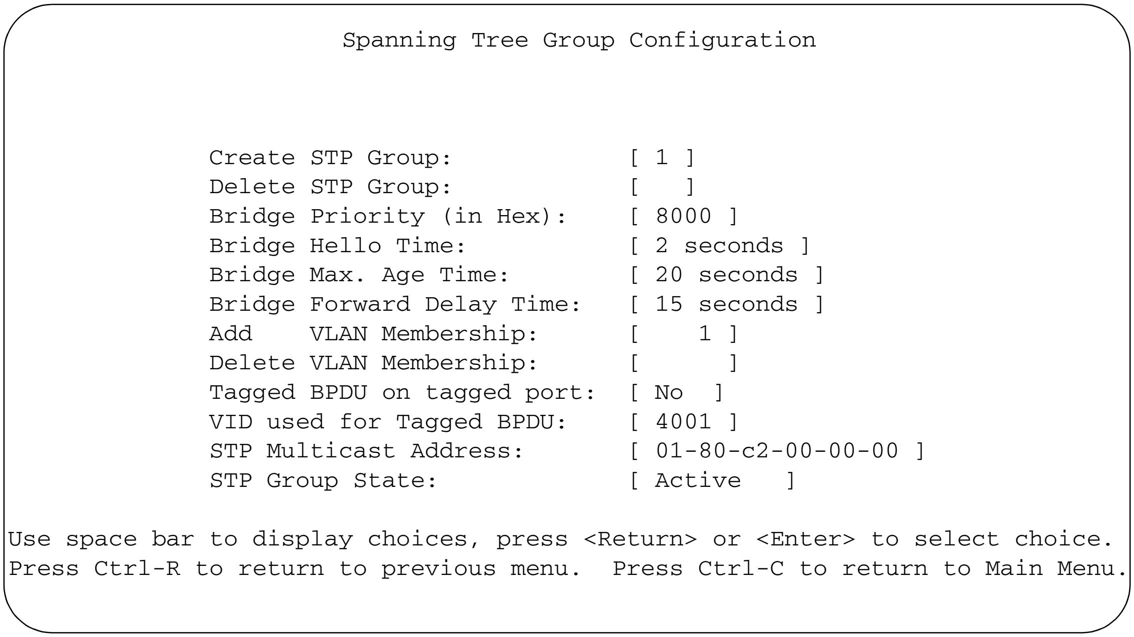

```
                        Spanning Tree Group Configuration



              Create STP Group:             [ 1 ]
              Delete STP Group:             [   ]
              Bridge Priority (in Hex):     [ 8000 ]
              Bridge Hello Time:            [ 2 seconds ]
              Bridge Max. Age Time:         [ 20 seconds ]
              Bridge Forward Delay Time:    [ 15 seconds ]
              Add    VLAN Membership:       [    1 ]
              Delete VLAN Membership:       [      ]
              Tagged BPDU on tagged port:   [ No  ]
              VID used for Tagged BPDU:     [ 4001 ]
              STP Multicast Address:        [ 01-80-c2-00-00-00 ]
              STP Group State:              [ Active   ]

Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 54 describes the Spanning Tree Group Configuration parameters.

Table 54 Spanning Tree Group Configuration parameters

Parameter	Description
Create STP Group	Allows you to create a spanning tree group. You can also use this field to select the STP Group information to display. Default Value 1 Range 1 to 8
Delete STP Group	Allows you to delete a spanning tree group. You cannot delete STP Group number 1, and you can delete only non-active STP Groups. Default Value Blank Range 2 to 8; only created STP Groups are available

Table 54 Spanning Tree Group Configuration parameters (continued)

Parameter	Description	
Bridge Priority	This option for the STP Group allows you to configure the priority value of the bridge ID in hexadecimal notation, which is the most significant byte of the bridge ID. The STA uses this parameter to determine the root bridge (or designated bridge). For example, the bridge with the lowest bridge ID becomes the root bridge, with Bridge Priority values.	
	Default Value	0x8000
	Range	0 to 0xFFFF
Bridge Hello Time	For the STP Group, allows you to configure the Hello Interval (the amount of time between transmissions of BPDUs). This parameter takes effect only when this bridge becomes the root bridge. Note that, although you can set the Hello Interval for a bridge using bridge management software, once the spanning tree computation process is complete, all bridges participating in the spanning tree network use the root bridge's Hello Interval parameter value. If any bridge becomes the root bridge, its Hello Interval parameter value becomes the Actual Hello Interval parameter value for all bridges participating in the spanning tree network. See also Hello Time.	
	Default Value	2 seconds
	Range	1 to 10 seconds
Bridge Max. Age Time	For the STP Group, allows you to configure the maximum age (in seconds) that a Hello message can attain before it is discarded. This parameter takes effect only when the bridge becomes the root bridge. Note that, if this bridge becomes the root bridge, its Maximum Age Time parameter value becomes the Actual Maximum Age Time parameter value for all bridges participating in the spanning tree network. See also Maximum Age Time.	
	Default Value	20 seconds
	Range	6 to 40 seconds
Bridge Forward Delay Time	For the STP Group, allows you to configure the Forward Delay parameter value for this bridge. This parameter takes effect only when this bridge becomes the root bridge. The Forward Delay parameter value specifies the amount of time that the bridge ports remain in the Listening and Learning states before entering the Forwarding state. Note that all bridges participating in the spanning tree network use the root bridge's Forward Delay parameter value. See also Forward Delay.	
	Default Value	15 seconds
	Range	4 to 30 seconds

Table 54 Spanning Tree Group Configuration parameters (continued)

Parameter	Description	
Add VLAN Membership	Allows you to add a VLAN to the specified spanning tree group.	
	Default Value	1
	Range	1 to 4094 **NOTE**: This field is updated with active VLANs currently defined in the system. A newly created and active VLAN is assigned to STP Group 1 by default.
Delete VLAN Membership	Allows you to delete a VLAN from the specified STP group.	
	Default Value	Blank
	Range	1 to 4094; but only configured ones are available **NOTE**: You cannot remove VLAN 1 from STP Group 1.
Tagged BPDU on tagged port	Allows you to choose to send either tagged or untagged BPDUs from a tagged port.	
	Default Value	STP Group 1: No; Other STP Groups: Yes
	Range	No or Yes
VID used for tagged BPDU	Allows you to select the VLAN ID (VID) for tagged BPDU for the specified spanning tree group. **NOTE**: You cannot use the same VID value on more than one STP Group. Also the VID cannot be an active VLAN.	
	Default Value	4001-4008 for STGs 1-8, respectively
	Range	1-4094
STP Multicast Address	Allows you to set the STP Group multicast MAC address. **NOTE**: you can configure the multicast address only for the *first four* STP Groups. You can use any multicast address that starts with 01.	
	Default Value	01-80-c2-00-00-00
STP Group State	Allows you to make the STP Group active or inactive. **NOTE:** You cannot set the default STG, STG1, to InActive. To enable a STP Group, at least one active VLAN must be assigned to that STP Group.	
	Default Value	Active for STG1; InActive for STGs 2 to 8.
	Range	Active or InActive

Spanning Tree Port Configuration Screen

The Spanning Tree Port Configuration screen allows you to set the STG participation for each switch port or all ports and to display spanning tree settings for individual switch ports or all switch ports.

Note: If spanning tree participation of any trunk member is changed (enabled or disabled), the spanning tree participation of all members of that trunk is changed similarly.

Note: Use either the Web-based management system, CLI, or DM to set the spanning tree path cost or priority for individual ports.

Figure 88 shows sample port displays for the two Spanning Tree Port Configuration screens.

➨ Choose Spanning Tree Port Configuration (or press c) from the Spanning Tree Configuration Menu to open the Spanning Tree Port Configuration screen.

Figure 88 Spanning Tree Port Configuration

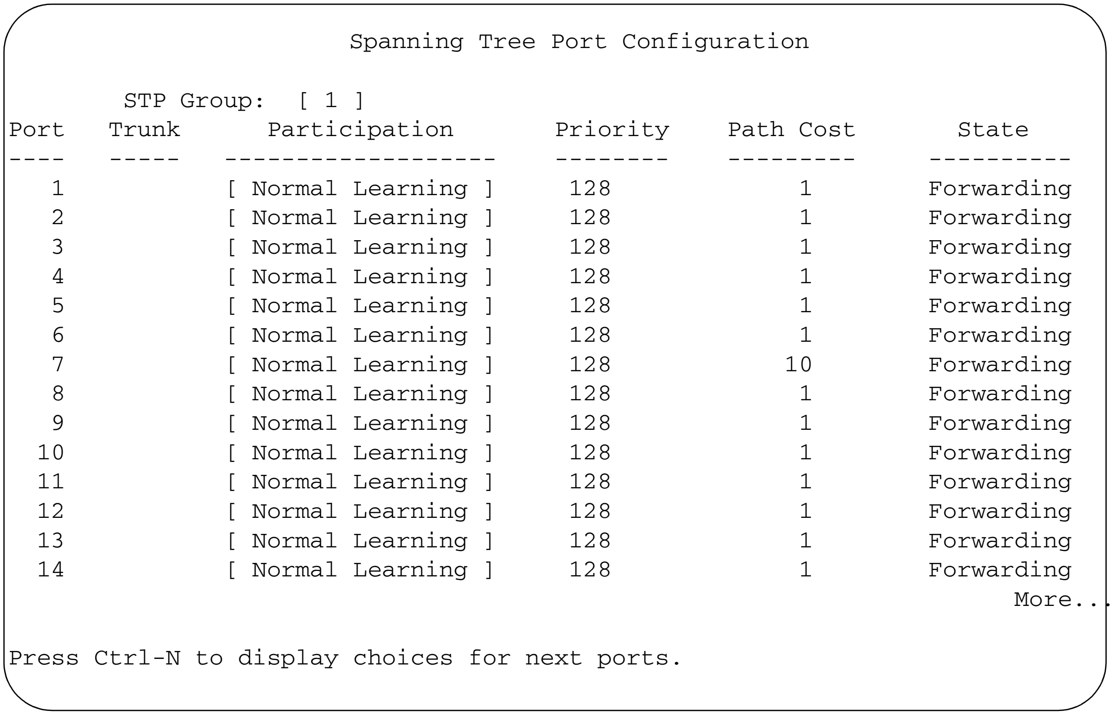

```
                        Spanning Tree Port Configuration

        STP Group:  [ 1 ]
Port   Trunk      Participation       Priority    Path Cost       State
----   -----   -------------------    --------    ---------     ----------
   1           [ Normal Learning ]     128            1         Forwarding
   2           [ Normal Learning ]     128            1         Forwarding
   3           [ Normal Learning ]     128            1         Forwarding
   4           [ Normal Learning ]     128            1         Forwarding
   5           [ Normal Learning ]     128            1         Forwarding
   6           [ Normal Learning ]     128            1         Forwarding
   7           [ Normal Learning ]     128           10         Forwarding
   8           [ Normal Learning ]     128            1         Forwarding
   9           [ Normal Learning ]     128            1         Forwarding
  10           [ Normal Learning ]     128            1         Forwarding
  11           [ Normal Learning ]     128            1         Forwarding
  12           [ Normal Learning ]     128            1         Forwarding
  13           [ Normal Learning ]     128            1         Forwarding
  14           [ Normal Learning ]     128            1         Forwarding
                                                                      More...

Press Ctrl-N to display choices for next ports.
```

Table 55 describes the Spanning Tree Port Configuration screen fields.

Table 55 Spanning Tree Port Configuration screen fields

Field	Description
STP Group	The field allows you to select the number of the spanning tree group (STG) you want to view. To view another STG, type that STG ID number and press [Enter], or press the spacebar on your keyboard to to toggle the STP Group numbers. Default Value 1 Range 1 to 8; only created STP Groups display
Port	Indicates the switch port numbers that correspond to the field values in that row of the screen (for example, the field values in row 2 apply to switch port 2). Note that the values in the *Switch* row affect all switch ports.
Trunk	The read-only data displayed in this column indicates if a particular port is a member of a MultiLink Trunking group.
Participation	Allows you to configure any (or all) of the switch ports for spanning tree participation. When an individual port is a trunk member (see Trunk field), changing this setting for one of the trunk members changes the setting for all members of that trunk. You should consider how this can change your network topology before you change this setting (see Chapters 1 and 2). The Fast Learning parameter is the same as Normal Learning, except that the state transition timer is shortened to 2 seconds. Default Value Normal Learning Range Normal Learning, Fast Learning, Disabled
Priority	This read-only field is a bridge spanning tree parameter that prioritizes the lowest path cost to the root. When one or more ports have the same path cost, spanning tree selects the path with the highest priority (lowest numerical value). See also Path Cost. Default Value 128 Range 0 to 240
Path Cost	This read-only field is a bridge spanning tree parameter that determines the lowest path cost to the root. Default Value 100 for 10 Mbps; 10 for 100 Mbps; 1 for 1000 Mbps Path Cost = 1000/LAN speed (in Mbps) Trunk Link Path Cost = 1000/Total link speed of all trunk links The higher the LAN speed, the lower the path cost. See also Priority. Range 1 to 65535

Table 55 Spanning Tree Port Configuration screen fields (continued)

Field	Description	
State	This read-only field indicates the current port state within the spanning tree network. Each port can transition to various states, as determined by the Participation field setting. For example, when the Participation field is set to disabled, the port does not participate in spanning tree and transitions to the Forwarding state (the default). When the Participation field is set to Normal Learning or Fast Learning, the port transitions from the Disabled state through the Blocking, Listening, and Learning states before entering the Forwarding state.	
	Default Value	Topology dependent
	Range	Disabled, Blocking, Listening, Learning, Forwarding

Note: You can remove a port from the specified STP Group by toggling the Participation field to Disabled.

Spanning Tree Switch Settings Screen

The Display Spanning Tree Switch Settings screen (Figure 89) allows you to view spanning tree parameter values for the selected STP Group. (STP Group 1 is the default STP group.)

To open the Spanning Tree Switch Settings screen:

➨ Choose Spanning Tree Switch Settings (or press d) from the Spanning Tree Configuration Menu screen.

Figure 89 Spanning Tree Switch Settings

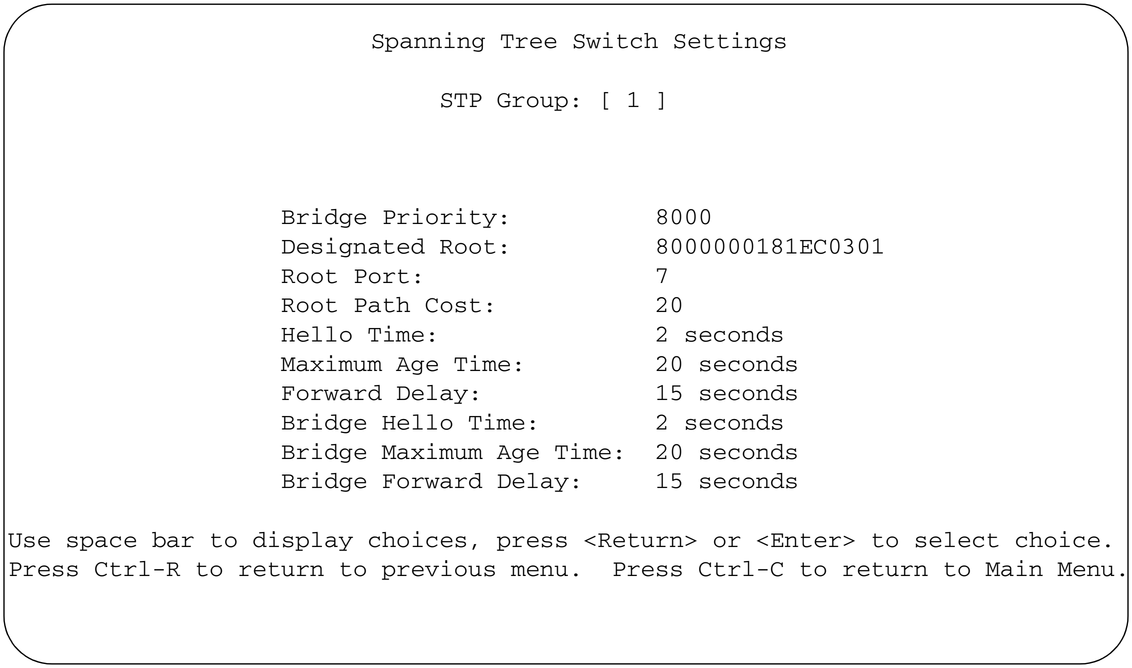

```
                         Spanning Tree Switch Settings

                              STP Group: [ 1 ]



                   Bridge Priority:          8000
                   Designated Root:          8000000181EC0301
                   Root Port:                7
                   Root Path Cost:           20
                   Hello Time:               2 seconds
                   Maximum Age Time:         20 seconds
                   Forward Delay:            15 seconds
                   Bridge Hello Time:        2 seconds
                   Bridge Maximum Age Time:  20 seconds
                   Bridge Forward Delay:     15 seconds

Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 56 describes the Spanning Tree Switch Settings parameters.

Table 56 Spanning Tree Switch Settings parameters

Parameter	Description
STP Group	The field allows you to select the number of the spanning tree group (STG) you want to view. To view another STG, type that STG ID number and press [Enter], or press the spacebar on your keyboard to to toggle the STP Group numbers. Default Value 1 Range 1 to 8; only created STP Groups display
Bridge Priority	This option for STP Group indicates the priority value of the bridge ID in hexadecimal notation, and is the most significant two bytes of the bridge ID. Spanning tree uses this parameter to determine the root bridge (or designated bridge). For example, the bridge with the lowest bridge ID becomes the root bridge, with Bridge Priority values compared first, followed by the hardware addresses.
Designated Root	This option for STP Group indicates the bridge ID of the root bridge, as determined by spanning tree.
Root Port	This option for STP Group indicates the switch port number that offers the lowest path cost to the root bridge.
Root Path Cost	This option for STP Group indicates the path cost to the root bridge.
Hello Time	This option for STP Group indicates the Actual Hello Interval, the amount of time between transmissions of configuration Bridge Protocol Data Units (BPDUs) that the root bridge is currently using. Note that all bridges participating in the spanning tree network use the root bridge's Hello Interval parameter value. See also Bridge Hello Time.
Maximum Age Time	This option for STP Group indicates the Maximum Age Time parameter value that the root bridge is currently using. This value specifies the maximum age that a Hello message can attain before it is discarded. Note that the root bridge's Maximum Age Time parameter value becomes the actual Maximum Age Time parameter value for all bridges participating in the spanning tree network. See also Bridge Maximum Age Time.
Forward Delay	This option for STP Group indicates the Forward Delay parameter value that the root bridge is currently using. This value specifies the amount of time that the bridge ports remain in the Listening and Learning states before entering the Forwarding state. Note that the root bridge's Forward Delay parameter value becomes the actual Forward Delay parameter value for all bridges participating in the spanning tree network. See also Bridge Forward Delay.

Table 56 Spanning Tree Switch Settings parameters (continued)

Parameter	Description
Bridge Hello Time	This option for STP Group indicates the Hello Interval (the amount of time between transmissions of BPDUs) specified by management for this bridge. This parameter takes effect only when this bridge becomes the root bridge. Note that, although you can set the Hello Interval for a bridge using bridge management software, once the spanning tree computation process is complete, all bridges participating in the spanning tree network use the root bridge's Hello Interval parameter value. If any bridge becomes the root bridge, its Hello Interval parameter value becomes the Actual Hello Interval parameter value for all bridges participating in the spanning tree network. See also Hello Time.
Bridge Maximum Age Time	This option for STP Group specifies the maximum age (in seconds) that a Hello message can attain before it is discarded. This parameter, specified by management for this bridge, takes effect only when the bridge becomes the root bridge. Note that, if this bridge becomes the root bridge, its Maximum Age Time parameter value becomes the Actual Maximum Age Time parameter value for all bridges participating in the spanning tree network. See also Maximum Age Time.
Bridge Forward Delay	This option for STP Group indicates the Forward Delay parameter value specified by management for this bridge. This parameter takes effect only when this bridge becomes the root bridge. The Forward Delay parameter value specifies the amount of time that the bridge ports remain in the Listening and Learning states before entering the Forwarding state. Note that all bridges participating in the spanning tree network use the root bridge's Forward Delay parameter value. See also Forward Delay.

Spanning Tree VLAN Membership Screen

The Spanning Tree VLAN Membership screen (Figure 90) allows you to view which VLANs belong to the selected STP Group. (STP Group 1 is the default STP group.)

To open the Spanning Tree VLAN Membership screen:

➨ Choose Spanning Tree VLAN Membership (or press v) from the Spanning Tree Configuration Menu screen.

Figure 90 Spanning Tree VLAN Membership screen

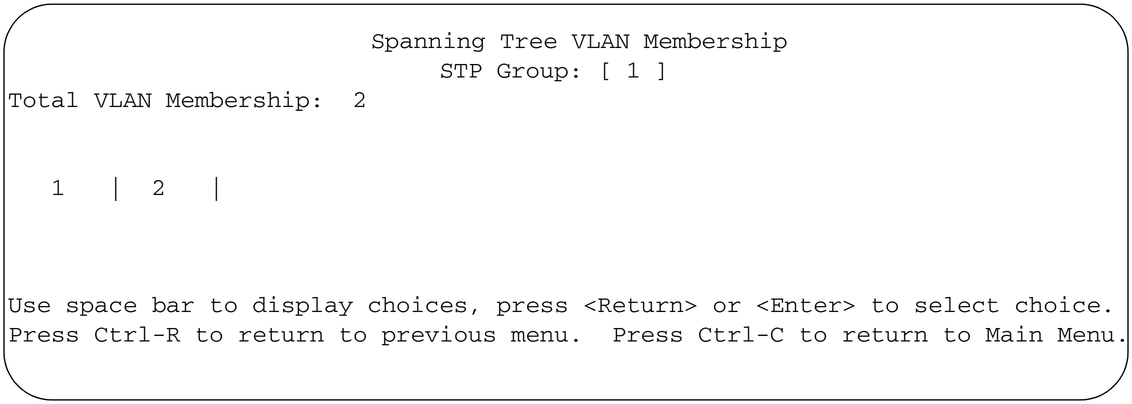

```
                          Spanning Tree VLAN Membership
                               STP Group: [ 1 ]
Total VLAN Membership:  2


   1   |  2   |




Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 57 describes the Spanning Tree VLAN Membership parameters.

Table 57 Spanning Tree VLAN Membership parameters

Parameter	**Description**
STP Group	The field allows you to select the number of the spanning tree group (STG) you want to view. To view another STG, type that STG ID number and press [Enter], or press the spacebar on your keyboard to to toggle the STP Group numbers. Default Value 1 Range 1 to 8; only created STP Groups display
VLAN Membership	Displays the total number of VLANs in the specified STP Group, as well as the VLAN IDs of the VLAN members.

TELNET/SNMP/Web Access Configuration Screen

The TELNET/SNMP/Web Access Configuration screen (Figure 91) allows a user at a remote console terminal to communicate with the Ethernet Routing Switch 3510-24T as if the console terminal were directly connected to it. You can have up to four active Telnet sessions at one time.

To open the TELNET/SNMP/Web Access Configuration screen:

➨ Choose TELNET/SNMP/Web Access Configuration (or press t) from the main menu

Figure 91 TELNET/SNMP/Web Access Configuration screen

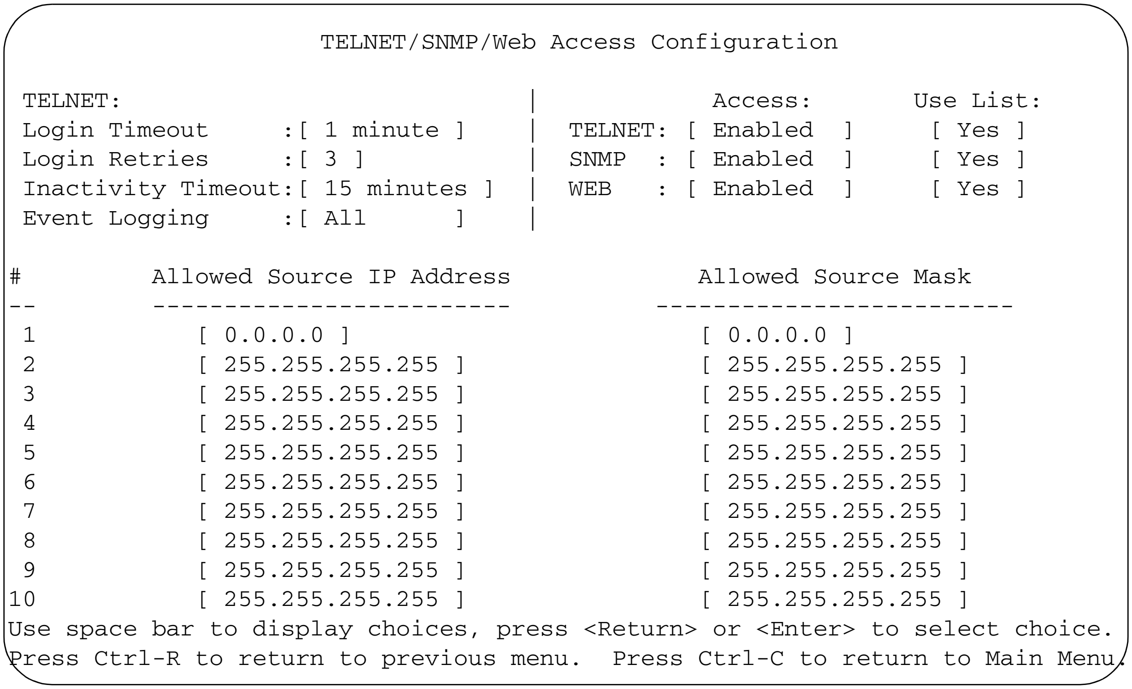

```
                      TELNET/SNMP/Web Access Configuration

 TELNET:                                |          Access:       Use List:
 Login Timeout      :[ 1 minute ]       | TELNET: [ Enabled  ]     [ Yes ]
 Login Retries      :[ 3 ]              | SNMP  : [ Enabled  ]     [ Yes ]
 Inactivity Timeout:[ 15 minutes ]      | WEB   : [ Enabled  ]     [ Yes ]
 Event Logging      :[ All       ]      |

#         Allowed Source IP Address           Allowed Source Mask
--        -------------------------         -------------------------
 1           [ 0.0.0.0 ]                       [ 0.0.0.0 ]
 2           [ 255.255.255.255 ]               [ 255.255.255.255 ]
 3           [ 255.255.255.255 ]               [ 255.255.255.255 ]
 4           [ 255.255.255.255 ]               [ 255.255.255.255 ]
 5           [ 255.255.255.255 ]               [ 255.255.255.255 ]
 6           [ 255.255.255.255 ]               [ 255.255.255.255 ]
 7           [ 255.255.255.255 ]               [ 255.255.255.255 ]
 8           [ 255.255.255.255 ]               [ 255.255.255.255 ]
 9           [ 255.255.255.255 ]               [ 255.255.255.255 ]
10           [ 255.255.255.255 ]               [ 255.255.255.255 ]
Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 58 describes the TELNET/SNMP/Web Access Configuration screen fields.

Table 58 TELNET/SNMP/Web Access Configuration screen fields

Field	Description
TELNET Access	Allows a user remote access to the management systems through a Telnet session. Default Value: Enabled Range: Enabled, Disabled
Login Timeout	Specifies the amount of time a user has to enter the correct password at the console-terminal prompt. Default Value: 1 minute Range: 0 to 10 minutes (0 indicates “no timeout”)
Login Retries	Specifies the number of times a user can enter an incorrect password at the console-terminal prompt before terminating the session. Default Value: 3 Range: 1 to 100
Inactivity Timeout	Specifies the amount of time the session can be inactive before it is terminated. Default Value: 15 minutes Range: 0 to 60 minutes (0 indicates “no timeout”)
Event Logging	Specifies the types of events that will be displayed in the Event Log screen (see “System Log Screen”). Default Value: All Range: All, None, Accesses, Failures Description: *All:* Logs the following Telnet events to the Event Log screen: • TELNET connect: Indicates the IP address and access mode of a Telnet session. • TELNET disconnect: Indicates the IP address of the remote host and the access mode, due to either a logout or inactivity. • Failed TELNET connection attempts: Indicates the IP address of the remote host whose IP address is not on the list of allowed addresses, or indicates the IP address of the remote host that did not supply the correct password. *None:* Indicates that no Telnet events will be logged in the Event Log screen. *Accesses*: Logs only Telnet connect and disconnect events in the Event Log screen. *Failures:* Logs only failed Telnet connection attempts in the Event Log screen.

Table 58 TELNET/SNMP/Web Access Configuration screen fields (continued)

Field	**Description**	
TELNET Access	Specifies if Telnet access is allowed and only to those on the list.	
	Default Value:	Access: Enabled; Use List: Yes
	Range:	Access: Enabled, Disabled; Use List: Yes, No
SNMP Access	Specifies if SNMP access is allowed and only to those on the list. (SNMP access includes the DM system.)	
	Default Value:	Access: Enabled; Use List: Yes
	Range:	Access: Enabled, Disabled; Use List: Yes, No
WEB Access	Specifies if access to the Web-based management system is allowed and only to those on the list.	
	Default Value:	Access: Enabled; Use List: Yes
	Range:	Access: Enabled, Disabled; Use List: Yes, No
Allowed Source IP Address	Specifies up to 10 user-assigned host IP addresses that are allowed Telnet access to the management systems.	
	Default Value:	0.0.0.0 (no IP address assigned)
	Range:	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point
Allowed Source Mask	Specifies up to 10 user-assigned allowed source address masks. The remote IP address is masked with the Allowed Source Mask and, if the resulting value equals the Allowed Source IP address, the connection is allowed. For example, a connection would be allowed with the following settings: Remote IP address = 192.0.1.5 Allowed Source IP Address = 192.0.1.0 Allowed Source Mask = 255.255.255.0	
	Default Value:	0.0.0.0 (no IP mask assigned)
	Range:	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point

Software Download Screen

The Software Download screen (Figure 92) allows you to revise the Ethernet Routing Switch 3510-24T software image that is located in nonvolatile flash memory.

Caution: Do not interrupt power to the device during the software download process. If the power is interrupted, the firmware image can become corrupted.

Achtung: Unterbrechen Sie die Stromzufuhr zum Gerät nicht, während die Software heruntergeladen wird. Bei Unterbrechung der Stromzufuhr kann das Firmware-Image beschädigt werden.

Attention: Ne pas couper l'alimentation de l'appareil pendant le chargement du logiciel. En cas d'interruption, le programme résident peut être endommagé.

Precaución: No interrumpa la alimentación del dispositivo durante el proceso de descarga del software. Si lo hace, puede alterar la imagen de la programación (firmware).

Attenzione: Non interrompere l'alimentazione elettrica al dispositivo durante il processo di scaricamento del software. In caso di interruzione, l'immagine firmware potrebbe danneggiarsi.

注意：ソフトウェアをダウンロードしているとき、ディバイスへの電源を切らないでください。電源を切ると、ファームウェアのイメージを損う恐れがあります。

To download the software image, you need a properly configured Trivial File Transfer Protocol (TFTP) server in your network, and an IP address for the switch. To learn how to configure the switch IP address, refer to “IP Configuration/Setup Screen”.

This section covers the following topics:

- “Using the Software Download Screen”, next
- “LED Indications During the Download Process”

Using the Software Download Screen

To open the Software Download screen:

➨ Choose Software Download (or press f) from the main menu.

The Software Download screen appears (Figure 92).

You can monitor the software download process by observing the LEDs (see “LED Indications During the Download Process”).

Figure 92 Software Download screen

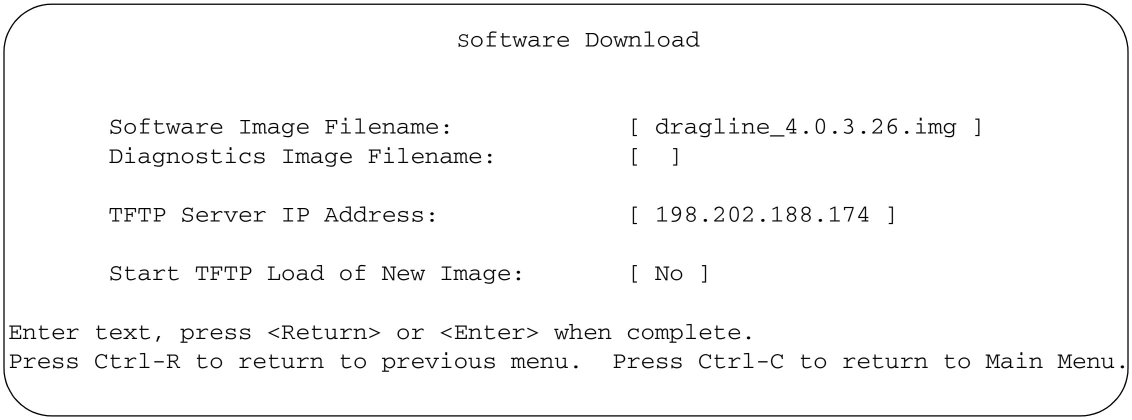

```
                              Software Download


       Software Image Filename:              [ dragline_4.0.3.26.img ]
       Diagnostics Image Filename:           [  ]

       TFTP Server IP Address:               [ 198.202.188.174 ]

       Start TFTP Load of New Image:         [ No ]

Enter text, press <Return> or <Enter> when complete.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 59 describes the Software Download screen fields.

Table 59 Software Download screen fields

Field	**Description**	
Software Image Filename	The Ethernet Routing Switch 3510-24T software image load file name.	
	Default Value	Zero-length string
	Range	An ASCII string of up to 30 printable characters
Diagnostics Filename	The Ethernet Routing Switch 3510-24T diagnostics file name.	
	Default Value	Zero-length string
	Range	An ASCII string of up to 30 printable characters
TFTP Server IP Address	The IP address of your TFTP load host.	
	Default Value	0.0.0.0 (no IP address assigned)
	Range	Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point
Start TFTP Load of New Image	Specifies whether to start the download of the switch software image (default is No). Use the spacebar to toggle the selection to the one you want. Press [Enter] to initiate the software download process.	
	NOTE: The software download process can take up to 60 seconds to complete (or more if the load host path is congested or there is a high volume of network traffic).	
	To ensure that the download process is not interrupted, do not power down the switch for approximately 10 minutes.	
	Default Value	No
	Range	No, Ethernet Routing Switch 3510-24T Image, Ethernet Routing Switch 3510-24T Diags, andEthernet Routing Switch 3510-24T If Newer

Note: If your station cannot ping the TFTP server during the downloading process, you may receive the following message:

`Image is Invalid`

Actually, the problem is that the TFTP server is not reachable, rather than any problems with the image.

LED Indications During the Download Process

When you download software to the Ethernet Routing Switch 3510-24T, the port LEDs light one after another in a chasing pattern (except the LEDs on ports 11, 12, 23, and 24 on a Ethernet Routing Switch 3510-24T).

While downloading the image, the pattern is fast, and then the pattern slows as the switch erases the flash memory. The pattern moves very fast as the switch programs the new image into the switch's memory. When the process is complete, the port LEDs are no longer lit and the switch resets.

The software download process automatically completes without user intervention. The process erases the contents of flash memory and replaces it with a new software image. Be careful not to interrupt the download process until after it runs to completion (the process can take up to 10 minutes, depending on network conditions).

Note: If problems occur during the software download process, refer to Chapter 6.

During the download process, the Ethernet Routing Switch 3510-24T is not operational. You can monitor the progress of the download process by observing the LED indications.

Configuration File Menu Screen

The Configuration File Menu screen (Figure 93) allows you to upload and download the configuration parameters of a Ethernet Routing Switch 3510-24T to a TFTP server. You can also download an ASCII configuration file from a TFTP server.

These options allow you to store your switch configuration parameters on a TFTP server. You can retrieve the configuration parameters of a standalone switch and use the retrieved parameters to automatically configure a replacement switch. You must set up the file on your TFTP server and set the filename read/write permission to enabled before you can save the configuration parameters.

To open the Configuration File Menu screen:

➨ Choose Configuration File Menu (or press g) from the main menu.

Figure 93 Configuration File Menu screen

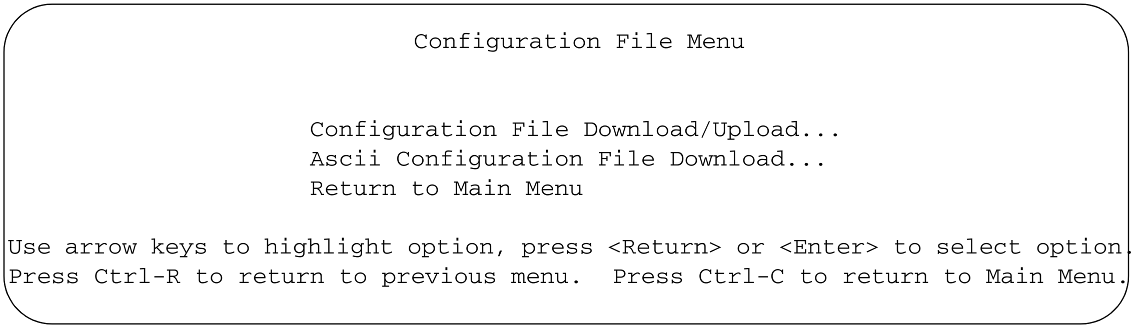

```
                              Configuration File Menu


                      Configuration File Download/Upload...
                      Ascii Configuration File Download...
                      Return to Main Menu

Use arrow keys to highlight option, press <Return> or <Enter> to select option.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 60 describes the Configuration File Menu screen options.

Table 60 Configuration File Menu screen options

Option	Description
Configuration File Download/Upload...	Displays the Configuration File Download/Upload screen (see “Configuration File Download/Upload Screen”” on page 240).
Ascii Configuration File Download...	Displays the ASCII Configuration File Download screen (see “ASCII Configuration File Download Screen”” on page 242).

Configuration File Download/Upload Screen

The Configuration File Download/Upload screen (Figure 94) allows you to store your switch configuration parameters on a TFTP server. Certain requirements apply when automatically configuring a switch using this feature (see “Requirements”” on page 242). Although most configuration parameters are saved to the configuration file, certain parameters are not saved (see Table 62 on page 242).

Choose Configuration File Download/Upload from the Configuration File Menu to open the Configuration File Download/Upload screen.

Figure 94 Configuration File Download/Upload screen

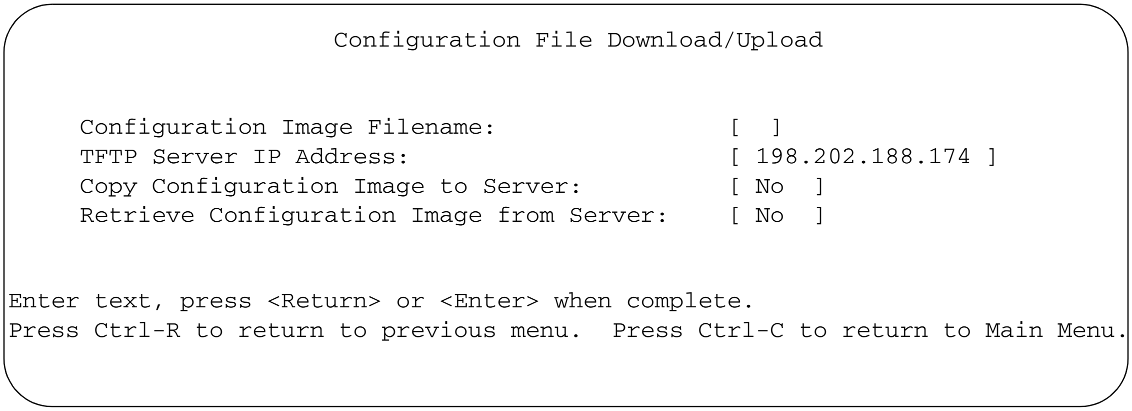

```
                      Configuration File Download/Upload


     Configuration Image Filename:                    [  ]
     TFTP Server IP Address:                          [ 198.202.188.174 ]
     Copy Configuration Image to Server:              [ No   ]
     Retrieve Configuration Image from Server:        [ No   ]


Enter text, press <Return> or <Enter> when complete.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 61 describes the Configuration File Download/Upload screen fields.

Table 61 Configuration File Download/Upload screen fields

Field	Description
Configuration Image Filename	The file name you have chosen for the configuration file. Choose a meaningful file name that will allow you to identify the file for retrieval when required. The file must already exist on your TFTP server and must be read/write enabled, as appropriate. Default Value Zero-length string Range An ASCII string of up to 30 printable characters
TFTP Server IP Address	The IP address of your TFTP load host. Default Value 0.0.0.0 (no IP address assigned) Range Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point
Copy Configuration Image to Server	Specifies whether to copy the presently configured switch parameters to the specified TFTP server (default is No). Use the spacebar to toggle the selection to Yes. Press [Enter] to initiate the process. Default Value No Range Yes, No
Retrieve Configuration Image from Server	Specifies whether to retrieve the stored switch configuration parameters from the specified TFTP server (default is No). If you choose Yes, the download process begins immediately and, when completed, causes the switch to reset with the new configuration parameters. Use the spacebar to toggle the selection to Yes. Press [Enter] to initiate the process. Default Value No Range Yes, No

Requirements

The following requirements apply to the Configuration File feature:

- The Configuration File feature can only be used to copy standalone switch configuration parameters to other standalone switches.
- A configuration file obtained from a standalone switch can only be used to configure other standalone switches that have the same firmware revision and model type as the donor standalone switch.

Table 62 describes Configuration File parameter information.

Table 62 Parameters not saved to the binary Configuration File

These parameters are not saved:	Used in this screen:	See page:
In-Band Switch IP Address		
In-Band Subnet Mask		
Default Gateway		
Console Read-Only Switch Password	Console/Comm Port Configuration	212
Console Read-Write Switch Password		
Configuration Image Filename	Configuration File Download/Upload	240
TFTP Server IP Address		

ASCII Configuration File Download Screen

The ASCII Configuration File Download screen (Figure 95) allows you to download an ASCII configuration file containing CLI commands from a TFTP server to configure the switch.

➨ Choose ASCII Configuration File Download (or press a) from the Configuration File Menu to open the ASCII Configuration File Download screen.

Figure 95 ASCII Configuration File Download screen

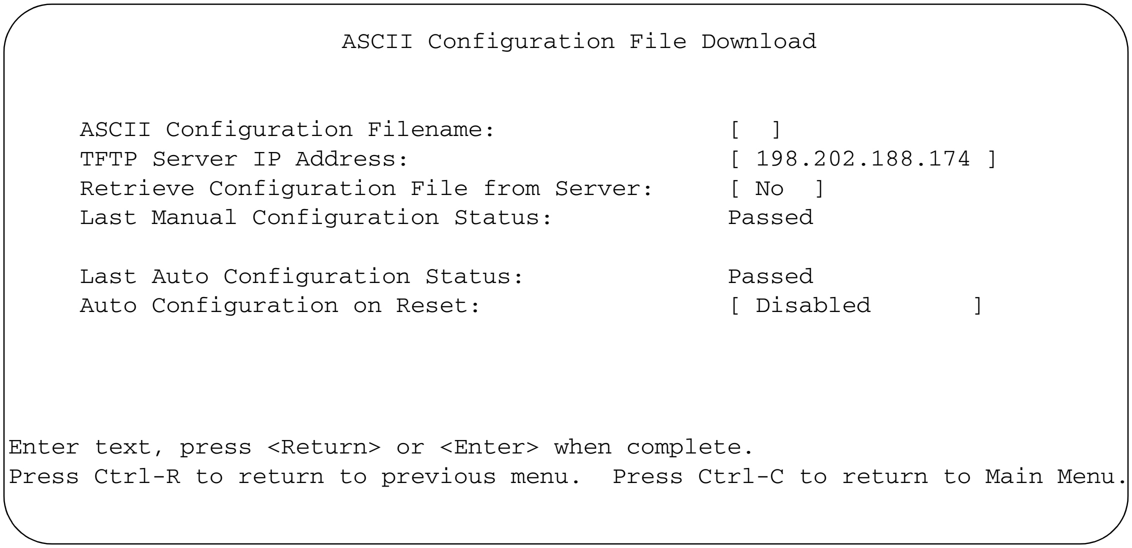

```
                         ASCII Configuration File Download


     ASCII Configuration Filename:                  [  ]
     TFTP Server IP Address:                        [ 198.202.188.174 ]
     Retrieve Configuration File from Server:       [ No  ]
     Last Manual Configuration Status:              Passed

     Last Auto Configuration Status:                Passed
     Auto Configuration on Reset:                   [ Disabled       ]




Enter text, press <Return> or <Enter> when complete.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 63 describes the ASCII Configuration File Download screen fields.

Table 63 ASCII Configuration File Download screen fields

Field	Description
ASCII Configuration Filename	Enter the file name you have chosen for the ASCII configuration file. Choose a meaningful file name that will allow you to identify the file for retrieval when required. The file must already exist on your TFTP server and must be read/write enabled. Default Value Zero-length string Range An ASCII string of up to 30 printable characters
TFTP Server IP Address	The IP address of your TFTP load host. Default Value 0.0.0.0 (no IP address assigned) Range Four-octet dotted-decimal notation, where each octet is represented as a decimal value, separated by a decimal point
Retrieve Configuration File from Server	Specifies whether to retrieve the stored switch ASCII configuration file from the specified TFTP server (default is No). If you choose Yes, the download process begins immediately and, when completed, causes the switch to be configured according to the CLI commands in the file. Use the spacebar to toggle the selection to Yes. Press [Enter] to initiate the process. Default Value No Range Yes, No
Last Manual Configuration Status	The system displays if the last manual configuration passed or failed. Default Value Passed Range Passed, Failed
Last Auto Configuration Status	The system displays if the last automatic configuration passed or failed. Default Value Passed Range Passed, Failed
Auto Configuration on Reset	Allows you to choose to Disabled, Use Configured, or Use BootP: • Disabled—Auto configuration on reset is disabled. • Use Configured—Use manually configured ASCII configuration filename and TFTP server address for auto configuration on reset. • Use BootP—Retrieve ASCII configuration filename, and optionally server address, using BootP, when BootP is enabled, and perform auto configuration on reset using these parameters. Note: Refer to Appendix E for a sample BootP configuration file. Default Value Disabled Range Disabled, Use Configured, Use BootP

System Log Screen

The System Log screen (Figure 96) displays or clears messages obtained from system nonvolatile random access memory (NVRAM) or dynamic random access memory (DRAM) and NVRAM.

System Log messages operate as follows:

- NVRAM messages are retrievable after a system reset.
- DRAM messages can be viewed while the system is operational.
- All NVRAM and DRAM messages are time stamped.
- When you restart your system after a reset, the DRAM messages are deleted.
- After a reset, all messages stored in NVRAM are copied to DRAM (DRAM messages are not copied to NVRAM). The messages copied to DRAM are time stamped to zero (0).

To open the System Log screen:

➨ Choose Display System Log (or press y) from the main menu.

Figure 96   System Log screen

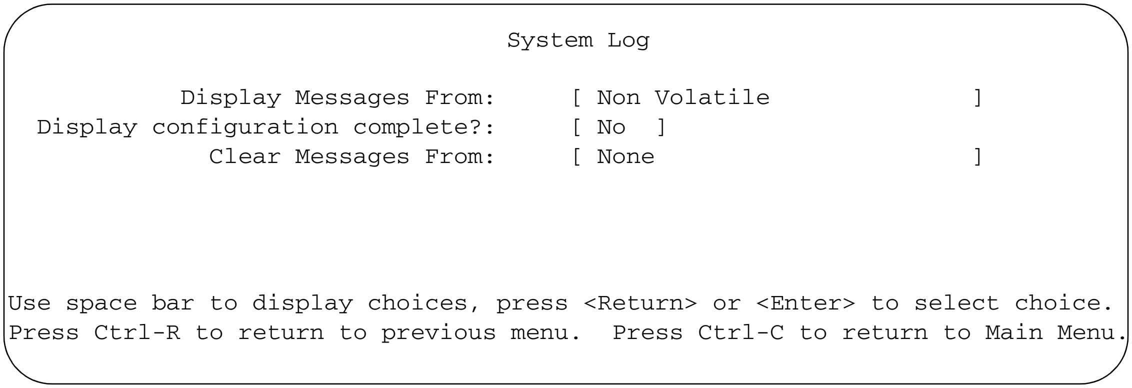

```
                                 System Log

           Display Messages From:      [ Non Volatile              ]
 Display configuration complete?:      [ No  ]
             Clear Messages From:      [ None                      ]




Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu.  Press Ctrl-C to return to Main Menu.
```

Table 64 describes the System Log screen fields.

Table 64 System Log screen fields

Field	Description
Display Messages From	This field allows you to select the RAM source your messages are obtained from. Choose Non Volatile (NVRAM), or Volatile (DRAM) + Non Volatile. Use the spacebar to toggle between the options. Default Non Volatile Range Non Volatile, Volatile + Non Volatile
Display configuration complete?	This field allows you to determine whether the configuration information received from NVRAM/DRAM (depending on what is selected in the Display Messages From field) is complete. Use the spacebar to toggle between the options. Default No Range No, Yes
Clear Messages From	This field allows you to clear the information messages from DRAM, NVRAM or both. If you clear DRAM messages, existing NVRAM messages are copied into DRAM. After a system reset, all existing NVRAM messages are copied to DRAM. Use the spacebar to toggle between the options. Default None Range Volatile, Volatile + Non Volatile.

Chapter 4
Policy-enabled Networks

This chapter provides an overview of the Differentiated Services (DiffServ) Quality of Service (QoS) network architecture. The Ethernet Routing Switch 3510-24T provides a Web-based management interface, a Nortel Command Line Interface (NCLI)**,** and the graphical user interface Device Manager (DM) to configure QoS. Refer to *Web Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3, Switch Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3, NCLI Configuration Guide for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3* for detailed information.

The components of QoS are discussed in the remainder of this chapter, which includes information about the following topics:

- "Summary", next
- "QoS Overview" on page 249
- "QoS Components" on page 250
- "Specifying Interface Groups" on page 252
- "Rules" on page 253
- "Specifying Actions" on page 257
- "Specifying Interface Action Extensions" on page 259
- "Specifying Meters" on page 260
- "Specifying Policies" on page 264
- "Packet Flow Using QoS" on page 266
- "Queue Sets" on page 268
- "Modifying Queue Set Characteristics" on page 269

Summary

Policy-enabled networks allow system administrators to prioritize the network traffic, thereby providing better service for selected applications. Using Quality of Service (QoS), the system administrators can establish service level agreements (SLAs) with customers of the network.

In general, QoS helps with two network problems: bandwidth and time-sensitivity. QoS can help you allocate bandwidth to critical applications, and you can limit bandwidth for less critical applications. Applications such as video and voice must have a certain amount of bandwidth to work correctly; using QoS, you can provide that bandwidth when necessary. Also, you can put a high priority on applications that are sensitive to timing or cannot tolerate delay by assigning that traffic to a high-priority queue.

Nortel uses Differentiated Services (DiffServ) to provide QoS functionality. A DiffServ architecture enables service discrimination of traffic flows by offering network resources to higher classes at the expense of lower classes of service. This architecture allows you to prioritize or to aggregate flows and provides Quality of Service (QoS) that is scalable.

Briefly, with DiffServ, you use policies to identify traffic to forward or drop, meter, re-mark, and assign to certain interfaces. The system marks the DiffServ (DS) field of IP packets to define how the packet is treated as it moves through the network. You identify, meter, and re-mark the traffic to facilitate flow prioritization. You can specify a number of policies, and each policy can match one or many flows—supporting complex classification scenarios.

QoS Overview

Differentiated services (DiffServ) is a Quality of Service (QoS) network architecture that offers varied levels of service for different types of data traffic. DiffServ lets you designate a specific level of performance on a packet-by-packet basis instead of using the “best-effort” model for your data delivery. You can give preferential treatment (prioritization) to applications that require high performance and reliable service, such as voice and video over IP.

To differentiate between traffic flows, the DiffServ (DS) field, as defined in RFCs 2474 and 2475, is marked. The DS field in the IP header is an octet, and the first six bits, called the DS codepoint (DSCP), are used in the DiffServ architecture. The DSCP marking dictates the forwarding treatment given to the packet at each hop. This marking occurs at the edge of the DiffServ domain and is based on the policy or filter for the particular microflow or an aggregate flow. The QoS system also can interact with 802.1p and Layer 2 QoS.

Within the DiffServ network, the marked packets are placed in a queue according to their marking, which in turn determines the per-hop-behavior (PHB) of that packet. For example, if a video stream is marked so that it receives the highest priority, then it is placed in a high-priority queue. As those packets traverse the DiffServ network, the video stream is forwarded before any other packets.

To ensure that the traffic stream conforms to the bandwidth assigned, policing within the network is necessary.

DiffServ Concepts

DiffServ is described in IETF RFCs 2474 and 2475. This architecture is flexible and allows for either end-to-end QoS or intradomain QoS by implementing complex classification and mapping functions at the network boundary or access points. Within a DiffServ domain, the packet treatment is regulated by this classification and mapping.

The DiffServ basic elements are implemented within the network and include:

- Packet classification functions
- A small set of per-hop forwarding behaviors
- Traffic metering and marking

Traffic is classified as it enters the DS network and is then assigned the appropriate PHB based on that classification. Within the IP packet, the 6 bits in the DSCP are marked to identify how the packet should be treated at each subsequent network node.

DiffServ assumes the existence of a Service Level Agreement (SLA) between DS domains that share a border. The SLA defines the profile for the aggregate traffic flowing from one network to the other based on policy criteria. In a given traffic direction, the traffic is expected to be metered at the ingress point of the downstream network.

As the traffic moves within the DiffServ network, policies ensure that traffic marked by the different DSCPs is treated according to that marking.

QoS Components

The Ethernet Routing Switch 3510-24T supports the following Nortel QoS classes:

- Critical and Network classes have the highest priority over all other traffic.
- Premium class is an end-to-end service functioning similarly to a virtual leased line. Traffic in this service class is normally guaranteed an agreed-upon peak bandwidth. Traffic requiring this service should be shaped at the network boundary in order to undergo a negligible delay and delay variance. This service class is suitable for real-time applications like video and voice over IP. The recommended PHB for this service is the Expedited Forwarding (EF) PHB.
- Platinum, Gold, Silver, and Bronze classes use the Assured Forwarding (AF) PHB. These classes are used for real-time, delay-tolerant traffic and non-real-time, mission-critical traffic.
- Standard class is the best-effort IP service with an additional, optional use of traffic classification that is used at the network boundary to request a better effort treatment for packets that are in-profile (packets that do not break the service agreements between the user and the service provider).

Table 65 describes the service classes and the required treatment.

Table 65 Service classes

Traffic category	Service class	Application type	Required treatment
Critical network control	Critical	Critical network control traffic	Highest priority over all other traffic. Guaranteed minimum bandwidth.
Standard network control	Network	Standard network control traffic	Priority over user traffic. Guaranteed minimum bandwidth.
Real time, delay intolerant, fixed bandwidth	Premium	Interhuman communications requiring interaction (such as VoIP).	Absolute bounded priority over user traffic. No packet loss for in-profile traffic. Virtual leased line with lowest amount of latency. Provisioned for peak rate.
Real time, delay tolerant, low variable bandwidth	Platinum	Interhuman communications requiring interaction with additional minimal delay (such as low-cost VoIP).	Higher-priority scheduling providing guaranteed minimum provisioned bandwidth. Competes for additional bandwidth.
Real time, delay tolerant, high variable bandwidth	Gold	Single human communication with no interaction (such as Web site streaming video).	High-priority scheduling providing guaranteed minimum provisioned bandwidth. Competes for additional bandwidth.
Non-real time, mission critical, interactive	Silver	Transaction processing (such as Telnet, Web browsing).	Medium priority scheduling providing guaranteed minimum provisioned bandwidth. Competes for additional bandwidth.
Non-real time, mission critical, non-interactive	Bronze	For example, E-mail, FTP, SNMP.	Lower-priority scheduling providing guaranteed minimum provisioned bandwidth. Competes for additional bandwidth.
Non-real time, non-mission critical	Standard	Bulk transfer (such as large FTP transfers, after-hours tape backup).	Best effort delivery. Uses remaining available bandwidth.

Specifying Interface Groups

Every port should be assigned to an interface group, which is used to apply policies to traffic received by this port. And, each port can belong to only *one* interface group. The Web-based interface for QoS uses the term “Interface Configurations” for this function. One policy references only one interface group, but you can configure several policies to reference the same interface group.

All ports that have the same interface group (role combination) have the same set of classification elements installed on them. When you move a port to another interface group (role combination), the classification elements associated with the previous interface group are removed and the classifications elements associated with the new interface group are installed on the port.

Note: If you assign a port that is part of a MultiLink Trunk (MLT) to an interface group, only that port joins the interface group. The other ports in the MLT do *not* become part of the interface group (role combination) automatically.

At factory default, ports are assigned to the default interface group (role combination), which is named allBayStackIfcs.

Each port is associated with the default interface group until you either associate the port with another interface group or remove the port from all interface groups. Ports that are *not* associated with any interface group are disabled for QoS; they remain disabled across reboots until you either assign that port to an interface group or reset the switch to factory defaults **(when it will be reassigned to allBayStackIfcs.)**

Note: You must remove all ports from an interface group in order to delete it.You cannot delete an interface group that is referenced by a policy.

Rules

Packet classifiers identify packets according to a particular content in the packet header such as the source address, destination address, source port number, destination port number, and other data. Packet classifiers identify flows for additional processing.

You can use two types of classifier elements to construct a classifier:

- Layer 2 (L2) classifier elements
- IP classifier elements

Classifier Definition

A classifier is made up of one or more classifier elements. The classifier elements dictate the classification criteria of the classifiers. Only one element of each type, IP or Layer 2, can be used to construct a classifier.

Figure 97 displays the relationship between the classifier elements, classifiers, and classifier blocks.

Figure 97 Relationship of classifier elements, classifiers, and classifier blocks

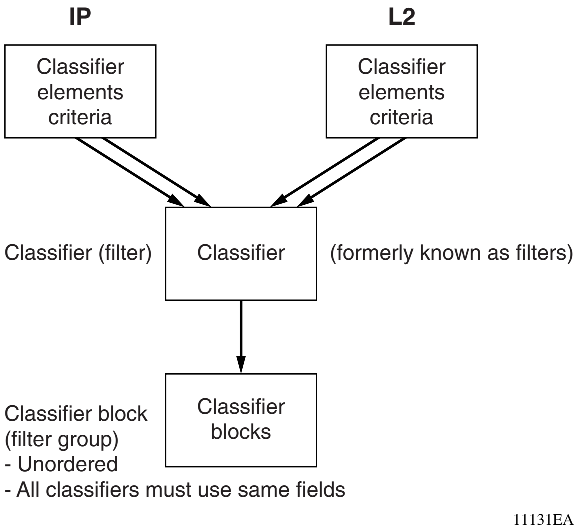
IP
L2
Classifier elements criteria
Classifier elements criteria
Classifier (filter)
Classifier
(formerly known as filters)
Classifier block (filter group)
- Unordered
- All classifiers must use same fields
Classifier blocks
11131EA

The system automatically creates some classifiers on trusted and untrusted ports. Additionally, you can define and create classifiers.

IP Classifier Elements

The Ethernet Routing Switch 3510-24T classifies packets based on various parameters in the IP header:

- IPv4/IPv6 address type
- IPv6 flow identifier
- IPv4/IPv6 source address/mask
- IPv4/IPv6 destination address/mask
- IPv4 protocol type/IPv6 next-header
- IPv4/IPv6 DSCP value
- IPv4/IPv6 Layer 4 source port number with TCP/UDP (range of)
- IPv4/IPv6 Layer 4 destination port number (range of)

Layer 2 Classifier Elements

The Ethernet Routing Switch 3510-24T classifies packets based on various parameters in the Layer 2 header:

- source MAC address/mask
- destination MAC address/mask
- VLAN ID number (range of)
- VLAN tag
- EtherType
- IEEE 802.1p user priority values

Note: Layer 2 classifier elements with an Ethernet Type of 0x0800 are treated as an IPv4 classifier and those with an Ether Type of 0x86DD are treated as an IPv6 classifier.

Classifiers and Classifier Blocks

You may combine classifier elements into classifiers, and you group classifiers into classifier blocks. You create classifiers by referencing an IP classifier element, an L2 classifier element, or one of each type.

Each classifier can have a maximum of a single IP classifier element plus a single L2 classifier element. You cannot put more than one IP classifier element or more than one L2 classifier element into one classifier. A classifier may contain one IP classifier element and one L2 classifier element, or one classifier element of each type —but no more. That is, you can have one (and only one) of *either*:

- one L2 classifier element
- one IP classifier element
- one L2 classifier element plus one IP classifier element

You combine classifiers into classifier blocks. Each classifier block has one or more classifiers.

As you plan your classifier blocks, keep in mind that you can have only a single IP classifier element plus a single L2 classifier element in each classifier. For example, to group five IP classifier elements, you create five separate classifiers, each with a unique IP classifier element, and then create a classifier block referencing those five classifiers.

All classifiers that are a part of a single classifier block—that is, with the same block number—must each filter on identically the same parameters at the packet level. This includes the same mask, range, and VLAN tag type. If this criterion is not met, you get an error message when you attempt to create the classifier block or to add a new member to an existing block. Also, if one of the classifier elements in a classifier block has associated actions or meters, then all classifier elements of that classifier block must also have associated actions or meters (not identical actions or meters, but also associated actions or meters).

A classifier or classifier block is associated through a policy with interface groups. Packets received from any port that is in an interface group are classified with the same filter criteria.

Each classifier/classifier block is associated with actions that are executed when the packet matches the filter criteria in the group. The filter criteria and the associated actions, metering criteria, and interface groups are referenced by a policy, which dictates the overall traffic treatment (refer to Figure 98 for an illustration of the traffic treatment).

Classifier elements via individual classifiers or a classifier block are associated with an interface group, action, and metering through a policy. Multiple policies may be applied to a given flow and the policy evaluation order is determined by the policy precedence. The order of precedence is from the highest precedence value to the lowest precedence (that is, a value of 8 is evaluated before a value of 7).

Note: Classifier blocks, not individual classifiers that comprise a block, can be associated with a meter/action.

In summary, classifiers combine different classifier elements. Classifier blocks combine classifiers to form an unordered set of classification data. Unordered data means that all classifiers associated with a policy are applied as if simultaneously, with no precedence.

Specifying Actions

Figure 98 summarizes how QoS matches packets with actions.

Figure 98 Flowchart of QoS actions

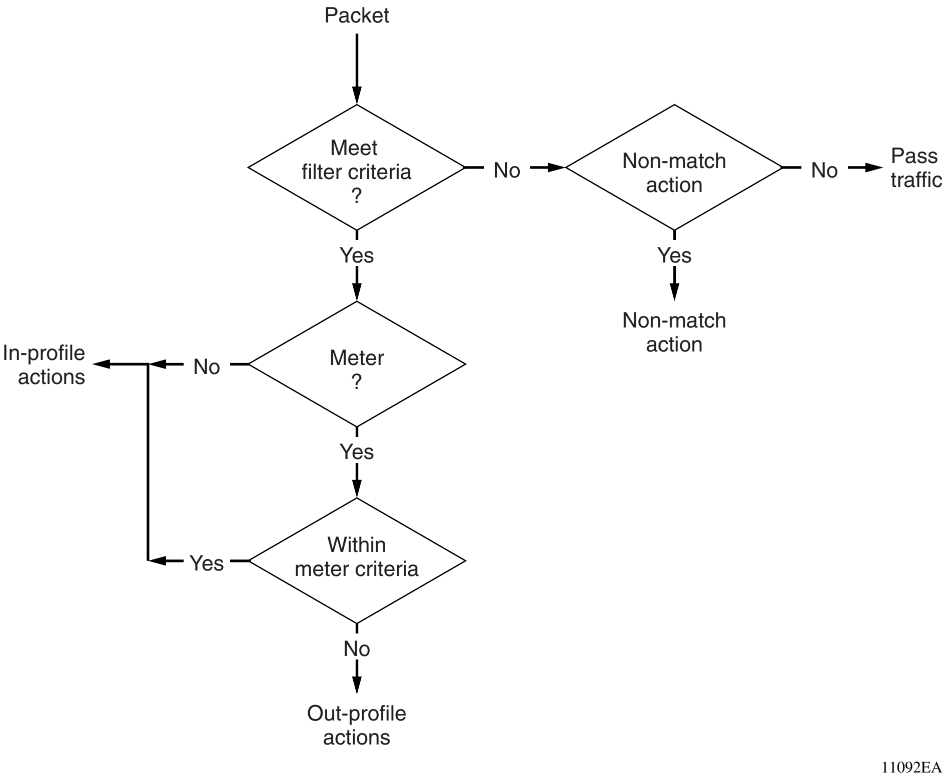

Table 66 shows a summary of the allowable actions for different matching criteria.

Table 66 Summary of allowable actions

Actions	In-Profile	Out-Of-Profile	Non-Matching
Drop/transmit	X	X	X
Update DSCP	X	X	X
Update 802.1p user priority	X		X
Set drop precedence	X	X	X

The Ethernet Routing Switch 3510-24T filters collectively direct the system to initiate the following actions on a packet, depending on your configuration:

- Drop
- Re-mark the packet
  - Re-mark a new DiffServ Codepoint (DSCP)
  - Re-mark the 802.1p field
  - Assign a drop precedence

Note: The 802.1p user priority value used for out-of-profile packets is derived from the associated in-profile action to prevent reordering at egress of packets from a single flow.

Packets received on an interface are matched against all policies associated with that interface. So, potentially, any number of policies—from none to many—are applied to the packet, depending on the policies associated with the specific interface. The set of actions applied to the packet is a result of the policies associated with that interface, ranging from no actions to many actions.

For example, if one policy associated with the specific interface only specifies a value updating the DSCP value while another policy associated with that same interface only specifies a value for updating the 802.1p user priority value, both of these actions occur. If conflicts among actions are detected—for example, if two policies on the specified interface request that the DSCP be updated but specify different values—the value from the policy with the higher precedence will be used.

The actions applied to packets include those actions defined from user-defined policies and those actions defined from system default policies. The user-defined actions always carry higher precedences than the system default actions. This means that, if user-defined policies do not specify actions that overlap with the actions associated with system default policies (for example, the DSCP and 802.1p update actions installed on untrusted interfaces), the lowest precedence, default policy actions will be included in the set of actions to be applied to the identified traffic.

Specifying Interface Action Extensions

The interface action extensions add to the base set of actions.

Table 67 shows a summary of the allowable interface action extensions for different matching criteria.

Table 67 Summary of allowable interface action extensions

Interface action extensions	**In-Profile**	**Out-Of-Profile**	**Non-Matching**
Set egress unicast port	X		X
Set egress non-unicast port	X		X

The Ethernet Routing Switch 3510-24T filters collectively direct the system to initiate the following interface action extensions on a packet, depending on your configuration:

- Set egress unicast interface—specifies redirection of normally switched known (with a previously learned destination address) unicast packets to a specific interface (port)
- Set egress non-unicast interface—specifies redirection of normally switched non-unicast (that is, broadcast, multicast, and flooding) packets to a specific interface (port)

Specifying Meters

QoS metering, which operates at ingress, provides different levels of service to data streams through user-configurable parameters. A meter is used to measure the traffic stream against a traffic profile, which you create. Thus, creating meters yields In-Profile and Out-of-Profile traffic.

You can associate different meters or match actions with different classifiers across a block of classifiers. You can configure policies without metering, or you can configure policies with a single meter or match action that applies to *all* the classifiers associated with that policy. You cannot define meter and action criteria in both the policy definition and the individual classifier definition.

A policy can be created with a meter that is applied to all classifiers, and a policy can be created that has meters applied to individual classifiers—but you cannot have both types in the same policy or action.

A meter applied to a policy will have that metering criteria applied to each port of the interface group (role combination). In other words, the specified bandwidth is allocated on each port, not distributed across all ports.

Using meters, you set a Committed Rate in Kb/s (1000 bits per second in each Kb/s). All traffic within this Committed Rate is In-Profile. Additionally, you set a Maximum Burst Rate that specifies an allowed data burst larger than the Committed Rate for a brief period. After you set the Maximum Burst Rate, the system helps you choose the Duration for this burst. Combined, these parameters define the In-Profile traffic.

Note: The range for the committed rate is 1000 to < 1023000 Kb/s**.** You set the rate in increments of 1000 Kb/s (1 megabit) each.

An example of traffic policing is limiting traffic entering a port to a specified bandwidth, such as 5000 Kb/s (Committed Rate). Instead of dropping all traffic that exceeds this threshold, you can configure a Maximum Burst Rate to exceed the threshold (Committed Rate), for a brief period of time (Duration), without being dropped.

Note: Burst rate and duration are used to determine burst size.

Note: Meter definitions where the committed burst size is too small based on the requested committed rate will be rejected.
The committed burst size can only be one of the following discrete values (in bytes): 16384 (16K), 20480 (20K), 32768 (32K), 45056 (44K), 77824 (76K), 143360 (140K), 274432 (268K), or 524288 (512K).

Trusted, Untrusted, and Unrestricted Interfaces

Ethernet Routing Switch 3510-24T ports are classified into three categories:

- trusted
- untrusted
- unrestricted

The classifications of trusted, untrusted, and unrestricted actually apply to *groups* of ports (interface groups). These three categories are also referred to as interface classes. In your network, trusted ports are usually connected to the core of the DiffServ network, and untrusted ports are typically access links that are connected to end stations. Unrestricted ports can be either access links or connected to the core network.

At factory default, all ports are considered untrusted (factory default value). However, for those interface groups you create, the default is unrestricted.

Because a port can belong to only one interface group, a port will be classified as trusted, untrusted, or unrestricted. These types are also referred to as interface classes.

Trusted and untrusted ports are automatically associated with policies that initiate default traffic processing. This default processing occurs if:

- no actions are initiated based on user-defined policy criteria that matches the traffic.

  OR
- the actions associated with the user-defined policy do not conflict with the default processing actions.

The default processing of trusted and untrusted interfaces is as follows:

- Trusted interfaces — IPv4 traffic received on trusted interfaces is re-marked at the layer 2 level, that is, the 802.1p user priority value is updated based on the DSCP value in the packet at ingress and the installed DSCP-to-CoS mapping data. The DSCP value is not updated. Remapping will only occur, by default, for standardized DSCP values (for example: EF, AFXX) and any proprietary Nortel values. The DSCP values that are remapped are associated with a non-zero 802.1p user priority value in the DSCP-to-COS Mapping Table.
- Untrusted interfaces—IPv4 traffic received on untrusted interfaces is re-marked at the layer 3 level—that is, the DSCP value is updated. The new DSCP value is determined differently depending on whether the packet is untagged or tagged:
    - Untagged frames

      The DSCP value is derived using the default port priority of the interface receiving the ingressing packet. This default port priority is used to perform a lookup in the installed CoS-to-DSCP mapping table.

      The 802.1p user priority value is unchanged—that is, the default port priority determines this value.

      (Thus the DSCP value on untagged frames on untrusted interfaces is updated using the default port priority of the ingress interface; the user sets the default port priority).
    - Tagged frames

      The DSCP value is re-marked to indicate best effort treatment is all that is required for this traffic.

      The 802.1p user priority value is updated based on the DSCP-to-CoS mapping data associated with the best effort DSCP, which is 0.

Table 68 shows the default guidelines the switch uses to re-mark various fields of IPv4 traffic (and layer 2 traffic matching IPv4) based on the class of the interface. These are the actions that occur if the user does not intervene at all; they are the default actions of the switch.

Table 68 Default with no user action re-marking of QoS fields by class of interface--IPv4

Type of filter	Action	Trusted	Untrusted	Unrestricted
IPv4 filter criteria or Layer 2 filter criteria matching IPv4	DSCP	Does not change	• Tagged—Updates to 0 (Standard) • Untagged—Updates using mapping table and port's default value	Does not change
	IEEE 802.1p	Updates based on DSCP mapping table value	Updates based on DSCP mapping table value	Does not change

Note: The default for layer 2 non-IP traffic is to pass the traffic through all interfaces classes with the QoS values for 802.1p and drop precedence unchanged.

The Ethernet Routing Switch 3510-24T does not trust the DSCP of IPv4 traffic received from an untrusted port, but it does trust the DSCP of IPv4 traffic received from a trusted port.

L2 non-IP traffic received on either a trusted port or an untrusted port traverses the switch with no change.

IPv4 traffic received on a trusted port has the 802.1p user priority value re-marked and the drop precedence set based on the DSCP in the received IP packet.

If an IPv4 packet is received from a trusted port and either it does not match any of the classifier elements installed by the user on this port or it does match a classifier element but is not dropped, the Ethernet Routing Switch 3510-24T uses default system classifiers to change the packet IEEE 802.1p and drop precedence based on the DSCP of the packet.

If an IPv4 packet is received from an untrusted port and it does not match any one of the classifier elements installed by the user on the port, the Ethernet Routing Switch 3510-24T uses default system classifiers to change the packet DSCP, IEEE 802.1p priority, and drop precedence as follows:

- If the packet is tagged, the 802.1p user priority value is derived from the DSCP-to-CoS mapping table using the best effort DSCP, which is 0.
- If an IPv4 packet is untagged, the Ethernet Routing Switch 3510-24T uses the default classifier to change the DSCP based on the default IEEE 802.1p priority of the ingress untrusted port to index into the DSCP-to-CoS mapping table to determine the DSCP value

Table 69 describes the default DSCP, QoS class, IEEE 802.1p, and egress queue assignment for packets in each traffic class.

Table 69 Default mapping of DSCP to QoS class and IEEE 802.1p

Incoming or re-marked DSCP (hex values)	QoS class	Number of queues (8)	Outgoing IEEE 802.1p user priority
CS7 (0x38)	Critical	1	7
CS6 (0x30)	Network	1	
EF(0x2E), CS5(0x28)	Premium	2	6
AF41(0x22), AF42(0x24), AF43(0x26), CS4(0x20)	Platinum	3	5
AF31(0x1A), AF32(0x1C), AF33(0x1E), CS3(0x18)	Gold	4	4
AF21(0x12), AF22(0x14), AF23(0x16), CS2(0x10)	Silver	5	3
AF11(0xA), AF12(0xC), AF13(0xE), CS1(0x8)	Bronze	6	2
DE(0x0), CS0(0x0), all undefined DSCPs	Standard	7	0

As displayed in Table 69, the traffic service class determines the IEEE 802.1p priority that determines the egress queue of the traffic. Non-IP traffic can be in the same IP service class if the non-IP packets are assigned the same IEEE 802.1p priority.

Specifying Policies

When network traffic attributes match those specified in a traffic pattern, the policy instructs the network device to perform a specified action on each packet that passes through it. A policy is a set of rules and actions that are applied to specific ports.

Among policies, the policy with the highest precedence is evaluated first, then the policy with the next-lowest precedence and so on. For example, with a precedence of 1 to 8, the system begins the evaluation with 8, moves onto 7, and so forth. This is important to remember when you configure policies. The valid precedence range is 1 - 10 if DHCP relay feature is enabled, or 1 - 11 if DHCP relay feature is disabled.A policy may reference an individual classifier or a classifier block.

A *policy* is a network traffic-controlling mechanism that monitors the characteristics of the traffic (for example, its source, destination, and protocol) and performs a controlling action on the traffic when certain user-defined characteristics are matched. A *policy action* is the effect a policy has on network traffic that matches the traffic profile of the policy.

The policies tie together:

- Actions
- Meters
- Classifier elements/classifiers/classifier blocks
- Interface groups

The policies, by connecting these user-defined configurations, control the traffic on the switch.

Note: You must configure interface groups (role combinations), classification criteria, actions and meters before you attempt to reference that data in a policy.

Ports are assigned to interface groups that are linked to policies. Although a single policy can reference only one interface group, you can configure several policies that reference the same interface group. The policies determine the traffic treatment of the flows.

Note: You can enable or disable policies; you do not have to delete a policy to disable it.
However, you must delete a policy to modify it.To modify a policy, you delete the current policy and create a new policy that is modified as you want. (A policy is automatically enabled when you create it.)

You can also track statistics for QoS. The Ethernet Routing Switch 3510-24T supports per policy and per policy/classifier/interface statistics tracking.

Packet Flow Using QoS

Using DiffServ and QoS, you can designate a specific performance level for packets. This system allows you to prioritize network traffic. However, it requires some thought to configure the prioritizations.You can specify a number of policies, and each policy can match one or many flows—supporting complex classification scenarios.

This section contains a very simplified introduction to the many ways to prioritize packets using QoS. In simple terms, the methods of prioritizing packets depend on the DSCP and the 802.1 priority level and drop precedence.

The QoS class basically directs which group of packets receives the best network throughput, which group of packets receives the next best throughput, and so on. The level of service for each packet is determined by the configurable DSCP.

The available levels of QoS classes are currently named Network, Premium, Platinum, Gold, Silver, Bronze, and Standard. The level of service for each packet is determined by the configurable DSCP.

Classifier elements, classifiers, and classifier blocks basically sort the packets by various configurable parameters. These parameters include VLAN IDs, IP source and subnet address, IP protocol, and many others.

The classifiers/classifier blocks are associated with policies, and policies are organized into a hierarchy. The policy with the highest precedence is evaluated first. The classifier elements, classifiers, and classifier blocks are associated with interface groups, in that packets from a specific port will have the same classification parameters as all others in the particular interface group (role combination).

Meters, operating at ingress, keep the sorted packets within certain parameters. You configure a committed rate of traffic, allowing a certain size for a temporary burst, as In-Profile traffic. All other traffic is configured as Out-of-Profile traffic. If you choose not to meter the flow, you do not configure meters.

Actions determine how the traffic is treated.

The overall total of all the interacting QoS factors on a group of packets is a policy. You configure policies that monitor the characteristics of the traffic and perform a controlling action on the traffic when certain user-defined characteristics are matched.

Figure 99 provides a schematic overview of QoS policies.

Figure 99 Schematic of QoS policy

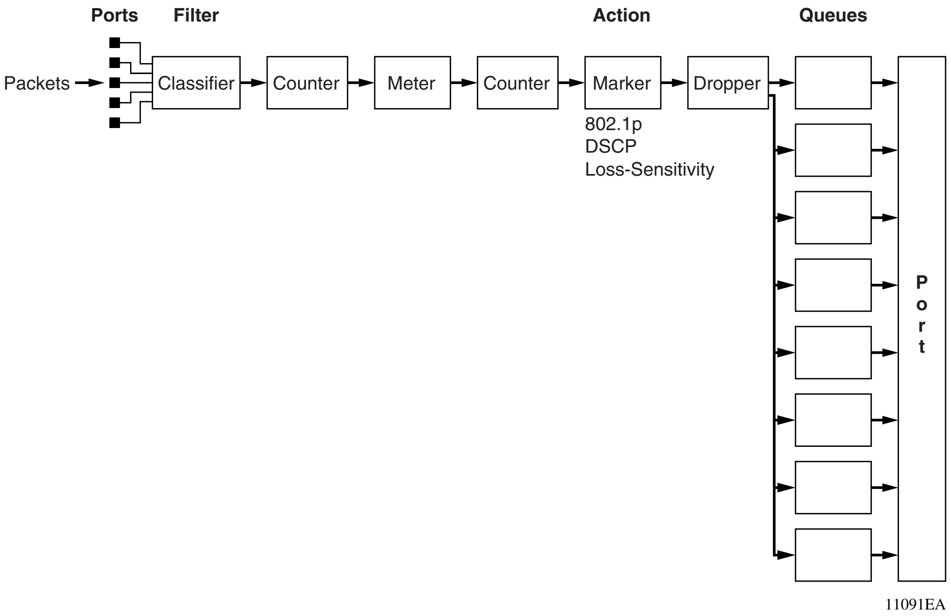

Queue Sets

A QoS queue set is used to logically represent the queuing capabilities that are associated with an egress QoS interface. A queue set is comprised of a number of related queuing components that dictate the queuing behavior supported by the set itself. These include:

- Queue count - the number of different CoS queues in the set.
- Queue service discipline - indicates the means through which queues (competing for limited transmission bandwidth) and the packets held in the queues are scheduled for transmission.
- Queue bandwidth allocation - indicates the absolute or relative amount of bandwidth that can be consumed by the queues in the set. When queues are serviced using a Weighted Round Robin (WRR) or Weighted Fair Queuing (WFQ) discipline, these values represent the weights associated with the queues.
- Queue service order - when multiple service disciplines are in use (e.g., strict priority and WRR), the service order indicates service precedence assigned to individual queues (strict priority) or clusters of queues (WRR).
- Queue size - indicates the maximum buffering resources that can be consumed by the individual queue.

Each QoS egress port has eight queue sets consisting of anywhere from 1 to 8 queues, depending on the queue set you assign to the QoS interfaces. Packets are assigned to a queue based on the IEEE 802.1p, or Class of Service (CoS), value associated with that packet. Depending on the queue set you configure, some queues are serviced in an absolute priority fashion and some queues may be serviced in a Weighted Round Robin (WRR) fashion.

Beginning with software version 4.0, you can configure the queue set, and hence the number of queues per QoS interface, the buffer allocation of the queue set, and the CoS-to-queue priority for each queue within the queue set.

Note: You configure these parameters for all QoS egress interfaces, not on a port-by-port basis. Thus, the egress queuing and buffering characteristics and the CoS-to-queue priorities are the same across all QoS ports. The Ethernet Routing Switch 3510-24T has factory default queue set and buffer allocation mode values. When your reset your system to defaults, these are the values the system has:

- factory default queue set: queue set 8

- buffer allocation mode: regular

Modifying Queue Set Characteristics

You can configure the following characteristics of the queue sets:

- The number of queues per egress QoS interface, their service discipline and relative weights - you select one of the 8 available pre-defined queue sets with the appropriate queue count, service discipline and weights for your specific application.
- The buffering resources consumed by the egress QoS interface - you select regular, large, or maximum to allocate the resources.

You cannot configure other queue characteristics (such as the service discipline or queue weights for WRR scheduler).

The user-configurable parameters for the queue sets take effect only after the next system reset. These configuration parameters are saved in NVRAM.

Although the CoS-to-queue assignments may be changed for all defined queue sets, only the assignments associated with the queue set currently in use affect the traffic processing.

The queues within a queue set are referred to as CoS queues because you map each queue within the queue set to a CoS priority value (refer to "Modifying CoS-to-queue priorities" on page 22). The 8 predefined queue sets contain a varying number of CoS queues, service disciplines, and queue weights. (The relative interface bandwidth consumption percentages for WRR queues are shown as percentages.)

To configure the queue set, choose one of the following 8 available queue sets, which will apply to all QoS egress interfaces, along with their characteristics:

- Queue set 1
  - default queue set
  - 8 CoS queues
  - 1 queue strict priority; 7 WRR queues
    - 7 WRR queues scheduled as 41%, 19%, 13%, 11%, 8%, 5%, and 3%
- Queue set 2
  - 7 CoS queues
  - 1 queue strict priority; 6 WRR queues
    - 6 WRR queues scheduled as 45%, 21%, 15%, 10%, 6%, and 3%
- Queue set 3
  - 6 CoS queues
  - 1 queue strict priority; 5 WRR queues
    - 5 WRR queues scheduled as 52%, 24%, 14%, 7%, and 3%
- Queue set 4
  - 5 CoS queues
  - 1 queue strict priority; 4 WRR queues
    - 4 WRR queues scheduled as 58%, 27%, 11%, and 4%
- Queue set 5
  - 4 CoS queues
  - 1 queue strict priority; 3 WRR queues
    - 3 WRR queues scheduled as 65%, 26%, and 9%
- Queue set 6
  - 3 CoS queues
  - 1 queue strict priority; 2 WRR queues
    - 2 WRR queues scheduled as 75% and 25%
- Queue set 7
  - 2 CoS queues
  - 2 strict priority queues
- Queue set 8

- 1 CoS queue
- 1 strict priority queue

➡ **Note:** Changes affecting the egress interface queue set do not take effect until you reset the system. However, if you query the default queue configuration after configuring a new queue set and prior to resetting the system, the system returns the newly configured (not yet effective) queue set.

You can also configure the buffer allocation (consumption) level for the configured queue set. You choose one from among regular, large, or maximum allocations.

➡ **Note:** You must reset the system for the modified buffer resource allocation to take effect. However, if you query the buffer resource after modifying the buffer resource allocation and prior to resetting the system, the system returns the newly configured (not yet effective) buffer resource.

You must use SNMP or the NCLI in order to modify these parameters.

Modifying CoS-to-queue Priorities

You can also modify the association of 802.1p, or CoS, values to each queue within the queue set. Within a given queue set, you can assign a value of 0 to 7 to each queue in that set.

➡ **Note:** Any modification to the CoS-to-queue values takes effect immediately; you do not have to reset the system to modify these values.

You can modify these values using SNMP, the NCLI, or the Web-based management system.

QoS Configuration Guidelines

You can install classifiers that will act on traffic destined for the switch itself, such as ICMP Echo Requests (ping) and SNMP messages. If the associated action is to drop the traffic, you can lock yourself out of the switch.

Using QoS on the Ethernet Routing Switch 3510-24T has the following limitations:

- You can configure up to 11 policies per interface group.
- You can configure up to 63 meters per interface (port).
- You can configure 100 filter components per interface (port).
- When you enable tracking statistics for the policies, the switch uses one counter for each classifier for each interface (port) of the policy or a counter for each policy. You can assign up to 32 counters to an interface (port).

Troubleshooting Tips

If you encounter problems configuring the queue sets, ensure that the modified queue set is associated with the QoS interfaces. It is important to note that the device must be reset for the changes to take effect.

Sometimes after modifying the default buffering resources, you do not see the queue sizes in the queue set updated. Again, the device must be reset for the changes to take effect.

Finally, modified CoS-to-queue assignments affect only the active queue; this may be why you do not see an effect immediately after modifying the values.

Chapter 5
Sample QoS Configuration

You can configure QoS using the Web-based management system, Nortel Command Line Interface (NCLI), SNMP, or Device Manager (DM). This section presents a sample QoS configuration using the Web-based management system using the QoS pages.

For more information on configuring QoS with the Web-based management system, refer to *Web Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3*. For information on configuring QoS with other management systems, refer to *NCLI Configuration Guide for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3* and *Switch Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3*.

It is important that you refer to *Web Management for Nortel Ethernet Routing Switch 3510-24T, Software Release 4.0.3* for details to access the Web-based management interface, directory and page navigation information, and field descriptions

Note: Nortel recommends that you configure classifier and interface parameters in the order in which the screens are presented in this example.

This chapter provides a sample configuration using the Web-based management interface QoS pages. You must define classifier elements before you define classifiers and Classifier blocks; you must define the interface action extensions before you define the actions; and you must define actions before you define the meters. The policy must be defined last, after the other parameters are configured. This chapter covers the following topics, using the QoS pages:

- "Creating Interface Groups", next
- "Accepting Default Mapping Values" on page 278
- "Setting up Classifiers and Classifier Blocks" on page 278

- "Defining an IP Classifier Element" on page 279
- "Defining a Layer 2 Classifier Element" on page 281
- "Creating Classifiers" on page 283
- "Creating Classifier Blocks" on page 285
- "Configuring Actions" on page 287
- "Configuring Interface Action Extensions" on page 289
- "Configuring Meters" on page 291
- "Configuring Policies" on page 293
- "Assigning Mapping Values" on page 296

➡ **Note:** You cannot modify the following configured items:
- interfaces
- classifier elements
- classifiers
- interface action extensions
- meters
- policies

You must first delete the current item and then enter a new one with the modifications.

Creating Interface Groups

To create an interface group:

1 In the Web-based management interface, click the Application > QoS menu option.

The QoS menu option expands to display:

- Devices
- Rules
- Actions
- Interface Actions Ext
- Meter
- Policy

- Agent

2 Click Devices.

The Devices menu option expands to display:

- Interface Config
- Priority Q Assign
- Priority Mapping
- DSCP Mapping

3 Click Interface Config.

The Interface Configuration page opens (Figure 100).

Figure 100 Interface Configuration page

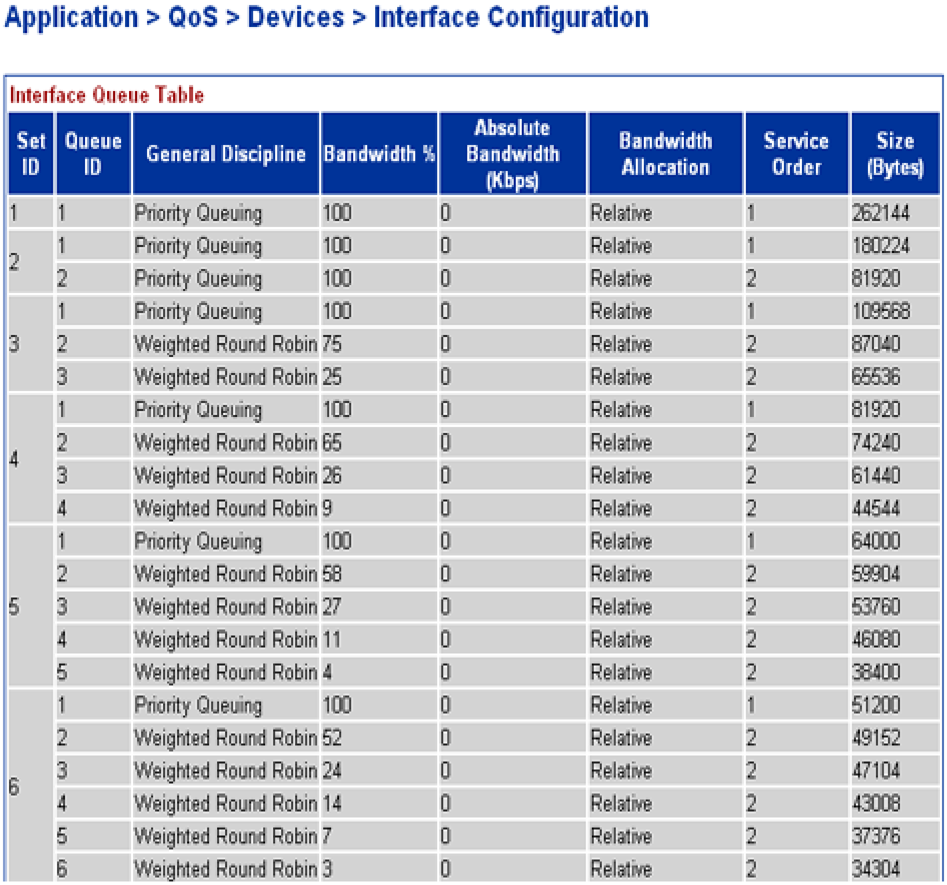

Application > QoS > Devices > Interface Configuration

Interface Queue Table

Set ID	Queue ID	General Discipline	Bandwidth %	Absolute Bandwidth (Kbps)	Bandwidth Allocation	Service Order	Size (Bytes)
1	1	Priority Queuing	100	0	Relative	1	262144
2	1	Priority Queuing	100	0	Relative	1	180224
	2	Priority Queuing	100	0	Relative	2	81920
3	1	Priority Queuing	100	0	Relative	1	109568
	2	Weighted Round Robin	75	0	Relative	2	87040
	3	Weighted Round Robin	25	0	Relative	2	65536
4	1	Priority Queuing	100	0	Relative	1	81920
	2	Weighted Round Robin	65	0	Relative	2	74240
	3	Weighted Round Robin	26	0	Relative	2	61440
	4	Weighted Round Robin	9	0	Relative	2	44544
5	1	Priority Queuing	100	0	Relative	1	64000
	2	Weighted Round Robin	58	0	Relative	2	59904
	3	Weighted Round Robin	27	0	Relative	2	53760
	4	Weighted Round Robin	11	0	Relative	2	46080
	5	Weighted Round Robin	4	0	Relative	2	38400
6	1	Priority Queuing	100	0	Relative	1	51200
	2	Weighted Round Robin	52	0	Relative	2	49152
	3	Weighted Round Robin	24	0	Relative	2	47104
	4	Weighted Round Robin	14	0	Relative	2	43008
	5	Weighted Round Robin	7	0	Relative	2	37376
	6	Weighted Round Robin	3	0	Relative	2	34304

The Interface Group Creation section of this page allows you to define groups of interfaces. You can view your interface configurations in the read-only Interface Queue Table and the Interface Group Table.

4 Use the Interface Group Creation section to create a new Role Combination. In the Role Combination field, enter `Webbrowsing`. (Remember, this is an example. You can enter any string in this field.)

Note: Do not use spaces in the naming field.

5 In the Interface Class field, choose **`unrestricted`**.

6 Click Submit.

The new entry appears in the Interface Group Table.

7 To assign interfaces to this role combination, click the modify icon in the left column of the Action field next to row of the role combination you just created.

The Interface Group Assignment page opens (Figure 101).

Figure 101 Interface Group Assignment page

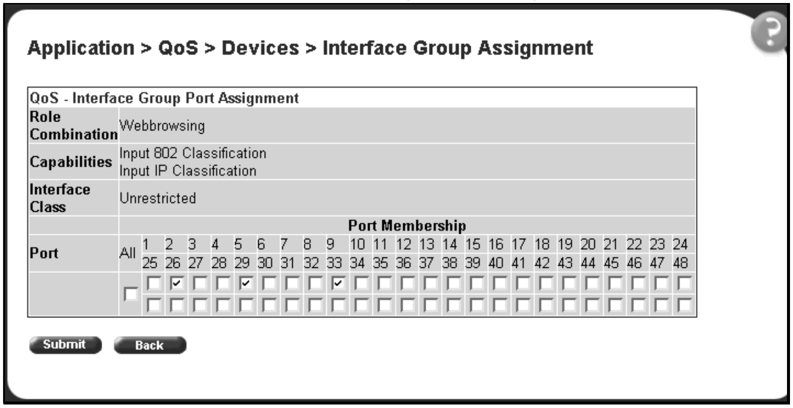

QoS - Interface Group Port Assignment

Role Combination	Webbrowsing
Capabilities	Input 802 Classification Input IP Classification
Interface Class	Unrestricted

Port Membership

Port	All	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
		25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48
	☐	☐	☑	☐	☐	☑	☐	☐	☐	☑	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐
		☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐

The Interface Group Assignment page displays the name of the interface group (role combination), the capabilities, and the interface class (or type of interface) in the group.

a Click the ports you want to add to the specified interface group, or click All to add all ports on the unit.

Note: Adding a port to an interface groups automatically deletes that port from any prior interface group assignment.

b Click Submit.

Note: If you delete a role combination, you must remove all ports in the Interface Group Assignment page first. A role combination cannot be deleted if it is referenced by an installed policy or interface.

Accepting Default Mapping Values

If you choose to accept the default values for IEEE 802.1p priority and DSCP values, skip this section and procede to "Setting up Classifiers and Classifier Blocks".

Note: Nortel recommends that you use the default mapping values to ensure end-to-end QoS connectivity across Nortel products.

To manually configure mapping values, refer to "Assigning Mapping Values" on page 296.

Setting up Classifiers and Classifier Blocks

You use classifier elements to create classifiers. You combine classifier elements into classifiers only in *one* of the following ways:

- one IP classifier element plus one L2 classifier element
- one IP classifier element
- one L2 classifier element

Classifiers allow you to classify packets by various parameters. (For more information on these parameters, refer to Chapter 4.) Classifiers are combined into classifier blocks. Classifiers or classifier blocks are then associated with an interface group via a policy.

You configure classifier elements. The QoS > Rules > IP Classifier Element page or the QoS > Rules > Layer 2 Classifier Element page allows you to enter matching conditions for an individual classifier element. You set up special conditions for packet processing. In order for packets to be processed, a packet has to match all the fields you specify.

Note: When you choose the value Ignore, the system matches all fields for that parameter.

Defining an IP Classifier Element

You create IP classifier elements for IP packets that are to be forwarded through the Ethernet Routing Switch 3510-24T on specific ingress ports. In each IP packet, there is a differentiated services (DiffServ) field in the packet header that you can mark for specific treatment. This field is called the DiffServ code point (DSCP). The DSCP has a specific value that determines how the packet is treated as it travels through the network. As each packet is examined it will be forwarded or dropped, depending on whether or not the classifier criteria is matched.

You use the IP Classifier Element Creation section of the Rules > IP Classifier Element page when defining your IP classifier elements.

To define an IP classifier element:

1. Click the Application > QoS > Rules > IP Classifier Element menu option.

   The IP Classifier Element page opens (Figure 102).

Figure 102 IP Classifier Element page

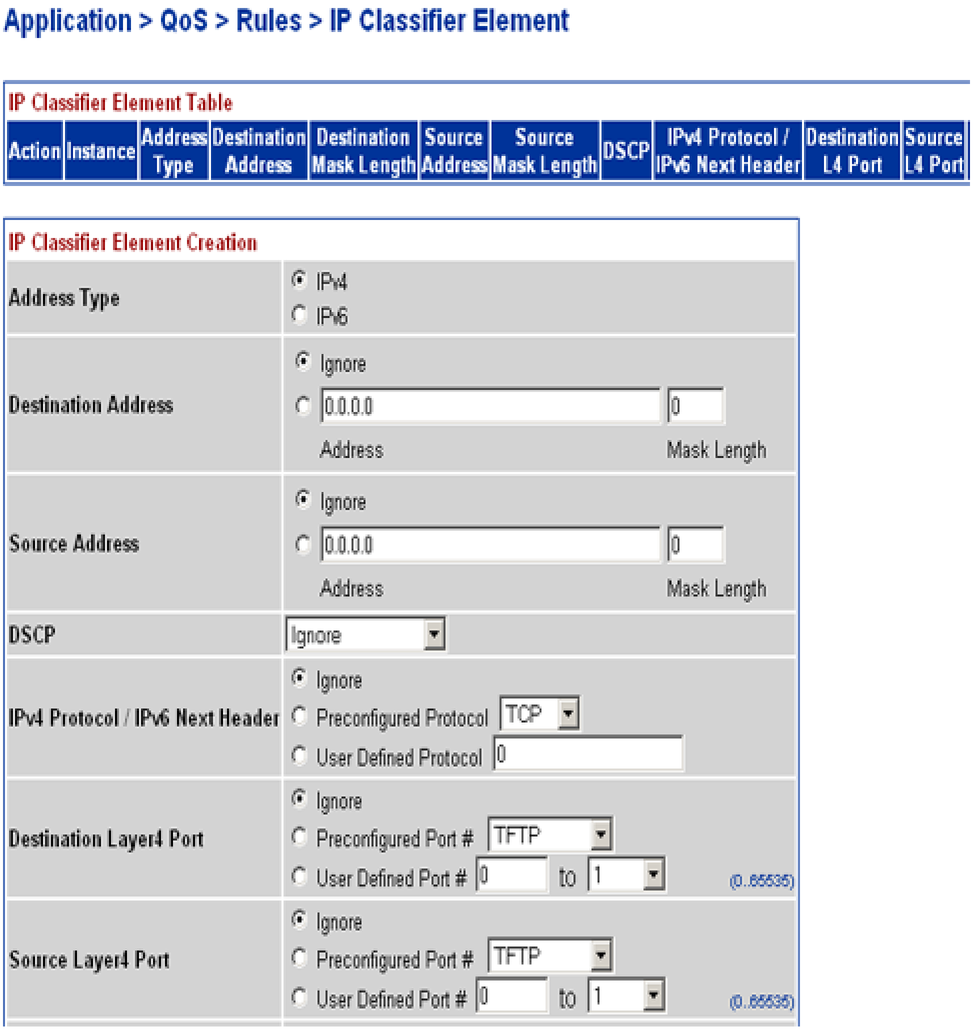

Application > QoS > Rules > IP Classifier Element

IP Classifier Element Table

Action	Instance	Address Type	Destination Address	Destination Mask Length	Source Address	Source Mask Length	DSCP	IPv4 Protocol / IPv6 Next Header	Destination L4 Port	Source L4 Port

IP Classifier Element Creation

Address Type	● IPv4 ○ IPv6
Destination Address	● Ignore ○ 0.0.0.0 Address; 0 Mask Length
Source Address	● Ignore ○ 0.0.0.0 Address; 0 Mask Length
DSCP	Ignore
IPv4 Protocol / IPv6 Next Header	● Ignore ○ Preconfigured Protocol TCP ○ User Defined Protocol 0
Destination Layer4 Port	● Ignore ○ Preconfigured Port # TFTP ○ User Defined Port # 0 to 1 (0..65535)
Source Layer4 Port	● Ignore ○ Preconfigured Port # TFTP ○ User Defined Port # 0 to 1 (0..65535)

2 In the Address Type box, click **`IPv4`**.

3 In the Destination Address field, in the Address field enter **`134.177.69.122.`**

This address is used to match the destination IP address in the packet's IP header.

4 In the Destination Address field, in the Mask Length, enter **`32`**.

5. In the Source Address field, in the Address field enter `134.177.69.173`.

   This is the IP address to match against the packet's source IP address.

6. In the Source Address field, in the Mask Length, enter `32`.
7. In the DSCP field, choose `32, 0x20,100000` from the list.

   This value matches packets with a DSCP of 0x20 (32 decimal value).

   If you choose Ignore, the DSCP value in the packet is ignored.

8. In the IPv4 Protocol/IPv6 Next Header field, choose `TCP` from the list.

   When you select TCP, you specify that only TCP packets be matched. If you select Ignore, all IP protocols are matched.

9. In the Destination Layer 4 Port field, click `Ignore`.
10. In the Source Layer 4 Port field, click `Ignore`.
11. In the IPv6 Flow Id field, click `Ignore`.
12. Click Submit.

    The new entry appears in the IP Classifier Element Table.

Defining a Layer 2 Classifier Element

You configure layer 2 classifier elements by defining IEEE 802-based parameters.

To configure a layer 2 classifier element:

1. Click the Application > QoS > Rules > Layer 2 Classifier Element menu option.

   The Layer2 Classifier Element page opens (Figure 103).

Figure 103 Layer 2 Classifier Element page

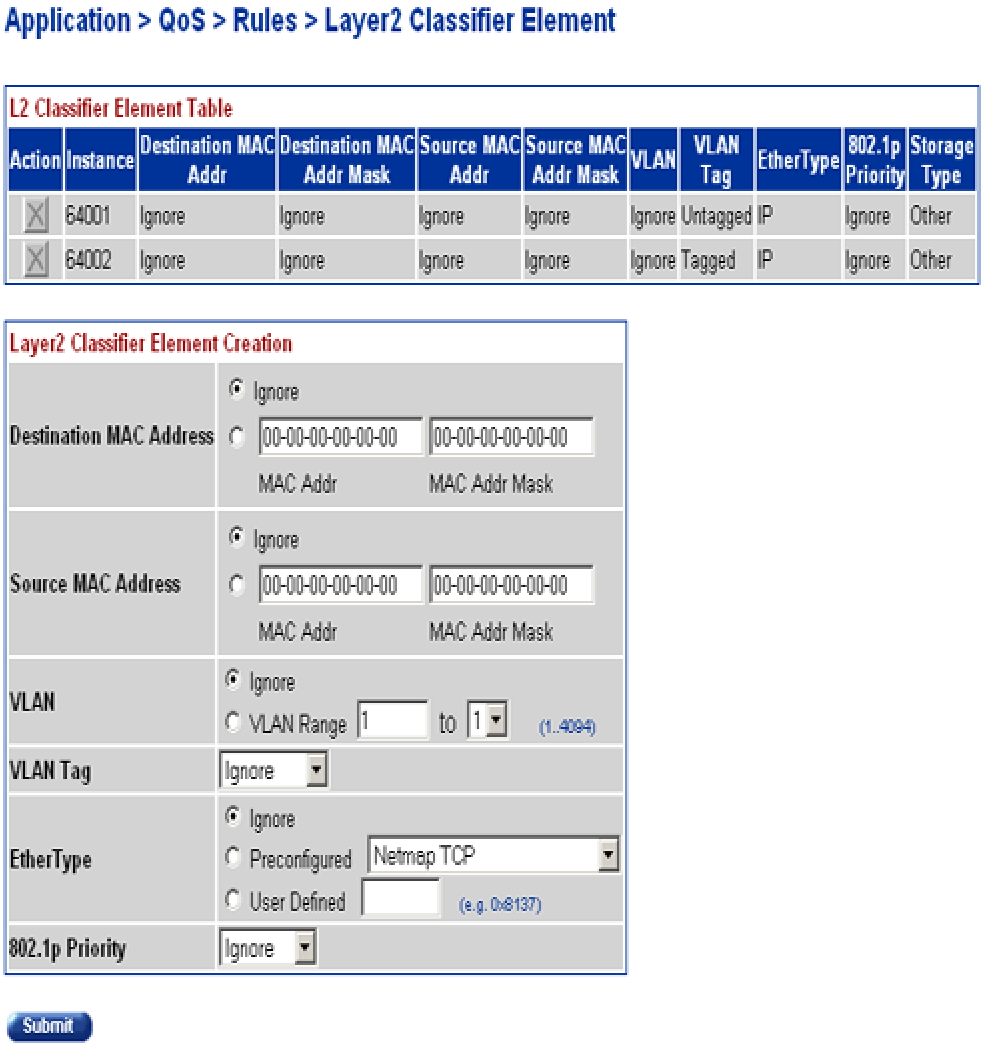

Application > QoS > Rules > Layer2 Classifier Element

L2 Classifier Element Table

Action	Instance	Destination MAC Addr	Destination MAC Addr Mask	Source MAC Addr	Source MAC Addr Mask	VLAN	VLAN Tag	EtherType	802.1p Priority	Storage Type
☒	64001	Ignore	Ignore	Ignore	Ignore	Ignore	Untagged	IP	Ignore	Other
☒	64002	Ignore	Ignore	Ignore	Ignore	Ignore	Tagged	IP	Ignore	Other

Layer2 Classifier Element Creation

Destination MAC Address	● Ignore ○ 00-00-00-00-00-00 (MAC Addr) 00-00-00-00-00-00 (MAC Addr Mask)
Source MAC Address	● Ignore ○ 00-00-00-00-00-00 (MAC Addr) 00-00-00-00-00-00 (MAC Addr Mask)
VLAN	● Ignore ○ VLAN Range 1 to 1 (1..4094)
VLAN Tag	Ignore
EtherType	● Ignore ○ Preconfigured Netmap TCP ○ User Defined (e.g. 0x8137)
802.1p Priority	Ignore

2 In the Destination MAC Address field, click **`Ignore`**.

3 In the Source MAC Address field, click **`Ignore`**.

4 In the VLAN field, click VLAN Range and leave at **1** for both.

5 In the VLAN Tag field, click **`Ignore`**.

6 In the EtherType field, click **`Preconfigured`**, and select **IP**.

Sets the EtherTypes for IP.

➡ **Note:** On the Layer 2 Classifier Element page, if you select any EtherType except IP4 or IPv6, you are creating a filter for non-IP packets.

---

7 In the 802.1p Priority field, choose **Ignore**.

8 Click Submit.

The new entry is displayed in the Layer2 Classifier Element Table.

Creating Classifiers

To create a classifier:

1 From the main menu, choose Application > QoS > Rules > Classifier.

The Classifier page opens (Figure 104).

Figure 104 Classifier page

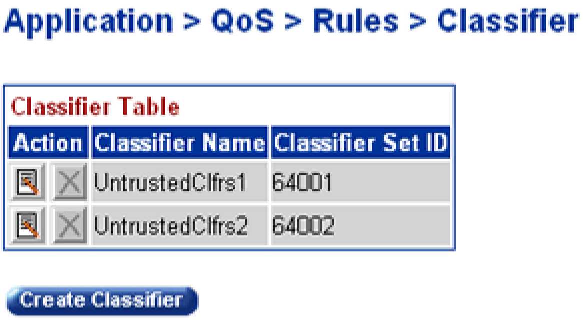

Application > QoS > Rules > Classifier

Classifier Table

Action	Classifier Name	Classifier Set ID
	UntrustedClfrs1	64001
	UntrustedClfrs2	64002

2 Click Create Classifier.

The Classifier Creation page opens (Figure 105).

Figure 105 Classifier Creation page

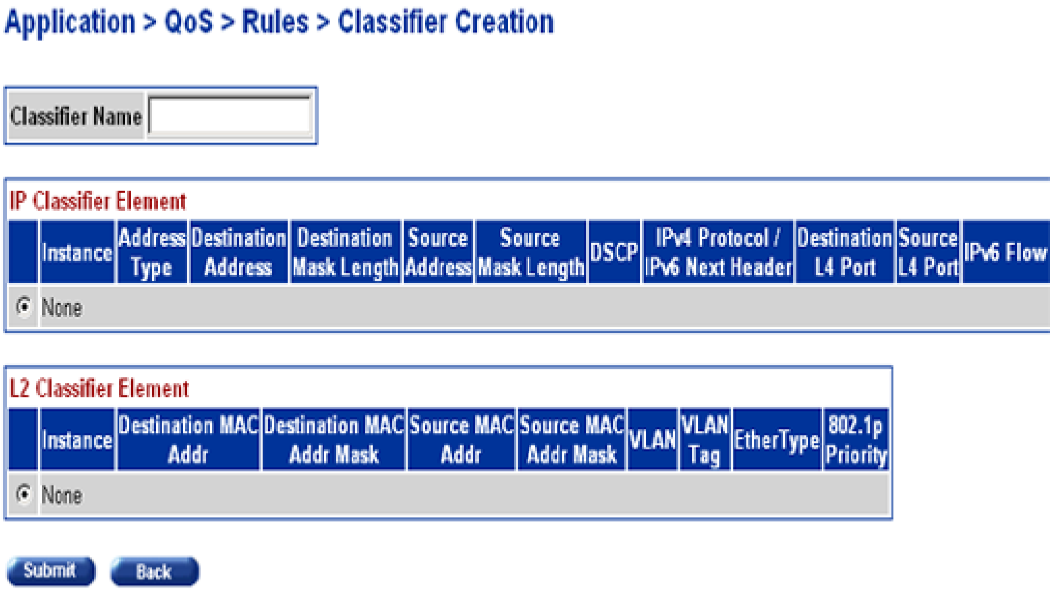

Application > QoS > Rules > Classifier Creation

Classifier Name	

IP Classifier Element

	Instance	Address Type	Destination Address	Destination Mask Length	Source Address	Source Mask Length	DSCP	IPv4 Protocol / IPv6 Next Header	Destination L4 Port	Source L4 Port	IPv6 Flow
●	None										

L2 Classifier Element

	Instance	Destination MAC Addr	Destination MAC Addr Mask	Source MAC Addr	Source MAC Addr Mask	VLAN	VLAN Tag	EtherType	802.1p Priority
●	None								

Note: Each classifier can have only a *single* IP classifier element *plus* a *single* L2 classifier element. You can, however, create a classifier using only one IP classifier element or only one L2 classifier element.
Each classifier in a classifier block must match the same parameters and the same mask, range, and packet format type. Additionally, all members of a classifier block must be configured consistently regarding meters and actions—that is, they must all specify meters or they must all not specify meters, and they must all specify actions or they must all not specify actions.

---

3 Click only one IP classifier element from those displayed on the IP Classifier Element table.

4 Click only one L2 classifier from those displayed on the LS Classifier Element table.

5 Click Submit.

The system returns you to the Classifier page and displays the new classifier on the Classifier Table (Figure 106).

Figure 106 Classifier page with new classifier

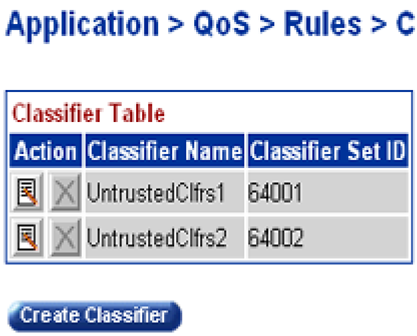

Classifier Table

Action	Classifier Name	Classifier Set ID
	UntrustedClfrs1	64001
	UntrustedClfrs2	64002

Note: You cannot modify a classifier. You must delete the one you wish to modify and create a new classifier as you want it. If you do not assign a name to the classifier, the system assigns one for you.

Creating Classifier Blocks

To create a classifier block:

1 From the main menu, choose Application > QoS > Rules > Classifier Block.

The Classifier Block page opens (Figure 107).

Figure 107 Classifier Block page

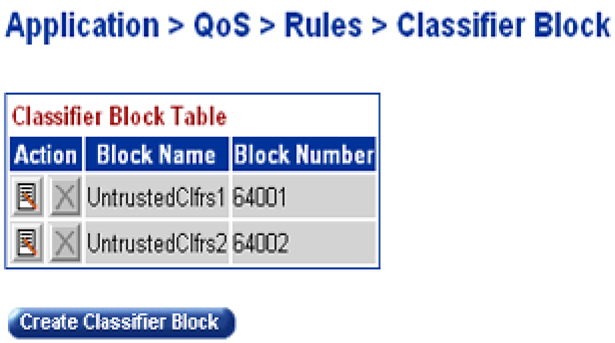

Application > QoS > Rules > Classifier Block

Classifier Block Table

Action	Block Name	Block Number
	UntrustedClfrs1	64001
	UntrustedClfrs2	64002

2 Click Create Classifier Block.

The Classifier Block Creation/Modification page opens (Figure 108).

Figure 108 Classifier Block Creation/Modification page

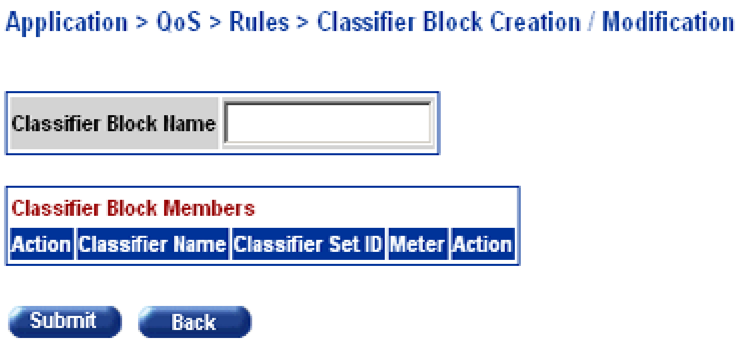

Application > QoS > Rules > Classifier Block Creation / Modification

Classifier Block Name	

Classifier Block Members

Action	Classifier Name	Classifier Set ID	Meter	Action

3 In the Classifier Block Name field, enter `Test`.

4 In the Classifier Block Members table, in the Action field, click next to the clfrComp1 entry.

5 Click Submit.

The system returns you to the Classifier Block page with your new entry displayed in the table.

➡ **Note:** Each classifier can have only a *single* IP classifier element *plus* a *single* L2 classifier element. You can, however, create a classifier using only one IP classifier element or only one L2 classifier element. Each classifier in a classifier block must match the same parameters and the same mask, range, and packet format type. Additionally, all members of a classifier block must be configured consistently regarding meters and actions—that is, they must all specify meters or they must all not specify meters, and they must all specify actions or they must all not specify actions.

Configuring Actions

When you assign actions to a policy or to members of a classifier block, you specify the type of behavior that you want a policy to apply to a flow of packets.You specify actions by associating them with a policy as a whole or by associating actions with each member of a classifier block. An action associated with a policy is applied to all classifier associated with the policy. An action associated with an individual classifier block member is only applied to traffic that specific classification criteria. When the specific classifier elements match incoming packets, the actions are performed on these packets.

To configure an action:

1. Click the Application > QoS > Actions menu option.

   The Actions page opens (Figure 109).

Figure 109 Actions page

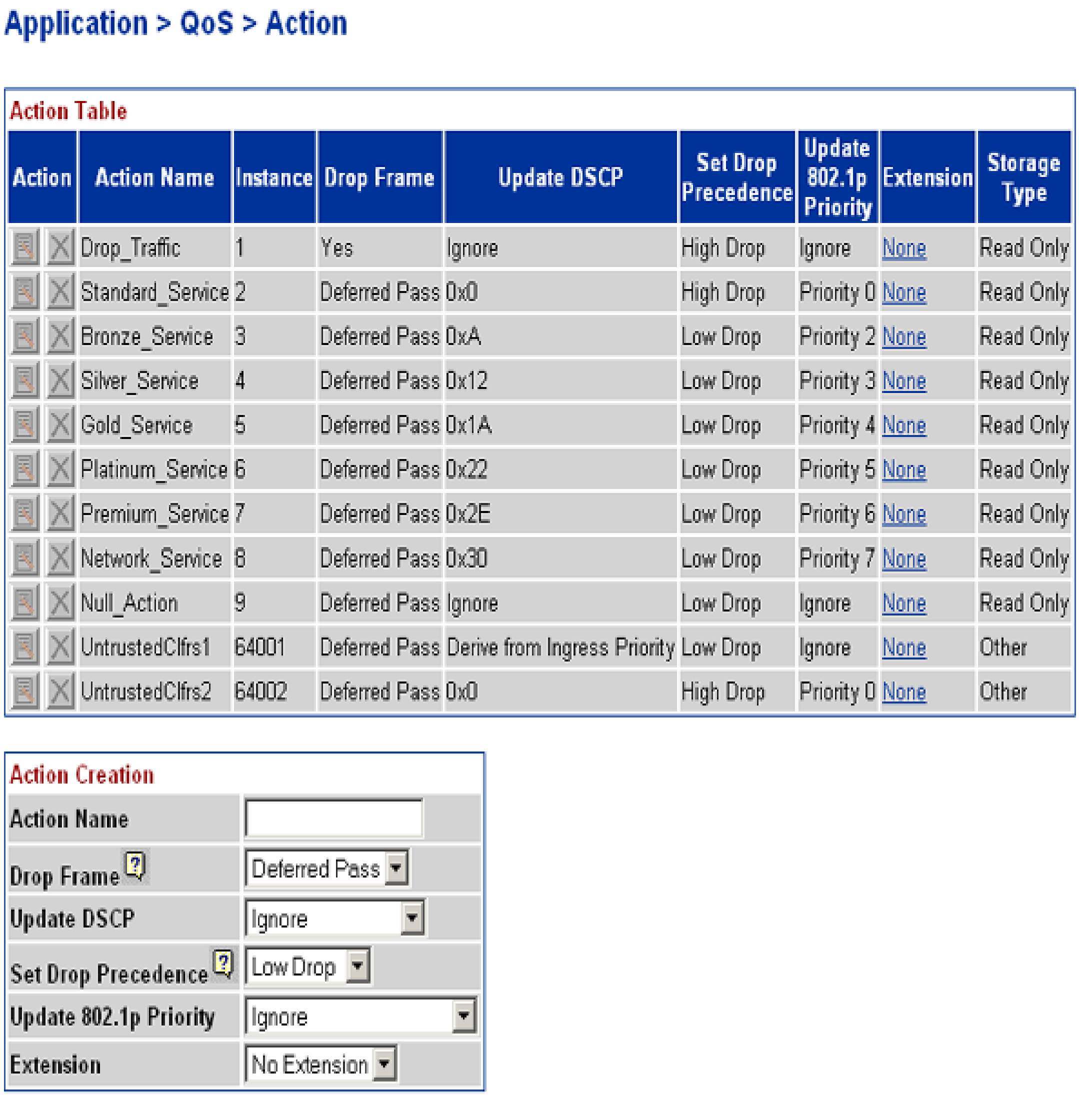

Application > QoS > Action

Action Table

Action	Action Name	Instance	Drop Frame	Update DSCP	Set Drop Precedence	Update 802.1p Priority	Extension	Storage Type
	Drop_Traffic	1	Yes	Ignore	High Drop	Ignore	None	Read Only
	Standard_Service	2	Deferred Pass	0x0	High Drop	Priority 0	None	Read Only
	Bronze_Service	3	Deferred Pass	0xA	Low Drop	Priority 2	None	Read Only
	Silver_Service	4	Deferred Pass	0x12	Low Drop	Priority 3	None	Read Only
	Gold_Service	5	Deferred Pass	0x1A	Low Drop	Priority 4	None	Read Only
	Platinum_Service	6	Deferred Pass	0x22	Low Drop	Priority 5	None	Read Only
	Premium_Service	7	Deferred Pass	0x2E	Low Drop	Priority 6	None	Read Only
	Network_Service	8	Deferred Pass	0x30	Low Drop	Priority 7	None	Read Only
	Null_Action	9	Deferred Pass	Ignore	Low Drop	Ignore	None	Read Only
	UntrustedClfrs1	64001	Deferred Pass	Derive from Ingress Priority	Low Drop	Ignore	None	Other
	UntrustedClfrs2	64002	Deferred Pass	0x0	High Drop	Priority 0	None	Other

Action Creation

Action Name	
Drop Frame	Deferred Pass
Update DSCP	Ignore
Set Drop Precedence	Low Drop
Update 802.1p Priority	Ignore
Extension	No Extension

2 In the Action Name field of the Action Creation section, enter **`Generic`**.

3 In the Drop Frame field, choose **`Yes`**.

4 In the Update DSCP field, choose **`47,0x2F,101111`**.

This entry changes the DSCP value to the decimal value 47 in the match packet.

5 In the Set Drop Precedence field, choose **`Low Drop`**.

6 In the Update 802.1p Priority field, select **`Priority 1`**.

Priority 1 specifies a low priority.

7 In the Extension Field, choose **`No Extension.`**

8 Click Submit.

The entry is displayed in the Action Table

In summary, you have configured a new action named Generic. This action specifies a low drop precedence, a user priority of 1, and a DSCP value of 0x2F for packets that match a classifier element associated with this action.

Configuring Interface Action Extensions

You create extensions to actions by using the Interface Action Extension page. These extensions allow you to filter on:

- Set an egress unicast
- Set an egress non-unicast

Note: An interface action extension can be referenced *only* by an action.

To create an interface action extension configuration:

1 From the main menu, choose Application > QoS >Interface Action Ext.

The Interface Action Extension page opens (Figure 110).

Figure 110 Interface Action Extension page

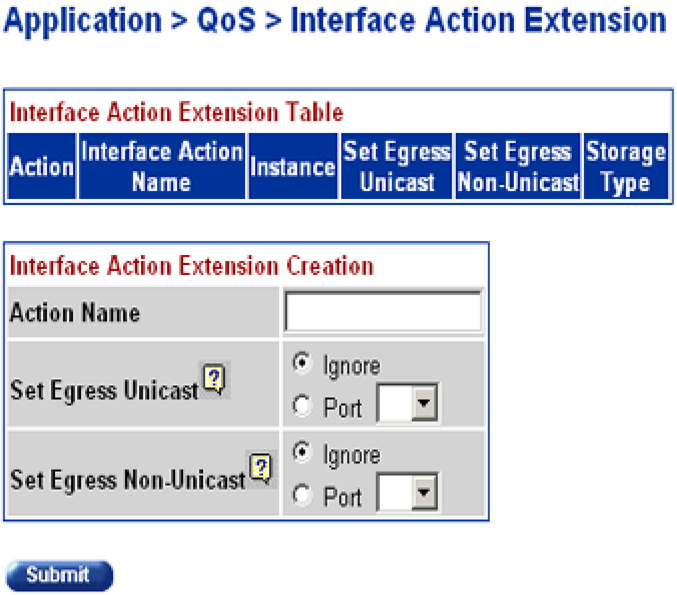

Application > QoS > Interface Action Extension

Interface Action Extension Table

Action	Interface Action Name	Instance	Set Egress Unicast	Set Egress Non-Unicast	Storage Type

Interface Action Extension Creation

Action Name	
Set Egress Unicast	● Ignore ○ Port
Set Egress Non-Unicast	● Ignore ○ Port

2 In the Action Name field, enter **`Sample`**.

3 In the Set Egress Unicast field, click **`Port`** and choose **`8`**.

4 In the Set Egress Non-Unicast field, click **`Ignore.`**

5 Click Submit.

The system returns you to the Interface Actions Extension page and displays your new entry on the Interface Action Extension Table.

Note: You cannot modify an interface action extension entry. You must delete that entry and create another one with the configuration you want.

Configuring Meters

Metering operates at ingress and provides different levels of service to data streams through user-configurable parameters. An example would be to limit traffic entering a port to a specified bandwidth, such as 2000 Kb/s (Committed Rate). Instead of dropping all traffic that exceeds this threshold, traffic policing allows you to configure a Committed Burst Rate to exceed the threshold (Committed Rate), for a brief period of time, without being dropped.

➡ **Note:** For more information on the Data metering, refer to "Configuring Policies" on page 293.

To configure a meter:

1 Click the Application > QoS > Meter menu option.

The Meter page opens (Figure 111).

Figure 111 Meter page

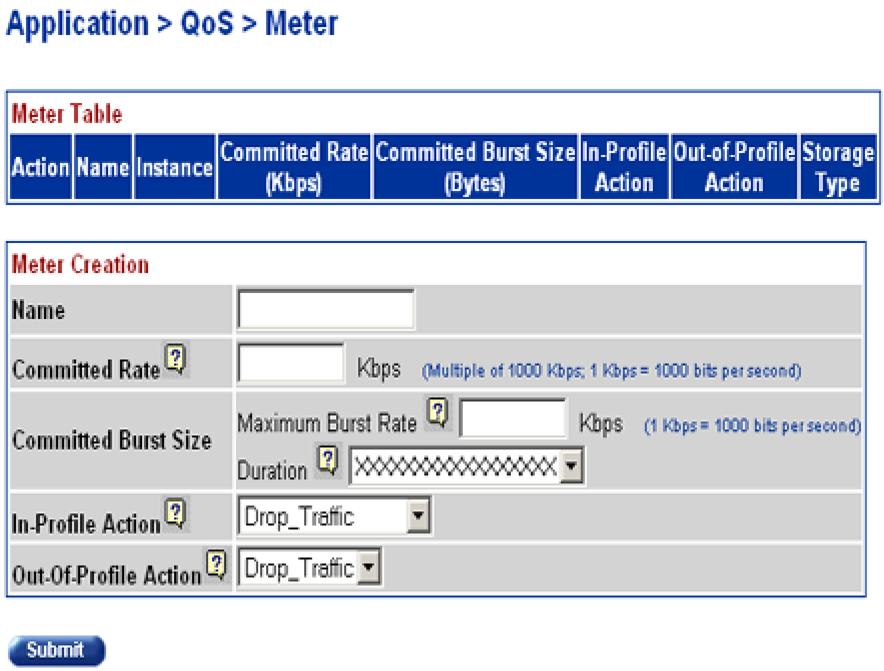

Application > QoS > Meter

Meter Table

Action	Name	Instance	Committed Rate (Kbps)	Committed Burst Size (Bytes)	In-Profile Action	Out-of-Profile Action	Storage Type

Meter Creation

Name	
Committed Rate	Kbps (Multiple of 1000 Kbps; 1 Kbps = 1000 bits per second)
Committed Burst Size	Maximum Burst Rate Kbps (1 Kbps = 1000 bits per second) Duration XXXXXXXXXXXXXXXXXX
In-Profile Action	Drop_Traffic
Out-Of-Profile Action	Drop_Traffic

2 In the Name field of the Meter Creation section, enter **`Practice`**.

3 In the Committed Rate field, enter **`3000`**.

4 In the Maximum Burst Rate field of the Committed Burst Size section, enter **`3500`**.

5 In the Duration field of the Committed Burst Size section, select **`262 milliseconds`** from the pull-down menu.

The switch calculates durations and presents the results to you in a pull-down menu. Choose the one you want.

6 In the In-Profile Action field, choose **`Generic`**.

7 In the Out-Of-Profile Action, choose **`Drop_Traffic`**.

8 Click Submit.

In summary, you have configured a new meter named Practice. This meter specifies committed data, with a committed rate of 3000 Kbps and a committed burst size of 16384 bytes, for packets that match a classifier associated with this meter.

Configuring Policies

Now you are ready to configure a policy. A policy is an interface group, a classifier or classifier block and the associated meter and action(s). Policies are applied according to the precedence order that you assign in the Policies page.

To configure a policy:

1 Click the Application > QoS > Policy menu option.

The Policy page opens (Figure 112).

Figure 112 Policy page

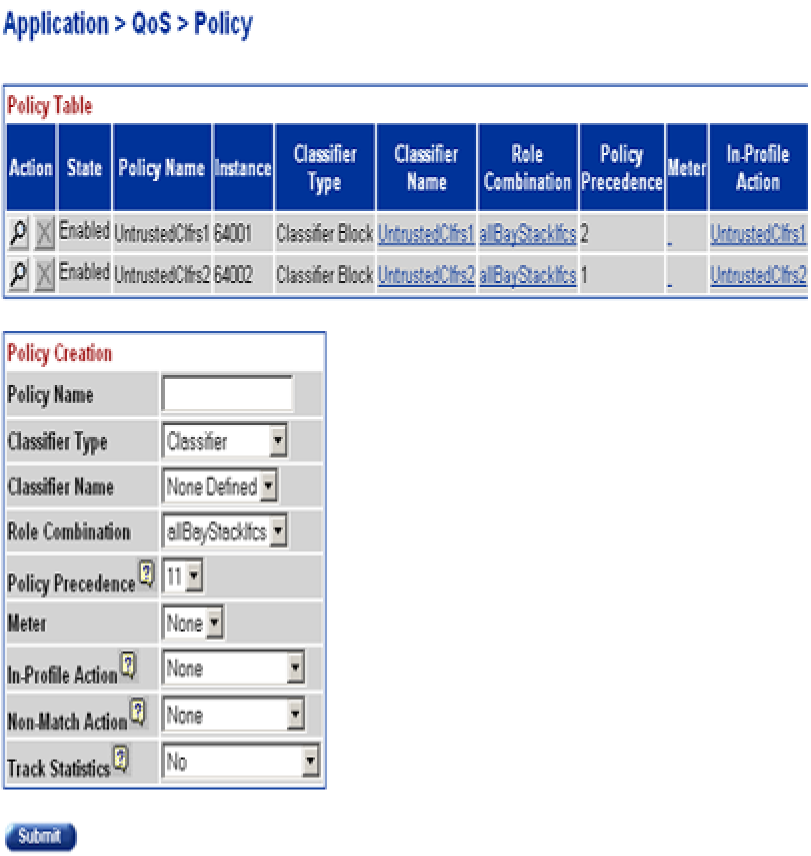

Application > QoS > Policy

Policy Table

Action	State	Policy Name	Instance	Classifier Type	Classifier Name	Role Combination	Policy Precedence	Meter	In-Profile Action
🔍 ✕	Enabled	UntrustedClfrs1	64001	Classifier Block	UntrustedClfrs1	allBayStackIfcs	2	_	UntrustedClfrs1
🔍 ✕	Enabled	UntrustedClfrs2	64002	Classifier Block	UntrustedClfrs2	allBayStackIfcs	1	_	UntrustedClfrs2

Policy Creation

Policy Name	
Classifier Type	Classifier
Classifier Name	None Defined
Role Combination	allBayStackIfcs
Policy Precedence	11
Meter	None
In-Profile Action	None
Non-Match Action	None
Track Statistics	No

2 In the Policy Name field of the Policy Creation area, enter **`Policy`**.

This entry is a unique name to identify this policy.

Note: Do not provide blank spaces in the naming field.

---

3 In the Classifier Type, choose **`Classifier Block`**.

4 In the Classifier Name field, choose **`Test`**.

5 In the Role Combination field, choose **`Webbrowsing`**.

This entry is the unique Role Combination that you created.

6 In the Policy Precedence field, enter **`6`**.

7 In the Meter field, choose **`Practice`**.

8 In the Non-Match Action field, choose **`Drop Traffic.`**

9 In the Track Statistics field, choose **`Aggregate Classifier`**.

10 Click Submit.

In summary, you configured a QoS policy called *Policy* (Figure 113). This policy applies a combination of packet classifier (matching) criteria and actions to individual interfaces (ports) in the hardware. You specified that this policy will use the *Test* classifier with the elements that you specified. *Policy* will use the Role Combination *Webbrowsing*, and the *Practice* meter. *Policy* specifies the type of behavior you want to apply to a flow of packets.

Note: A policy is enabled as soon as you create it.

Figure 113 Policy page with newly created policy

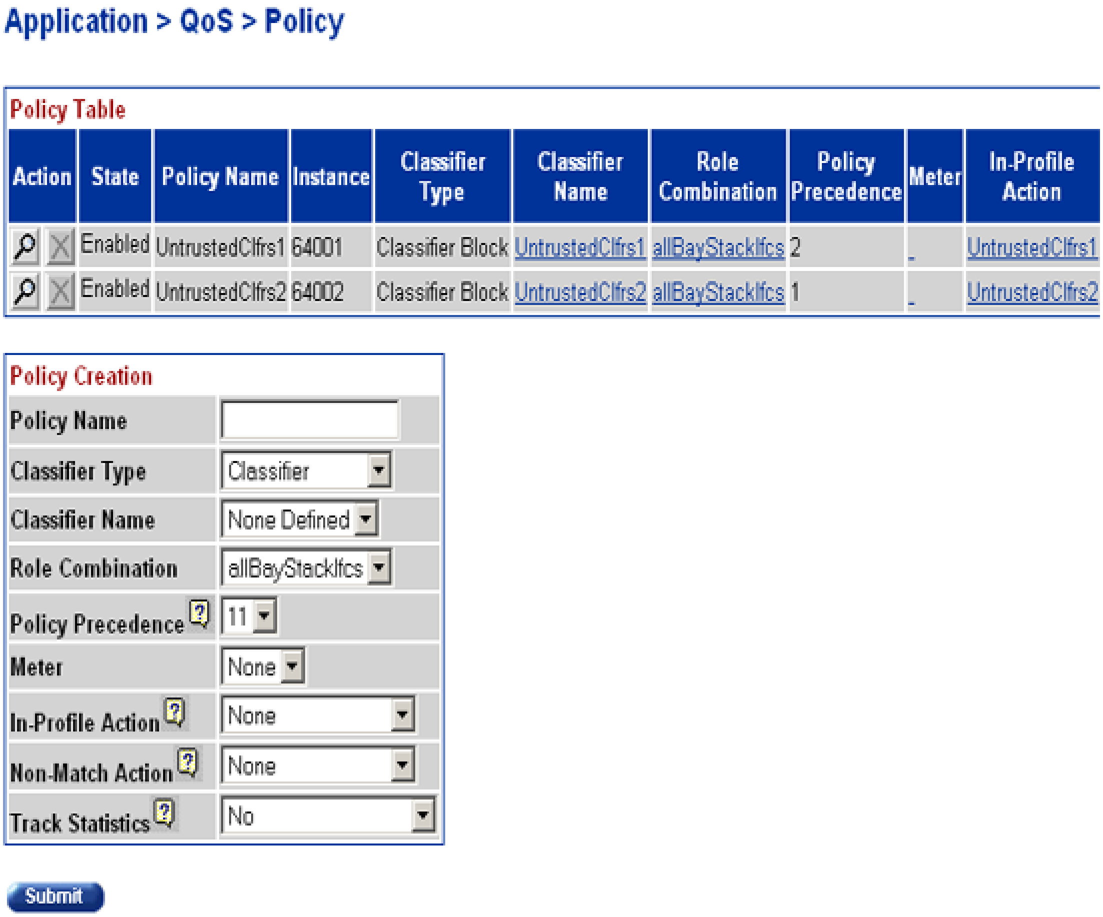

Application > QoS > Policy

Policy Table

Action	State	Policy Name	Instance	Classifier Type	Classifier Name	Role Combination	Policy Precedence	Meter	In-Profile Action
🔍 ✕	Enabled	UntrustedClfrs1	64001	Classifier Block	UntrustedClfrs1	allBayStackIfcs	2	_	UntrustedClfrs1
🔍 ✕	Enabled	UntrustedClfrs2	64002	Classifier Block	UntrustedClfrs2	allBayStackIfcs	1	_	UntrustedClfrs2

Policy Creation

Policy Name	
Classifier Type	Classifier ▾
Classifier Name	None Defined ▾
Role Combination	allBayStackIfcs ▾
Policy Precedence ?	11 ▾
Meter	None ▾
In-Profile Action ?	None ▾
Non-Match Action ?	None ▾
Track Statistics ?	No ▾

Assigning Mapping Values

To manually configure the mapping among 802.1p priority values, priority, and DSCP mapping, you must use with the following QoS pages:

- "Assigning 802.1p Priority Queue Assignment" on page 297, next
- "Assigning 802.1p User Priority Mapping" on page 298

- "Verifying DSCP Mapping" on page 300

Note: Nortel recommends that you use the default mapping values to ensure end-to-end QoS connectivity across Nortel products.

Assigning 802.1p Priority Queue Assignment

You assign IEEE 802.1p priority values to a queue. There are 8 different queue sets for which assignments can be made. This information is used for assigning egress traffic to outbound queues.

To configure 802.1p priority:

1. Click the Application > QoS > Devices > Priority Q Assign menu option.

   The 802.1p Priority Queue Assignment page opens (Figure 114).

Figure 114 802.1p Priority Queue Assignment page

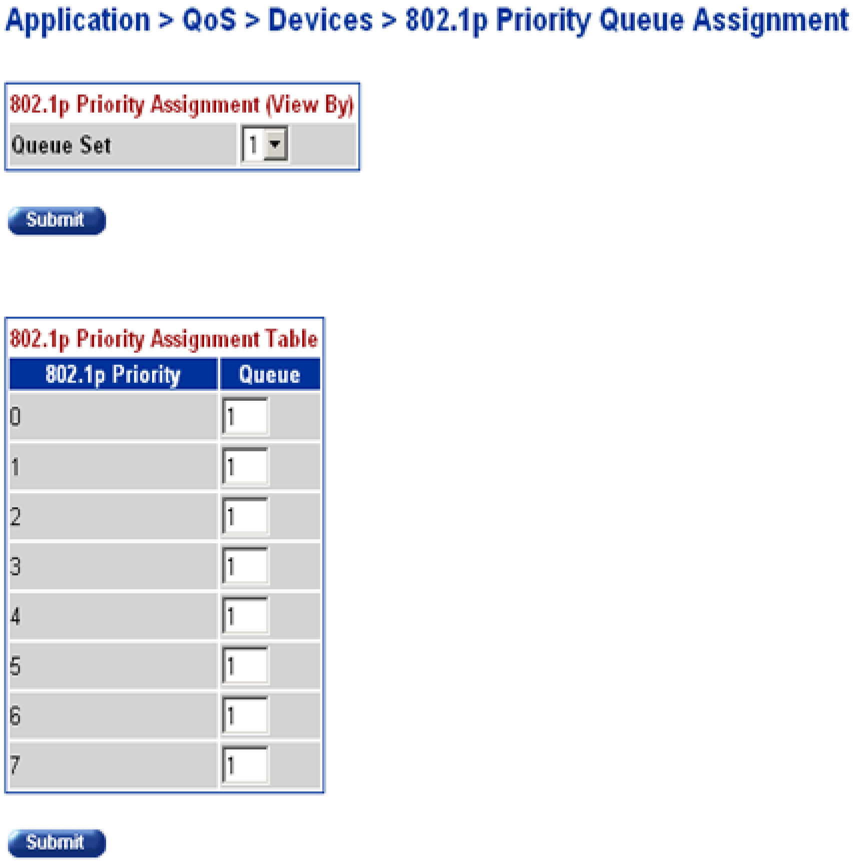

Application > QoS > Devices > 802.1p Priority Queue Assignment

802.1p Priority Assignment (View By)

Queue Set	1

802.1p Priority Assignment Table

802.1p Priority	Queue
0	1
1	1
2	1
3	1
4	1
5	1
6	1
7	1

2 In the Queue field, enter `1` next to the user priority of 6.

3 Click Submit.

The system returns you to the 802.1p Priority Queue Assignment page and displays the updated value.

Assigning 802.1p User Priority Mapping

To configure IEEE 802.1p user priority to DSCP mapping:

1 Click the Application > QoS > Devices > Priority Mapping menu option.

The 802.1p Priority Mapping page opens (Figure 115).

Figure 115 802.1p Priority Mapping page

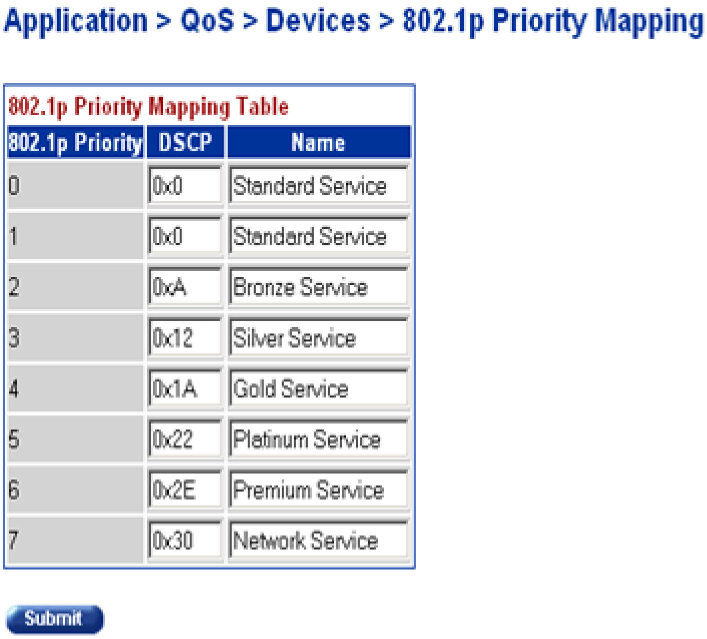

Application > QoS > Devices > 802.1p Priority Mapping

802.1p Priority Mapping Table

802.1p Priority	DSCP	Name
0	0x0	Standard Service
1	0x0	Standard Service
2	0xA	Bronze Service
3	0x12	Silver Service
4	0x1A	Gold Service
5	0x22	Platinum Service
6	0x2E	Premium Service
7	0x30	Network Service

2 Enter the DSCP value for specific 802.1p. Priority 2 as `0x0`.

3 Click Submit.

Verifying DSCP Mapping

Next, verify the mapping of the DSCP to an IEEE 802.1p priority, drop precedence, and service class.

➨ Click the Application > QoS >Devices > DSCP Mapping menu option.

The DSCP Mapping page opens (Figure 116).

Figure 116 DSCP Mapping page

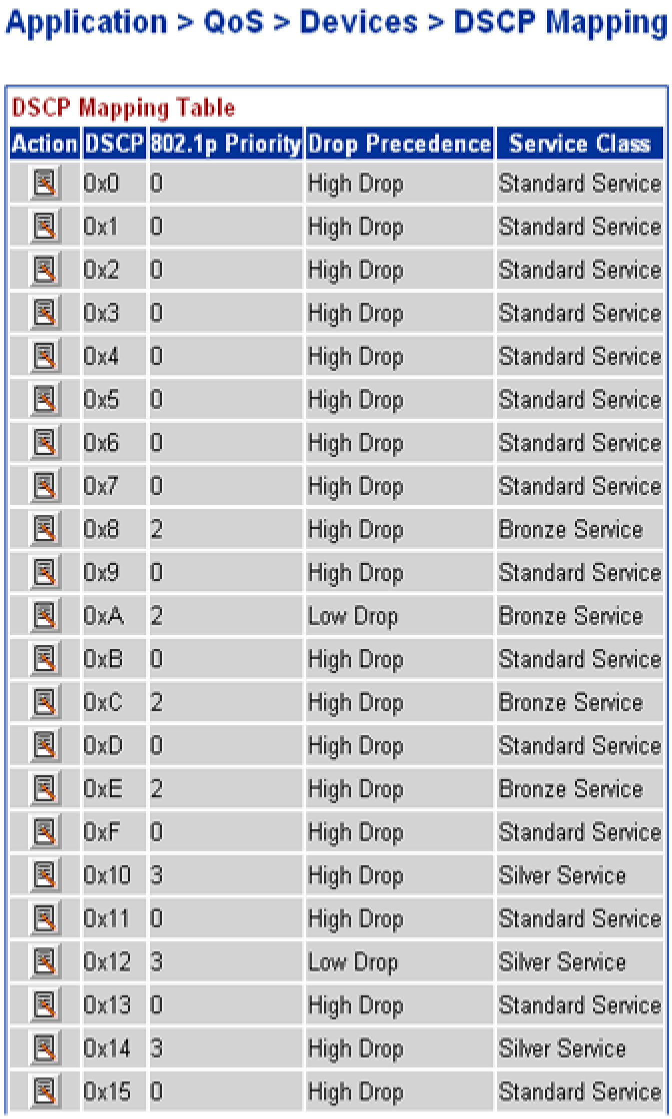

Application > QoS > Devices > DSCP Mapping

DSCP Mapping Table

Action	DSCP	802.1p Priority	Drop Precedence	Service Class
✎	0x0	0	High Drop	Standard Service
✎	0x1	0	High Drop	Standard Service
✎	0x2	0	High Drop	Standard Service
✎	0x3	0	High Drop	Standard Service
✎	0x4	0	High Drop	Standard Service
✎	0x5	0	High Drop	Standard Service
✎	0x6	0	High Drop	Standard Service
✎	0x7	0	High Drop	Standard Service
✎	0x8	2	High Drop	Bronze Service
✎	0x9	0	High Drop	Standard Service
✎	0xA	2	Low Drop	Bronze Service
✎	0xB	0	High Drop	Standard Service
✎	0xC	2	High Drop	Bronze Service
✎	0xD	0	High Drop	Standard Service
✎	0xE	2	High Drop	Bronze Service
✎	0xF	0	High Drop	Standard Service
✎	0x10	3	High Drop	Silver Service
✎	0x11	0	High Drop	Standard Service
✎	0x12	3	Low Drop	Silver Service
✎	0x13	0	High Drop	Standard Service
✎	0x14	3	High Drop	Silver Service
✎	0x15	0	High Drop	Standard Service

To change the DSCP to an 802.1p priority:

1. Click the Application > QoS > Devices > DSCP Mapping menu option.

   The DSCP Mapping page opens (Figure 116).

2 Click the Modify icon of DSCP 0x1.

The DSCP Mapping page opens (Figure 117) for DSCP 0x1.

Figure 117 DSCP Mapping page to modify mapping

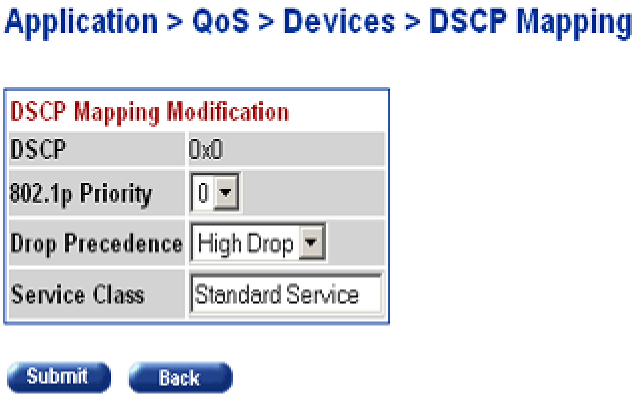

Application > QoS > Devices > DSCP Mapping

DSCP Mapping Modification

DSCP	0x0
802.1p Priority	0
Drop Precedence	High Drop
Service Class	Standard Service

3 In the 802.1 User Priority field, choose **`1`**.

4 In the Drop Precedence field, choose **`Low Drop`**.

5 In the Service Class field, enter the name for the service **`New Service`**.

6 Click Submit.

Chapter 6
Troubleshooting

This chapter describes how to isolate and diagnose problems with your Ethernet Routing Switch 3510-24T and covers the following topics:

- "Interpreting the LEDs", next
- "Diagnosing and Correcting Problems" on page 306

The chapter topics lead you through a logical process for troubleshooting the Ethernet Routing Switch 3510-24T. For example, to understand the various states that your switch LEDs can exhibit during normal operation, see "Interpreting the LEDs".

For more help in determining the problem, "Diagnosing and Correcting Problems" describes symptoms and corrective actions you can perform to resolve specific problems. Subsequent sections give step-by-step procedures to correct the problems.

Interpreting the LEDs

Figure 118 shows the Ethernet Routing Switch 3510-24T LED display panel. See Table 70 for a description of the Ethernet Routing Switch 3510-24T LEDs.

Figure 118 Ethernet Routing Switch 3510-24T LED display panel

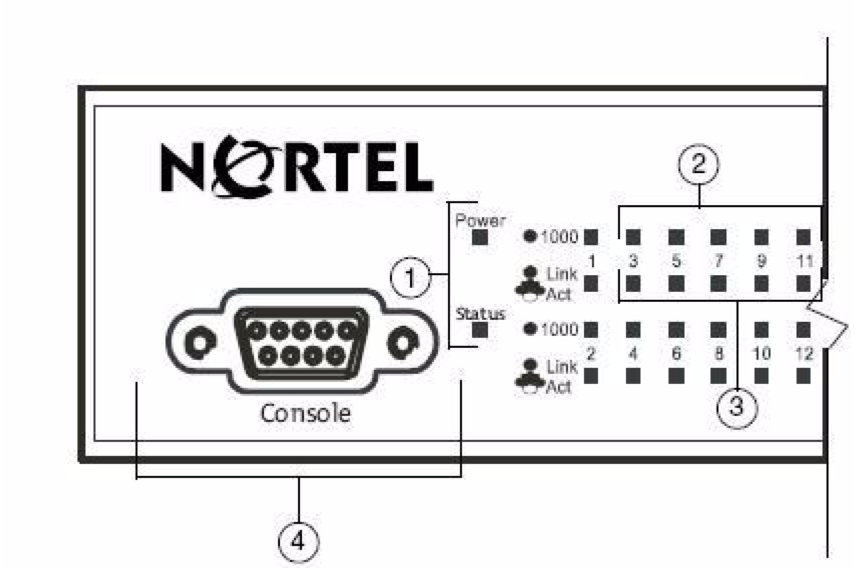

1 = Switch LEDs
2 = 10/100/1000-Mbps LEDs
3 = Link/Activity LEDs
4 = Console Port

Table 70 Ethernet Routing Switch 3510-24T LED descriptions

Label	State	Meaning
Power	On	The switch is connected to AC power and is receiving power.
	Off	The switch is not connected to AC power, or the AC power is not supplying power.
Status	Steady	The power-on self-test is complete, and the switch is operating normally.
	Blinking	A nonfatal error occurred during the self-test.
	Off	The switch failed the self-test.

Table 71 10/100/100 Port LEDs on the Ethernet Routing Switch 3510-24T

Label	Color/State	Meaning
1000	On	The port is set to operate at 1000 Mbps.
	Off	The port is set to operate at 10 or 100 Mbps (and LNK/ACT is green). When the LED is off, refer to the LNK/ACT Section.

Table 71 10/100/100 Port LEDs on the Ethernet Routing Switch 3510-24T (continued)

Label	**Color/State**	**Meaning**
LNK/ACT	Steady	The link is good.
	Blinking	There is activity on this port. The blinking rate indicates the level of activity.
	Slow blinking	This port has been disabled by software.
	Off	This port has no link or activity.

➡ **Note:** The SFP GBIC Ports are shared with the last four RJ-45 Ports.

LED Indications during Software Download Process

When you download software to the Ethernet Routing Switch 3510-24T, the port LEDs light one after another in a chasing pattern (except the LEDs on ports 11, 12, 23, and 24 on a Ethernet Routing Switch 3510-24T).

While downloading the image, the pattern is fast, and then the pattern slows as the switch erases the flash memory. The pattern moves very fast as the switch programs the new image into the switch's memory. When the process is complete, the port LEDs are no longer lit and the switch resets.

Diagnosing and Correcting Problems

Before you perform the problem-solving steps in this section, power the Ethernet Routing Switch 3510-24T (disconnect and then reconnect the AC power cord); then check if the switch follows the normal power-up sequence.

Warning: To avoid bodily injury from hazardous electrical current, never remove the top cover of the device. There are no user-serviceable components inside.

Vorsicht: Um Verletzungsgefahr durch einen elektrischen Stromschlag auszuschließen, nehmen Sie niemals die obere Abdeckung vom Gerät ab. Im Geräteinnern befinden sich keine Komponenten, die vom Benutzer gewartet werden können.

Avertissement: Pour éviter tout risque d'électrocution, ne jamais retirer le capot de l'appareil. Cet appareil ne contient aucune pièce accessible par l'utilisateur.

Advertencia: A fin de evitar daños personales por corrientes eléctricas peligrosas, no desmonte nunca la cubierta superior de este dispositivo. Los componentes internos no son reparables por el usuario.

Avvertenza: Per evitare lesioni fisiche dovute a scariche pericolose di corrente, non rimuovere mai il coperchio superiore del dispositivo. I componenti interni non possono essere manipolati dall'utente.

警告: 危険な電流から身体を保護するために、ディバイスの上部カバーを決して取り外さないでください。内部には、ユーザが扱うコンポーネントはありません。

Normal power-up sequence

In a normal power-up sequence, the LEDs appear as follows:

1. After power is applied to the switch, the Power LED turns on within 5 seconds.
2. The switch initiates a self-test, during which the port LEDs display various patterns to indicate the progress of the self-test.
3. After the self-test, the remaining port LEDs indicate their operational status.

Port connection problems

Table 72 Corrective actions

Symptom	Probable cause	Corrective action
All LEDs are off.	The switch is not receiving AC power.	Verify that the AC power cord is fastened securely at both ends and that power is available at the AC power outlet.
	The fans are not operating or the airflow is blocked, causing the unit to overheat.	Verify that there is sufficient space for adequate airflow on both sides of the switch.
		NOTE: Operating temperature for the switch must not exceed 45°C (113°F). Do not place the switch in areas where it can be exposed to direct sunlight or near warm air exhausts or heaters.
The Act LED for a connected port is off or does not blink (and you have reason to believe that traffic is present).	The switch is experiencing a port connection problem. The switch's link partner is not autonegotiating properly.	See "Port connection problems" next.

You can usually trace port connection problems to either a poor cable connection or an improper connection of the port cables at either end of the link. To remedy these types of problems, make sure that the cable connections are secure and that the cables connect to the correct ports at both ends of the link

Port connection problems are also traceable to the autonegotiation mode or the port interface.

Autonegotiation Modes

Port connection problems can occur when a port (or station) is connected to another port (or station) that is not operating in a compatible mode (for example, connecting a full-duplex port on one station to a half-duplex port on another station).

The Ethernet Routing Switch 3510-24T negotiates port speeds according to the IEEE 802.3u, IEEE 802.3z, and IEEE 802.3ab autonegotiating standards. The switch adjusts (autonegotiates) its port speed and duplex mode to match the best service provided by the connected station, up to 1000 Mbps in full-duplex mode as follows:

- If the connected station uses a form of autonegotiation that is not compatible with the IEEE autonegotiating standard, the Ethernet Routing Switch 3510-24T cannot negotiate a compatible mode for correct operation.
- If the autonegotiation feature is not present or not enabled at the connected station, the Ethernet Routing Switch 3510-24T may not be able to determine the correct duplex mode.

In both situations, the Ethernet Routing Switch 3510-24T “autosenses” the speed of the connected station and, by default, reverts to half-duplex mode. If the connected station is operating in full-duplex mode, communication errors may occur.

To correct this mode mismatch problem:

1 Use the Port Configuration screen to disable autonegotiation for the suspect port.

2 Manually set the Speed/Duplex field to match the speed/duplex mode of the connected station.

You may have to try several settings before you find the correct speed/duplex mode of the connected station.

If the problem persists:

1 Disable the autonegotiation feature at the connected station.

2 Manually set the speed/duplex mode of the connected station to the same speed/duplex mode you have manually set for the Ethernet Routing Switch 3510-24T port.

Port Interface

Note: You can use Category 5 copper unshielded twisted pair (UTP) cable when you are running at speeds of 10 Mbps, 100 Mbps, or 1000 Mbps. You can use Category 3 cable when you are running at speeds of 10 Mbps.
If you use Category 3 cable, you must disable autonegotiation when the device connected to the Ethernet Routing Switch 3510-24T supports a 100-Mbps or 1000-Mbps connection. If the connected device supports only 10-Mbps connection, autonegotiation may be enabled.

Ensure that the devices are connected using the appropriate crossover or straight-through cable (see Appendix C), or that autonegotiation is active.

Note: IEEE 1000BASE-TX requires full-duplex mode operation with autonegotiation enabled.

Note: Auto-MDI-X and auto-polarity both require that auto-negotiation be enabled.

The Ethernet Routing Switch has four shared front-panel ports: 21, 22, 23 and 24 on the Ethernet Routing Switch 3510-24T. If you insert an SFP GBIC into one on these ports, that port handles gigabit Ethernet speed only. With no optional SFP GIBCs inserted, these ports function as the other 10/100/1000 front-panel ports. The autonegotiation configuration settings are set only once for each port, so these settings are the same for these two sockets that share a port. Autonegotiation is enabled or disabled based on the configuration for these shared ports. However, the speed and duplex settings are ignored on that port once you insert an SFP GBIC.

Note: Ensure that you have both sides configured identically when you are using the SFP GBIC, or you may lose connectivity.

Appendix A
Technical Specifications

This appendix provides technical specifications for the Ethernet Routing Switch 3510-24T.

Environmental

Table 73 lists environmental specifications.

Table 73 Environmental specifications

Parameter	Operating specification	Storage specification
Temperature	0° to 45°C (32° to 113°F)	-25° to 70°C (-13° to 158°F)
Humidity	85% maximum relative humidity, noncondensing	95% maximum relative humidity, noncondensing
Altitude	3024 m (10,000 ft)	3024 m (10,000 ft)

Electrical

Table 74 lists power electrical parameters for the Ethernet Routing Switch 3510-24T.

Table 74 Electrical parameters

Parameter	Electrical specification
Input Voltage	100 to 240 VAC @ 47 to 63 Hz
Input Power Consumption	135 W maximum

Table 74 Electrical parameters (continued)

Input current	1.3 A @ 115 VAC 0.65 A @ 230 VAC
Maximum thermal output	460 BTU/hr

Physical Dimensions

Table 75 lists physical dimensions for Ethernet Routing Switch 3510-24T.

Table 75 Physical dimensions for Ethernet Routing Switch 3510-24T

Parameter	Specifications
Height	4.45 cm (1.75 in.)
Width	48.32 cm (17.25 in.)
Depth	38.74 cm (15.25 in)
Weight for Ethernet Routing Switch 3510-24T	3.0 kg (7 lbs)

Fan Noise

The fan noise is less than 56 dB for the Ethernet Routing Switch 3510-24T.

Performance Specifications

Table 76 lists performance specifications.

Table 76 Performance specifications

Parameter	Specifications
Frame Forward Rate (64-byte packets)	• Ethernet Routing Switch 3510-24T: 35,714,285 pps maximum
Port Forwarding/Filtering Performance (64-byte packets)	• For 10 Mbps: 14,880 pps maximum • For 100 Mbps: 148,810 pps maximum • For 1 Gbps: 1,488,095 pps maximum
Address Database Size	• Ethernet Routing Switch 3510-24T: 8000 entries

Table 76 Performance specifications (continued)

Addressing	48-bit MAC address
Frame Length	64 to 1518 bytes (IEEE 802.1Q Untagged) 68 to 1522 bytes (IEEE 802.1Q Tagged) 1519 to 9216 bytes (jumbo frame IEEE 802.1Q Untagged) 1523 to 9216 bytes (jumbo frame IEEE 802.1Q Tagged)

Interface Options

The Ethernet Routing Switch 3510-24T has 10/100/1000BASE-TX switch ports with RJ-45 (8-pin modular) connectors for MDI-X interfaces.

Note: In **Autonegotiation**, mode, the Ethernet Routing Switch 3510-24T automatically provides the proper MDI/MDI-X connection on the RJ-45ports, thereby eliminating the need for crossover cables. in **Forced Negotiation** mode, the RJ-45 ports provide an MDI-X connection, which allows end-station equipment to be connected using straight-through cables. To connect other MDI-X port devices, such as another switch or a hub, a crossover cable must be used.
The Ethernet Routing Switch 3510-24T supports Auto MDI/MDI-X only when autonegotation is enabled.

Refer to *Installing Gigabit Interface Converters and Small Form Factor Pluggable Interface Converters* for information on the interface connectors on available uplink modules.

Safety Agency Certification

The safety certifications follow:

- UL EN60950 (UL Listed or CSA 22.2 No. 60950 (CUL))
- IEC 60950/EN 60950, CB report and certificate with all national deviations
- C22.2 No. 950 (CUL) with all national deviations

- UL-94-V1 flammability requirements for PC board
- NOM-019 (NOM)

Electromagnetic Emissions

The Ethernet Routing Switch 3510-24T meets the following standards:

- US. CFR47, Part 15, Subpart B, Class A
- Canada. ICES-003, Issue 3, Class A
- Australia/New Zealand. AS/NZS 3548:1995, A1:1997/A2:1997 class A
- Japan. VCCI-V-3/02.04 class A
- Taiwan. CNS 13438, Class A
- EN55022:1998/A1:2000
- EN61000-3-2:2000
- EN61000-3-3:1995/A1:2001

Electromagnetic Immunity

The module meets the EN55024:1998/A1:2001 standard.

Declaration of Conformity

The Declaration of Conformity for the Ethernet Routing Switch 3510-24T complies with ISO/IEC Guide 22 and EN45014. The declaration identifies the product models, Nortel name and address, and the specifications recognized by the European community.

As stated in the Declaration of Conformity, the Ethernet Routing Switch 3510-24T complies with the provisions of Council Directives 89/336/EEC and 73/23/EEC.

Appendix B
Quick Steps to Features

If you are a system administrator with experience configuring Ethernet Routing Switch 3510-24T VLANs, MultiLink Trunking, Port Mirroring, IGMP Snooping, and EAPOL authentication processes, use the flowcharts on the following pages as quick configuration guides. The flowcharts refer you to the "configuration rules" appropriate for each feature.

The flowcharts cover the following features:

- "Configuring 802.1Q VLANs" on page 315, next
- "Configuring MultiLink Trunks" on page 319
- "Configuring Port Mirroring" on page 320
- "Configuring IGMP Snooping" on page 322

Configuring 802.1Q VLANs

To create or modify an 802.1Q VLAN, follow the flowcharts in Figure 119, Figure 120, and Figure 121.

To open the VLAN Configuration screen:

➨ Choose VLAN Configuration (or press v) from the VLAN Configuration Menu screen.

Figure 119 Configuring 802.1Q VLANs (1 of 3)

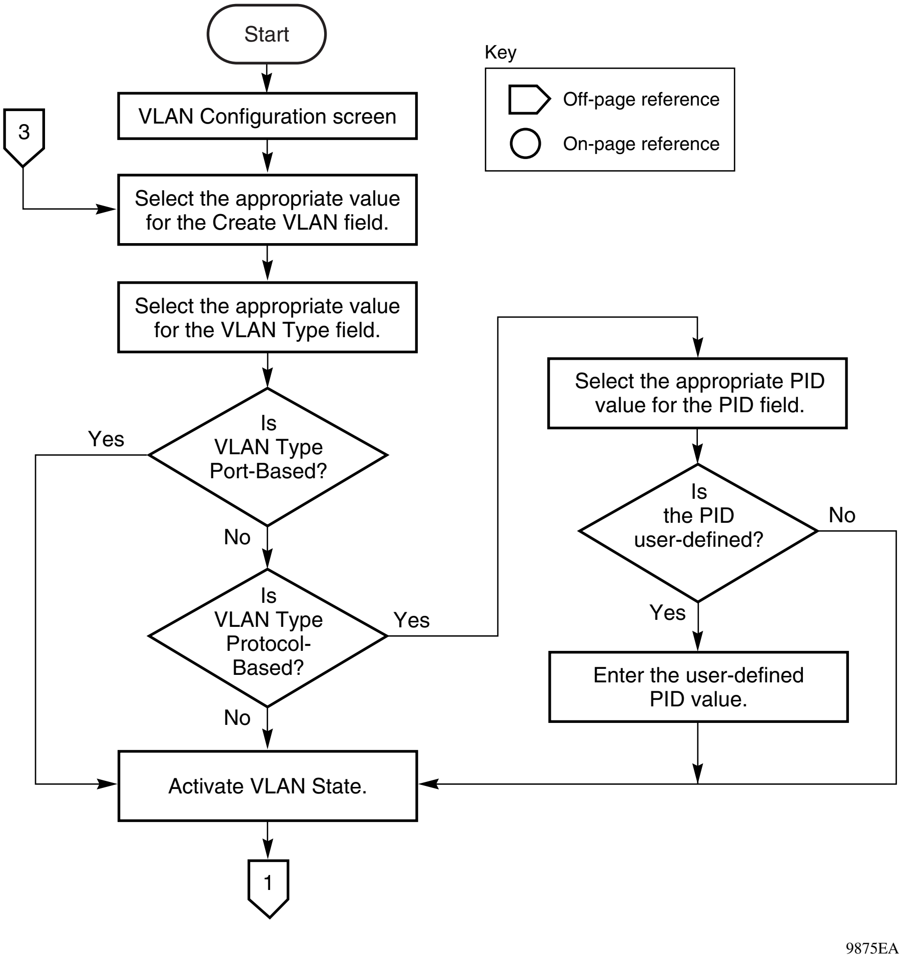
Start
Key
Off-page reference
On-page reference
VLAN Configuration screen
3
Select the appropriate value for the Create VLAN field.
Select the appropriate value for the VLAN Type field.
Select the appropriate PID value for the PID field.
Yes
Is VLAN Type Port-Based?
No
Is the PID user-defined?
No
Yes
Is VLAN Type Protocol-Based?
Yes
No
Enter the user-defined PID value.
Activate VLAN State.
1
9875EA

Figure 120 Configuring 802.1Q VLANs (2 of 3)

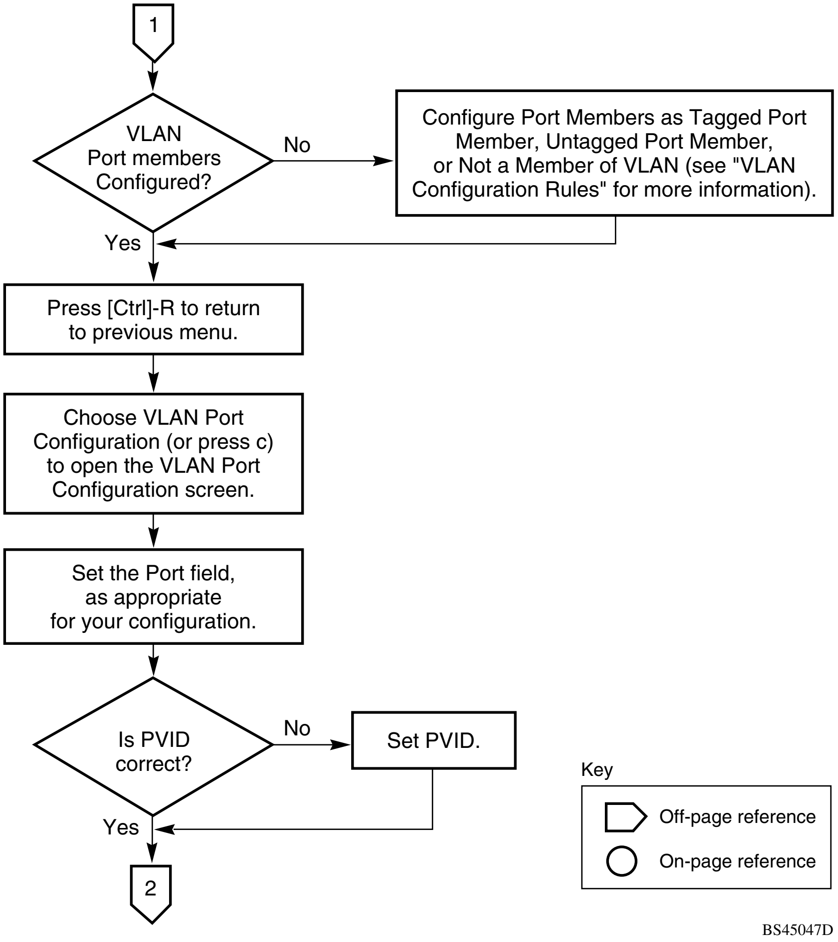
1
VLAN Port members Configured?
No
Configure Port Members as Tagged Port Member, Untagged Port Member, or Not a Member of VLAN (see "VLAN Configuration Rules" for more information).
Yes
Press [Ctrl]-R to return to previous menu.
Choose VLAN Port Configuration (or press c) to open the VLAN Port Configuration screen.
Set the Port field, as appropriate for your configuration.
Is PVID correct?
No
Set PVID.
Yes
2
Key
Off-page reference
On-page reference
BS45047D

Figure 121 Configuring 802.1Q VLANs (3 of 3)

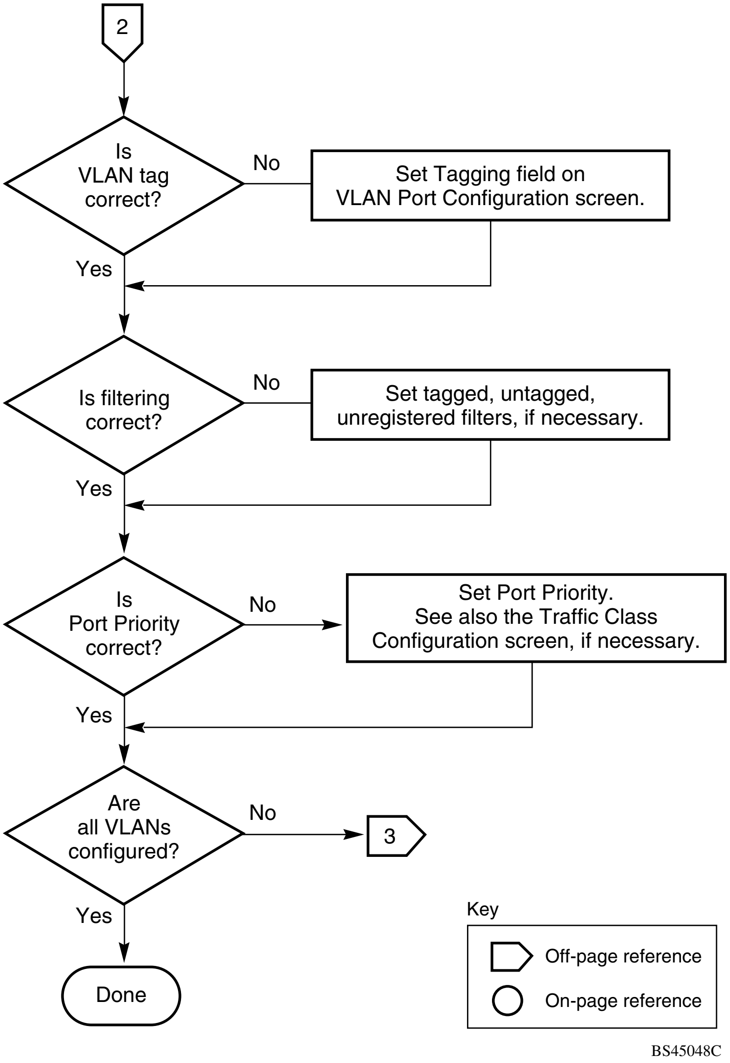
2
Is VLAN tag correct?
No
Set Tagging field on VLAN Port Configuration screen.
Yes
Is filtering correct?
No
Set tagged, untagged, unregistered filters, if necessary.
Yes
Is Port Priority correct?
No
Set Port Priority. See also the Traffic Class Configuration screen, if necessary.
Yes
Are all VLANs configured?
No
3
Yes
Done
Key
Off-page reference
On-page reference
BS45048C

Configuring MultiLink Trunks

To create or modify a MultiLink Trunk, follow the flowchart in Figure 122.

To open the MultiLink Trunk Configuration screen:

➨ Choose MultiLink Trunk Configuration (or press t) from the MultiLink Trunk Configuration Menu screen

Figure 122 Configuring MultiLink Trunks

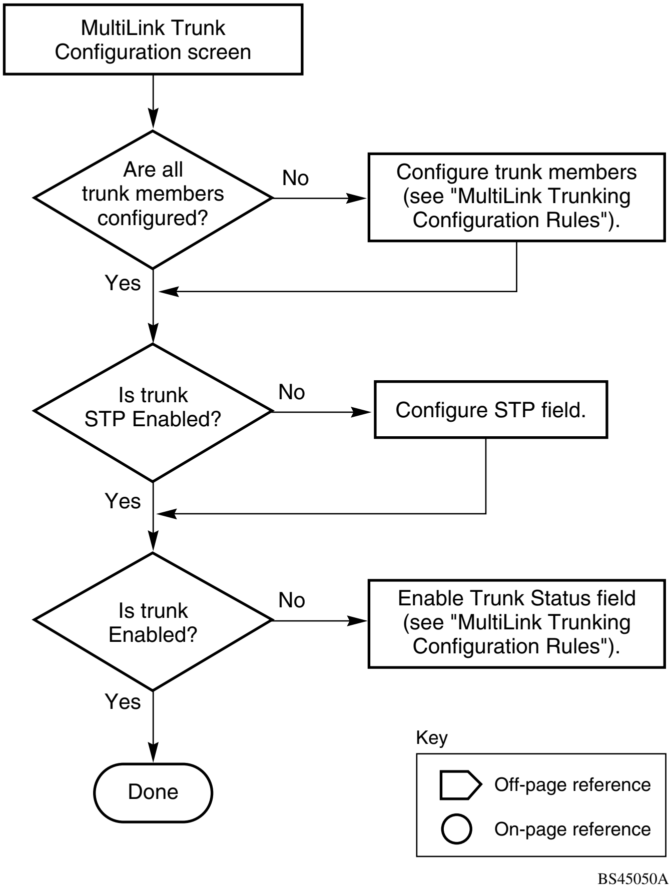
MultiLink Trunk Configuration screen
Are all trunk members configured?
No
Configure trunk members (see "MultiLink Trunking Configuration Rules").
Yes
Is trunk STP Enabled?
No
Configure STP field.
Yes
Is trunk Enabled?
No
Enable Trunk Status field (see "MultiLink Trunking Configuration Rules").
Yes
Done
Key
Off-page reference
On-page reference
BS45050A

Configuring Port Mirroring

To create or modify port-mirroring ports, follow the flowcharts in Figure 123 and Figure 124).

To open the Port Mirroring Configuration screen:

➨ Choose Port Mirroring Configuration (or press i) from the Switch Configuration Menu screen

Figure 123 Configuring Port Mirroring (1 of 2)

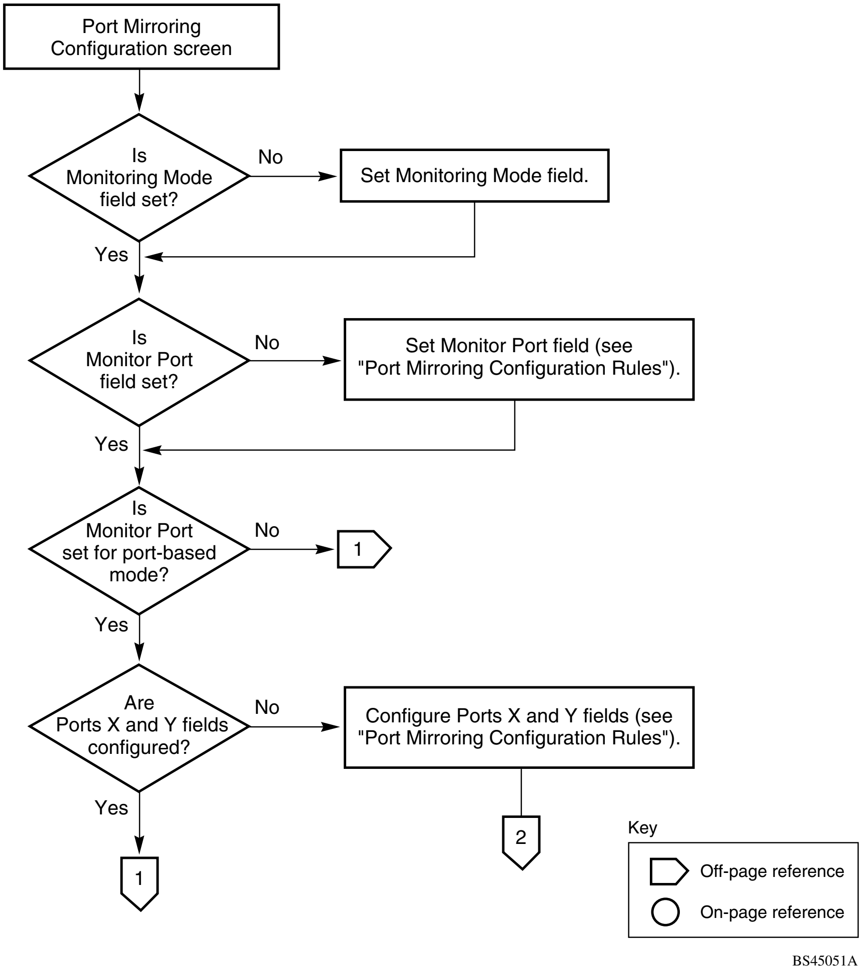
Port Mirroring Configuration screen
Is Monitoring Mode field set?
No
Set Monitoring Mode field.
Yes
Is Monitor Port field set?
No
Set Monitor Port field (see "Port Mirroring Configuration Rules").
Yes
Is Monitor Port set for port-based mode?
No
1
Yes
Are Ports X and Y fields configured?
No
Configure Ports X and Y fields (see "Port Mirroring Configuration Rules").
Yes
1
2
Key
Off-page reference
On-page reference
BS45051A

Figure 124 Configuring Port Mirroring (2 of 2)

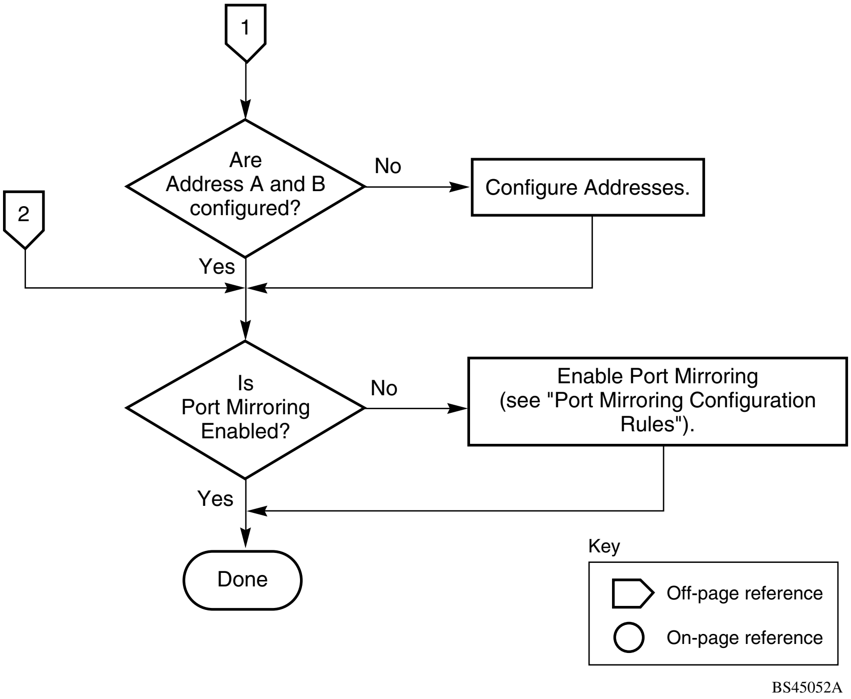
1
Are Address A and B configured?
No
Configure Addresses.
2
Yes
Is Port Mirroring Enabled?
No
Enable Port Mirroring (see "Port Mirroring Configuration Rules").
Yes
Done
Key
Off-page reference
On-page reference
BS45052A

Configuring IGMP Snooping

To create or modify IGMP Snooping ports, follow the flowcharts in Figure 125 to Figure 127.

To open the IGMP Configuration screen:

➨ Choose IGMP Configuration (or press g) from the Switch Configuration Menu screen.

Figure 125 Configuring IGMP Snooping (1 of 3)

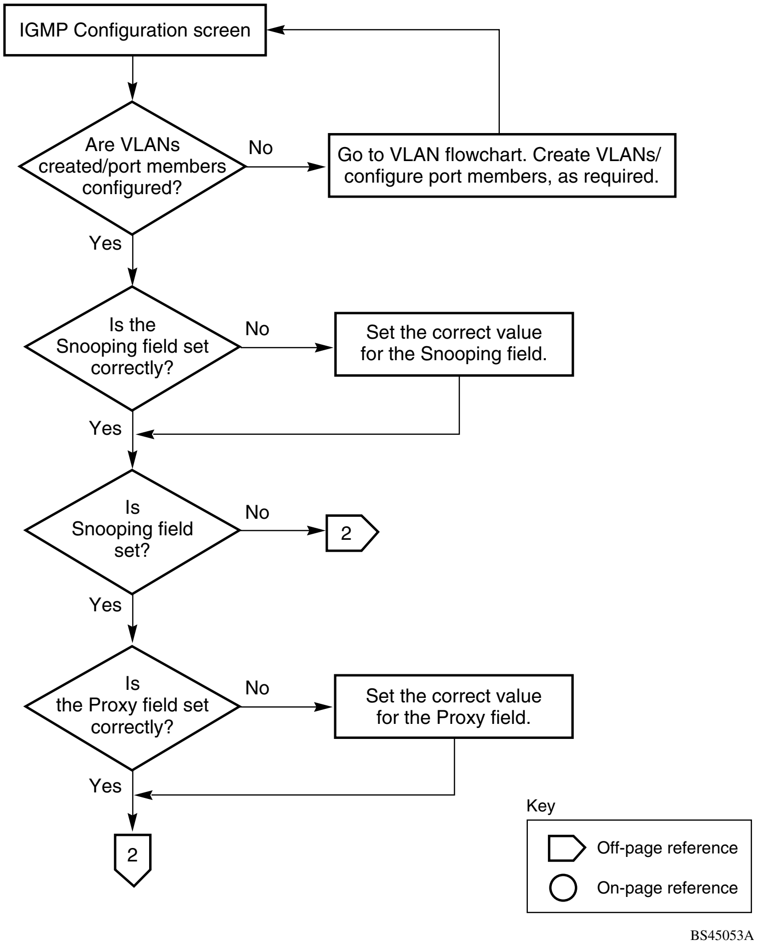
IGMP Configuration screen
Are VLANs created/port members configured?
No
Go to VLAN flowchart. Create VLANs/ configure port members, as required.
Yes
Is the Snooping field set correctly?
No
Set the correct value for the Snooping field.
Yes
Is Snooping field set?
No
2
Yes
Is the Proxy field set correctly?
No
Set the correct value for the Proxy field.
Yes
2
Key
Off-page reference
On-page reference
BS45053A

Figure 126 Configuring IGMP Snooping (2 of 3)

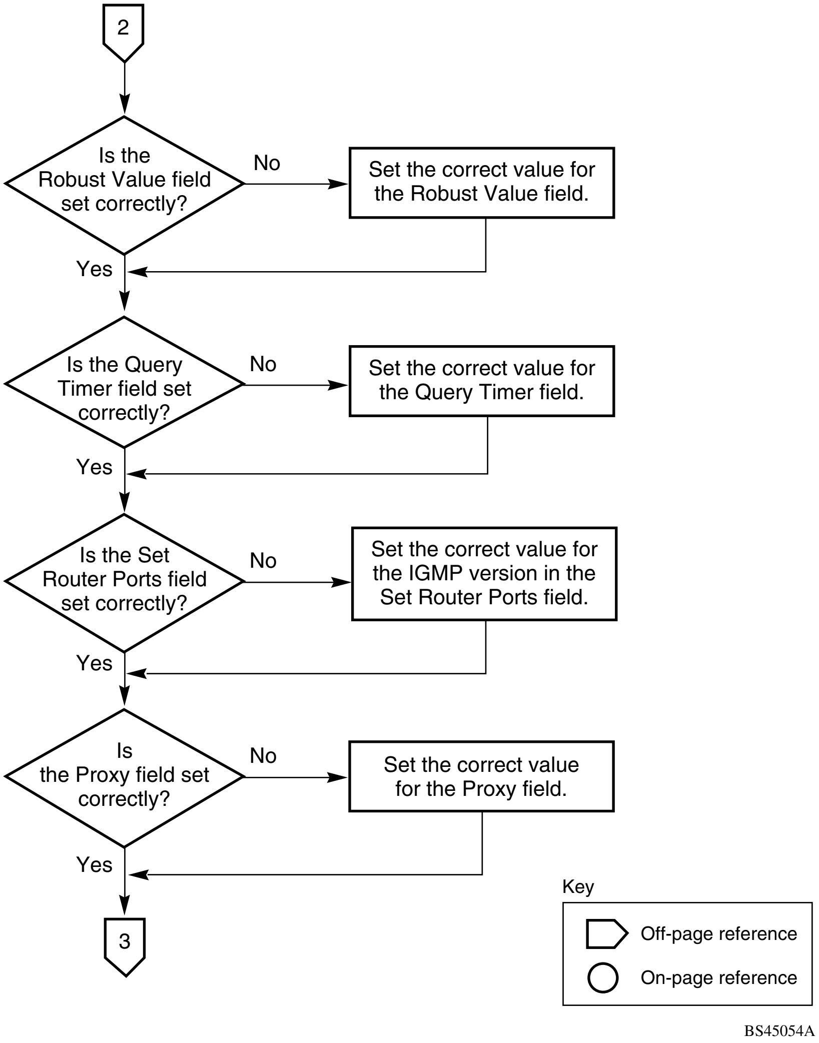
2
Is the Robust Value field set correctly?
No
Set the correct value for the Robust Value field.
Yes
Is the Query Timer field set correctly?
No
Set the correct value for the Query Timer field.
Yes
Is the Set Router Ports field set correctly?
No
Set the correct value for the IGMP version in the Set Router Ports field.
Yes
Is the Proxy field set correctly?
No
Set the correct value for the Proxy field.
Yes
3
Key
Off-page reference
On-page reference
BS45054A

Figure 127 Configuring IGMP Snooping (3 of 3)

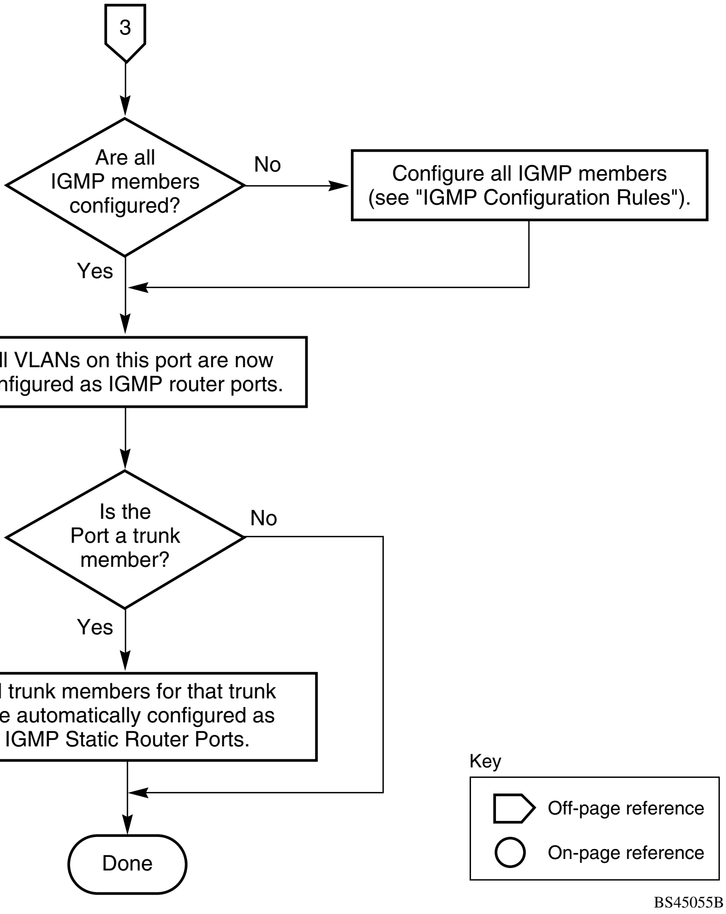
3
Are all IGMP members configured?
No
Configure all IGMP members (see "IGMP Configuration Rules").
Yes
All VLANs on this port are now configured as IGMP router ports.
Is the Port a trunk member?
No
Yes
All trunk members for that trunk are automatically configured as IGMP Static Router Ports.
Done
Key
Off-page reference
On-page reference
BS45055B

Appendix C
Connectors and Pin Assignments

This appendix describes the Ethernet Routing Switch 3510-24T port connectors and pin assignments.

RJ-45 (10BASE-T/100BASE-TX/1000BASE-TX) Port Connectors

Note: You can use Category 5 copper Unshielded Twisted Pair (UTP) cable when you are running at speeds of 10 Mbps, 100 Mbps, or 1000 Mbps. You can use Category 3 cable when you are running at speeds of 10 Mbps.

If you use Category 3 cable, you must disable autonegotiation when the device connected to the Ethernet Routing Switch 3510-24T supports a 100-Mbps or 1000-Mbps connection. If the connected device supports only 10-Mbps connection, autonegotiation may be enabled.

The RJ-45 port connectors (Figure 128) are wired as MDI-X ports to connect end stations without using crossover cables. (See "MDI and MDI-X Devices" on page 327 for information about MDI-X ports.) For 10BASE-T connections, use Category 3 (or higher) UTP cable. For 100BASE-TX/1000BASE-TX connections, use only Category 5 UTP cable.

Figure 128 RJ-45 (8-pin modular) port connector

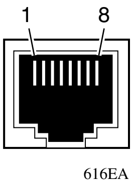

Table 77 lists the RJ-45 (8-pin modular) port connector pin assignments.

Table 77 RJ-45 port connector pin assignments

Pin	Signal	Description
1	RX+	Receive Data +
2	RX-	Receive Data -
3	TX+	Transmit Data +
4	Not applicable	Not applicable
5	Not applicable	Not applicable
6	TX-	Transmit Data -
7	Not applicable	Not applicable
8	Not applicable	Not applicable

For 1000BASE-T, all 8 pins are used for four pairs of bi-directional data.

Table 78 lists the types of bi-directional data for each of the 1000BASE-T pin connectors.

Table 78 1000BASE-T Pin Connectors

Pin	Type of Data
1	Bi-directional data A+
2	Bi-directional data A-
3	Bi-directional data B+
4	Bi-directional data C+
5	Bi-directional data C-
6	Bi-directional data B-
7	Bi-directional data D+
8	Bi-directional data D-

MDI and MDI-X Devices

Media Dependent Interface (MDI) is the IEEE standard for the interface to UTP cable.

For two devices to communicate, the transmitter of one device must connect to the receiver of the other device. The connection is established through a crossover function, which can be a crossover cable or a port that implements the crossover function internally.

Ports that implement the crossover function internally are known as MDI-X ports, where X refers to the crossover function.

Note: For the transmitter of one device to connect to the receiver of another device, the total number of crossovers must always be an odd number.

The following sections describe the use of straight-through and crossover cables for connecting MDI and MDI-X devices.

MDI-X to MDI Cable Connections

The Ethernet Routing Switch 3510-24T features Auto-MDI/MDI-X detection. With auto-negotiation enabled, you can use straight Category 5 cables for MDI to MDI-X connections.

Auto-polarity

The Ethernet Routing Switch 3510-24T features auto-polarity. With autonegotiation enabled, auto-polarity automatically reverses the polarity of a pair of pins from positive to negative or negative to positive. This corrects the polarity of the received data if the port detects that the polarity of the data has been reversed due to a wiring error.

DB-9 (RS-232-D) Console/Comm Port Connector

The DB-9 Console/Comm Port connector (Figure 129) is configured as a data communications equipment (DCE) connector. The DSR and CTS signal outputs are always asserted; the CD, DTR, RTS, and RI signal inputs are not used. This configuration enables a management station (a PC or console terminal) to connect directly to the switch using a straight-through cable.

Figure 129 DB-9 Console port connector

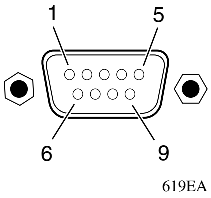

Table 79 lists the DB-9 Console port connector pin assignments.

Table 79 DB-9 Console port connector pin assignments

Pin	Signal	Description
1	CD	Not used
2	TXD	Transmit data (output)
3	RXD	Receive data (input)
4	DSR	
5	GND	Signal ground
6	DSR	Not used
7	CTS	
8	RTS	Not used
9	RI	Not used
Shell		Chassis ground

Appendix D
Default Settings

Table 80 lists the factory default settings for the Ethernet Routing Switch 3510-24T according to the console interface (CI) screens and fields for the settings.

Table 80 Factory default settings

Field	Default setting	Appears in this CI screen
Unit	1	"IP Configuration/Setup Screen" on page 148
BootP Request Mode	BootP When Needed	
In-Band Subnet Mask	0.0.0.0 (no subnet mask assigned)	
Default Gateway	0.0.0.0 (no IP address assigned)	
Read-Only Community String	public	"SNMP Configuration Screen" on page 154
Read-Write Community String	private	
Trap IP Address	0.0.0.0 (no IP address assigned)	
Community String	Zero-length string	
Authentication Trap	Enabled	
Link Up/Down Trap	Enabled	
sysContact	Zero-length string	"System Characteristics Screen" on page 156
sysName	Zero-length string	
sysLocation	Zero-length string	

Table 80 Factory default settings (continued)

Field	**Default setting**	**Appears in this CI screen**
Aging Time Find an Address Port Mirroring Address A:	300 seconds 00-00-00-00-00-00 (no MAC address assigned) 00-00-00-00-00-00 (no MAC address assigned)	"MAC Address Table Screen" on page 161
Port Mirroring Address B:	00-00-00-00-00-00 (no MAC address assigned)	
MAC Address Security	Disabled	"MAC Address Security Configuration Menu Screen" on page 162
MAC Address Security SNMP-Locked	Disabled	
Partition Port on Intrusion Detected:	Disabled	
DA Filtering on Intrusion Detected:	Disabled	
Generate SNMP Trap on Intrusion	Disabled	
Clear by Ports	NONE	
Learn by Ports	NONE	
Current Learning Mode	Disabled	
Trunk	blank field	"MAC Address Security Port Configuration Screen" on page 168
Security	Disabled	
Port List	blank field	"MAC Address Security Port Lists Screens" on page 170
Find an Address	blank field	"MAC Address Security Table Screens" on page 171
MAC Address	- - - - - (no address assigned)	
Allowed Source	- (blank field)	
Create VLAN	1	"VLAN Configuration Menu Screen" on page 177
Delete VLAN	blank field	
VLAN Name	VLAN # (*VLAN number*)	
Management VLAN	Yes, VLAN #1	

Table 80 Factory default settings (continued)

Field	Default setting	Appears in this CI screen
VLAN Type	Port-based	
Protocol ID (PID)	None	
User-Defined PID	0x0000	
VLAN State	Inactive	
Port Membership	All ports assigned as members of VLAN 1	
Port	1	"VLAN Port Configuration Screen" on page 184
Filter Untagged Frames	No	
Filter Unregistered Frames	No	
Port Name	Port 1	
PVID	1	
Port Priority	0	
Tagging	Untag All	
AutoPVID	Enabled	
Port	1	
PVID	1 (read only)	
Port Name	Port 1 (read only)	
Status	Enabled (for all ports)	"Port Configuration Screen" on page 188
Autonegotiation	Enabled (for all ports)	
Speed/Duplex	(Refer to Autonegotiation.)	
Trunk	1 to 6 (depending on configuration status)	"MultiLink Trunk Configuration Menu screen" on page 193
Trunk Members (Port)	Blank field	
STP Learning	Normal	
Trunk Mode	Basic	
Trunk Status	Disabled	
Trunk Name	Trunk #1 to Trunk #6	
Traffic Type	Rx and Tx	"MultiLink Trunk Utilization Screen" on page 197
Port	Blank Field	

Table 80 Factory default settings (continued)

Field	Default setting	Appears in this CI screen
Monitoring Mode	Disabled	"Port Mirroring Configuration Screen" on page 199
Monitor Port	Zero-length string	
Port X	Zero-length string	
Port Y	Zero-length string	
Address A	00-00-00-00-00-00 (no MAC address assigned)	
Address B	00-00-00-00-00-00 (no MAC address assigned)	
Packet Type	Both	"Rate Limiting Configuration Screen" on page 202
Limit	None	
VLAN	1	"IGMP Configuration Menu Screen" on page 205
Snooping	Disabled	
Proxy	Disabled	
Robust Value	2	
Query Time	125 seconds	
Set Router Ports	Version 1	
Static Router Ports	- (for all ports)	
Port	1	
Console Port Speed	9600 Baud	"Console/Comm Port Configuration Screen" on page 212
Console Switch Password	None	
Telnet/Web Switch Password	None	
Console Read-Write Switch Password	secure	
Create STP Group	1	"Spanning Tree Group Configuration Screen" on page 220
Delete STP Group	Blank Field	
Bridge Priority	8000	
Bridge Hello Time	2 seconds	
Bridge Maximum Age Time	20 seconds	

Table 80 Factory default settings (continued)

Field	**Default setting**	**Appears in this CI screen**
Bridge Forward Delay	15 seconds	
Add VLAN Membership	1	
Delete VLAN Membership	Blank Field	
Tagged BPDU on tagged port	• STP Group 1—No • Other STP Groups—Yes	
STP Group State	• STP Group 1—Active • Other STP Groups—InActive	
VID used for tagged BPDU	4001-4008 for STGs 1-8, respectively	
STP Group	1	"Spanning Tree Port Configuration Screen" on page 224
Participation	Normal Learning	
Priority	128	
Path Cost	1	
State	Forwarding	
STP Group	1	"Spanning Tree Switch Settings Screen" on page 228
STP Group	1	"Spanning Tree VLAN Membership Screen" on page 231
TELNET Access	Enabled	"TELNET/SNMP/Web Access Configuration Screen" on page 232
Login Timeout	1 minute	
Login Retries	3	
Inactivity Timeout	15 minutes	
Event Logging	All	
Allowed Source IP Address (10 user-configurable fields)	First field: 0.0.0.0 (no IP address assigned)	
	Remaining nine fields: 255.255.255.255 (any address is allowed)	
Allowed Source Mask (10 user-configurable fields)	First field: 0.0.0.0 (no IP address assigned)	

Table 80 Factory default settings (continued)

Field	Default setting	Appears in this CI screen
	Remaining nine fields: 255.255.255.255 (any address is allowed)	
Image Filename	Zero-length string	"Software Download Screen" on page 235
TFTP Server IP Address	0.0.0.0 (no IP address assigned)	
Start TFTP Load of New Image	No	
Configuration Image Filename	Zero-length string	"Configuration File Download/Upload Screen" on page 240
TFTP Server IP Address	0.0.0.0 (no IP address assigned)	
Copy Configuration Image to Server	No	
Retrieve Configuration Image from Server	No	
ASCII Configuration Filename	Zero-length string	"ASCII Configuration File Download Screen" on page 242
TFTP Server IP Address	0.0.0.0 (no IP address assigned)	
Retrieve Configuration file from Server	No	
Last Manual Configuration Status	Passed	
Last Auto Configuration Status	Passed	
Auto Configuration on Reset	Disabled	

Appendix E
Sample BootP Configuration File

This appendix provides a sample BootP configuration file. The BootP server searches for this file, called bootptab (or BOOTPTAB.TXT, depending on your operating system), which contains the site-specific information (including IP addresses) needed to perform the software download and configuration. You can modify this sample BootP configuration file or create one of your own.

A sample BootP configuration file follows:

```
# The following is a sample of a BootP configuration file that was extracted
# from a Nortel EZ LAN network management application.  Note that other
BootP daemons can use a configuration file with a different format.
#
# Before using your switch BootP facility, you must customize your BootP
# configuration file with the appropriate data.
#
# Blank lines and lines beginning with '#' are ignored.
#
# Legend:
#
#       first field -- hostname
#                ht -- hardware type
#                ha -- host hardware address
#                tc -- template host (points to similar host entry)
#                ip -- host IP address
#                hd -- bootfile home directory
#                bf -- bootfile
# EZ             dt -- device type
# EZ             fv -- firmware version
# EZ             av -- agent version
#                cs – TFTP server address for ASCII config file (optional)
#
# Fields are separated with a pipe (|) symbol. Forward slashes (/) are
# required to indicate that an entry is continued to the next line.
#
```

```
# Caution
#
#     Omitting a Forward slash (/) when the entry is continued to the next
#     line, can cause the interruption of the booting process or the
#     incorrect image file to download.  Always include forward slashes
#     where needed.
#
# Important Note:
#
#     If a leading zero (0) is used in the IP address it is calculated as an
#     octal number.  If the leading character is "x" (upper or lower case),
#     it is calculated as a hexadecimal number. For example, if an IP address
#     with a base 10 number of 45 is written as .045 in the BOOTPTAB.TXT file,
#     the Bootp protocol assigns .037 to the client.
#
# Global entries are defined that specify the parameters used by every device.
# Note that hardware type (ht) is specified first in the global entry.
#
# The following global entry is defined for an Ethernet device. Note that this
# is where a client's subnet mask (sm) and default gateway (gw) are defined.
#
global1|/
      |ht=ethernet|/
      |hd=c:\opt\images|/
      |sm=255.255.255.0|/
      |gw=192.0.1.0|
#
# The following sample entry describes a BootP client:

bay1|ht=ethernet|ha=0060fd000000|ip=192.0.0.1|hd=c:\ezlan\images|bf=bs5510.txt

# Where:
#    host name:                    bay1
#    hardware type:                Ethernet
#    MAC address:                  00-60-FD-00-00-00
#    IP address:                   192.0.0.0
#    home directory of boot file:  c:\ezlan\images
#    ASCII config file:            bs5510.txt
# When ASCII configuration download is configured to perform auto configuration
# on reset using BootP, the filename must be specified using the ‘bf’ keyword.
# If the ASCII configuration file is not resident on the BootP server, the
# server address can be specified using the ‘cs’ keyword.
```

Index