




SUPERSTACK™ II SWITCH 2200 ADMINISTRATION CONSOLE USER GUIDE



Part No. 801-00310-000
Published September 1996
Revision 01

3Com Corporation ■ 5400 Bayfront Plaza ■ Santa Clara, California ■ 95052-8145

© 3Com Corporation, 1996. All rights reserved. No part of this documentation may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without permission from 3Com Corporation.

3Com Corporation reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of 3Com Corporation to provide notification of such revision or change.

3Com Corporation provides this documentation without warranty of any kind, either implied or expressed, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. 3Com may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

UNITED STATES GOVERNMENT LEGENDS:

If you are a United States government agency, then this documentation and the software described herein are provided to you subject to the following restricted rights:

For units of the Department of Defense:

Restricted Rights Legend: Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) for restricted Rights in Technical Data and Computer Software clause at 48 C.F.R. 52.227-7013. 3Com Corporation, 5400 Bayfront Plaza, Santa Clara, California 95052-8145.

For civilian agencies:

Restricted Rights Legend: Use, reproduction, or disclosure is subject to restrictions set forth in subparagraph (a) through (d) of the Commercial Computer Software - Restricted Rights Clause at 48 C.F.R. 52.227-19 and the limitations set forth in 3Com Corporation's standard commercial agreement for the software. Unpublished rights reserved under the copyright laws of the United States.

If there is any software on removable media described in this documentation, it is furnished under a license agreement included with the product as a separate document, in the hardcopy documentation, or on the removable media in a directory file named LICENSE.TXT. If you are unable to locate a copy, please contact 3Com and a copy will be provided to you.

Unless otherwise indicated, 3Com registered trademarks are registered in the United States and may or may not be registered in other countries.

3Com, LANplex, LinkBuilder, NETBuilder, NETBuilder II, ViewBuilder, EtherDisk, EtherLink, EtherLink II, and Transcend are registered trademarks of 3Com Corporation. 3TECH, FDDLLink, SmartAgent, and Star-Tek are trademarks of 3Com Corporation. 3ComFacts is a service mark of 3Com Corporation.

IBM and Netview AIX are registered trademarks of International Business Machines Corporation. Apple, AppleTalk, and Macintosh are trademarks of Apple Computer, Inc. CompuServe is a registered trademark of CompuServe, Inc. MS-DOS and Windows are registered trademarks of Microsoft Corporation. OpenView is a registered trademark of Hewlett-Packard Co. Sniffer is a registered trademark of Network General Corp. SunNet Manager, SunOS, and OpenWindows are trademarks of Sun Microsystems, Inc. UNIX is a registered trademark of Novell Inc.

Other brand and product names may be registered trademarks or trademarks of their respective holders.

Guide written, edited, and illustrated by Beth Britt, Patricia Crawford, Lynne Gelfand, Michael Jenness, Patricia L. Johnson, Michael Taillon, and Iain Young. Edited by Bonnie Jo Collins.

CONTENTS

ABOUT THIS GUIDE

- Introduction 1
- How to Use This Guide 2
- Conventions 3
- Switch 2200 Documentation 4
- Documentation Comments 5

PART I INTRODUCTION

1 SUPERSTACK™ II SWITCH 2200 ADMINISTRATION OVERVIEW

- About Switch 2200 Administration 1-1
- Configuration Tasks 1-1

2 HOW TO USE THE ADMINISTRATION CONSOLE

- Initial User Access 2-1
- Levels of User Access 2-1
 - Administer Access Example 2-2
 - Write Access Example 2-2
 - Read Access Example 2-3
- Using Menus to Perform Tasks 2-3
 - Administration Console Menu Structure 2-4
 - System Menu 2-4
 - Ethernet Menu 2-4
 - FDDI Menu 2-5
 - Bridge Menu 2-5
 - IP Menu 2-6
 - SNMP Menu 2-7
 - Analyzer Menu 2-7
 - Selecting Menu Options 2-8
 - Entering Values 2-9
 - Getting Out 2-9

Administration Console Interface Parameters	2-10
Adjusting the Screen Height	2-10
Disabling the Reboot and Abort Keys	2-11
Remote Access Parameters	2-11
Preventing Disconnections	2-11
Enabling Timeout of Remote Sessions	2-12
Setting Timeout Interval for Remote Sessions	2-13
Running Scripts of Administration Console Tasks	2-13
Getting Help in the Administration Console	2-16
Online Help	2-16
Viewing More Levels of Menu Options	2-16
Exiting the Administration Console	2-17

PART II SYSTEM-LEVEL FUNCTIONS

3 CONFIGURING MANAGEMENT ACCESS TO THE SYSTEM

About Management Access	3-1
Using a Serial Connection	3-1
Using an IP Interface	3-1
In-band or Out-of-band?	3-2
Setting Up the Console Serial Port	3-2
Setting Up an IP Interface for Management	3-3
General Setup Process	3-3
Administering Interfaces	3-3
Displaying Interfaces	3-4
Defining an Interface	3-5
Modifying an Interface	3-6
Removing an Interface	3-7
Administering Routes	3-7
Displaying the Routing Table	3-8
Defining a Static Route	3-9
Removing a Route	3-9
Flushing a Route	3-10
Setting the Default Route	3-10
Removing the Default Route	3-10
Administering the ARP Cache	3-11
Displaying the ARP Cache	3-11
Removing an ARP Cache Entry	3-11
Flushing ARP Cache Entries	3-12
Setting the RIP Mode	3-12
Pinging an IP Station	3-12
Displaying IP Statistics	3-14

- Setting Up SNMP on Your System 3-15
 - Displaying SNMP Settings 3-15
 - Configuring Community Strings 3-15
 - Administering SNMP Trap Reporting 3-16
 - Displaying Trap Information 3-16
 - Configuring Trap Reporting 3-17
 - Removing Trap Destinations 3-18
 - Flushing Trap Destinations 3-19
 - Setting Up SMT Event Proxying 3-19

4 ADMINISTERING YOUR SYSTEM ENVIRONMENT

- Displaying the System Configuration 4-1
- Setting Passwords 4-2
- Setting the System Name 4-3
- Changing the Date and Time 4-3
- Rebooting the System 4-4

5 BASELINING STATISTICS

- About Setting Baselines 5-1
- Displaying the Current Baseline 5-1
- Setting Baselines 5-2
- Enabling or Disabling Baselines 5-2

6 SAVING, RESTORING, AND RESETTNG NONVOLATILE DATA

- About Working with Nonvolatile Data 6-1
- Saving NV Data 6-2
- Restoring NV Data 6-3
- Examining a Saved NV Data File 6-5
- Resetting NV Data to Defaults 6-6

PART III ETHERNET AND FDDI PARAMETERS

7 ADMINISTERING ETHERNET PORTS

- Displaying Ethernet Port Information 7-1
- Labeling a Port 7-8
- Setting the Port State 7-8

8 ADMINISTERING FDDI RESOURCES

- Administering FDDI Stations 8-1
 - Displaying Station Information 8-2
 - Setting the Connection Policies 8-3
 - Setting Neighbor Notification Timer 8-5
 - Enabling and Disabling Status Reporting 8-5
- Administering FDDI Paths 8-6
 - Displaying Path Information 8-6
 - Setting tvxLowerBound 8-7
 - Setting tmaxLowerBound 8-8
 - Setting maxT-Req 8-9
- Administering FDDI MACs 8-9
 - Displaying MAC Information 8-10
 - Setting the Frame Error Threshold 8-16
 - Setting the Not Copied Threshold 8-17
 - Enabling and Disabling LLC Service 8-18
 - Setting the MAC Paths 8-18
- Administering FDDI Ports 8-19
 - Displaying Port Information 8-19
 - Setting IerAlarm 8-20
 - Setting IerCutoff 8-21
 - Setting Port Labels 8-22
 - Setting the Port Paths 8-23

9 SETTING UP THE SYSTEM FOR ROVING ANALYSIS

- About Roving Analysis 9-1
- Displaying the Roving Analysis Configuration 9-2
- Adding an Analyzer Port 9-3
- Removing an Analyzer Port 9-4
- Starting Port Monitoring 9-5
- Stopping Port Monitoring 9-6

PART IV BRIDGING PARAMETERS

10 ADMINISTERING THE BRIDGE

- Displaying Bridge Information 10-1
- Enabling and Disabling IP Fragmentation 10-5
- Enabling and Disabling IPX Snap Translation 10-5
- Setting the Address Threshold 10-6
- Setting the Aging Time 10-6

- Administering STP Bridge Parameters 10-7
 - Enabling and Disabling STP on a Bridge 10-7
 - Setting the Bridge Priority 10-7
 - Setting the Bridge Maximum Age 10-8
 - Setting the Bridge Hello Time 10-9
 - Setting the Bridge Forward Delay 10-9
 - Setting the STP Group Address 10-10
-

11 ADMINISTERING BRIDGE PORTS

- Displaying Bridge Port Information 11-1
 - Setting the Multicast Limit 11-7
 - Administering STP Bridge Port Parameters 11-8
 - Enabling and Disabling STP on a Port 11-8
 - Setting the Port Path Cost 11-9
 - Setting the Port Priority 11-10
 - Administering Port Addresses 11-11
 - Listing Addresses 11-11
 - Adding New Addresses 11-12
 - Removing Addresses 11-12
 - Flushing All Addresses 11-13
 - Flushing Dynamic Addresses 11-13
 - Freezing Dynamic Addresses 11-13
-

12 CREATING AND USING PACKET FILTERS

- About Packet Filtering 12-1
- Listing Packet Filters 12-2
- Displaying Packet Filters 12-3
- Creating Packet Filters 12-3
 - Concepts for Writing a Filter 12-4
 - How the Packet Filter Language Works 12-4
 - Basic Elements of a Packet Filter 12-6
 - Implementing Sequential Tests in a Packet Filter 12-8
 - Preprocessed and Run-time Storage 12-9
 - Procedure for Writing a Filter 12-10
 - Examples of Creating Filters 12-11
 - Filtering Problem 12-11
 - Packet Filter Solution 12-12
 - Tools for Writing a Filter 12-17
 - Using the Built-in Line Editor 12-17
 - Using an External Text Editor 12-20
- Deleting Packet Filters 12-20
- Editing, Checking and Saving Packet Filters 12-20

Loading Packet Filters 12-22
Assigning Packet Filters to Ports 12-22
Unassigning Packet Filters from Ports 12-24

13 CONFIGURING ADDRESS AND PORT GROUPS TO USE IN PACKET FILTERS

Using Groups in Packet Filters 13-1
Listing Groups 13-2
Displaying Groups 13-3
Creating New Groups 13-4
Deleting Groups 13-6
Adding Addresses and Ports to Groups 13-7
Removing Addresses or Ports from a Group 13-9
Loading Groups 13-11

PART APPENDIXES

A PACKET FILTER OPCODES, EXAMPLES, AND SYNTAX ERRORS

Opcodes A-1
Packet Filter Examples A-9
 Destination Address Filter A-9
 Source Address Filter A-9
 Length Filter A-9
 Type Filter A-10
 Ethernet Type IPX and Multicast Filter A-10
 Multiple Destination Address Filter A-10
 Source Address and Type Filter A-11
 Accept XNS or IP Filter A-11
 XNS Routing Filter A-11
 Address Group Filter A-12
 Port Group Filter A-12
Common Syntax Errors A-13

B TECHNICAL SUPPORT

- Online Technical Services B-1
 - 3Com Bulletin Board Service B-1
 - Access by Modem B-1
 - Access by ISDN B-2
 - World Wide Web Site B-2
 - 3ComForum on CompuServe® B-2
 - 3ComFactsSM Automated Fax Service B-3
- Support from Your Network Supplier B-3
- Support from 3Com B-4
- Returning Products for Repair B-4

INDEX

ABOUT THIS GUIDE

Introduction

The *SuperStack™ II Switch 2200 Administration Console User Guide* provides all the information you need to configure and manage your Switch 2200 once it is installed and the system is attached to the network. Prior to using this guide, you should have already installed and set up your system using the *SuperStack™ II Switch 2200 Getting Started* guide.

Audience description

This guide is intended for the system or network administrator who is responsible for configuring, using, and managing the Switch 2200 system. It assumes a working knowledge of local area network (LAN) operations and a familiarity with communications protocols that are used on interconnected LANs.



If the information in the Release Notes shipped with this product differs from the information in this guide, follow the Release Notes.

How to Use This Guide

This guide is organized by types of tasks you may need to perform on the Switch 2200. The parts of the guide are described in Table 1.

Table 1 Description of Guide Parts

Part	Contents
I: Introduction	<p>Introducing Switch 2200 administration</p> <p>Learning about the various system configurations and the quick commands to perform them</p> <p>Learning about password access to the Console</p> <p>Learning about the Administration Console menu structure and maneuvering within the Console (using commands and moving between menus)</p> <p>Setting interface parameters (screen height and control keys)</p> <p>Running scripts of Console tasks</p> <p>Getting help</p>
II: System-Level Functions	<p>Setting up the system for management access (through serial ports or using IP and setting up SNMP)</p> <p>Configuring SNMP community strings</p> <p>Setting up trap reporting</p> <p>Configuring system parameters, such as name, date/time, and passwords</p> <p>Baselining statistics</p> <p>Saving, restoring, and resetting nonvolatile data</p>
III: Ethernet and FDDI Parameters	<p>Displaying statistics for and labeling Ethernet ports</p> <p>Displaying statistics for and configuring various parameters for FDDI stations, ports, MACs, and paths</p> <p>Setting up the system to monitor Ethernet port activity using roving analysis</p>

(continued)

Table 1 Description of Guide Parts (continued)

Part	Contents
IV: Bridging	Configuring bridge and bridge port parameters Administering the Spanning Tree Protocol bridge and bridge port parameters Displaying and configuring bridge port addresses Creating and using packet filters Creating address groups and port groups and using them as filtering criteria
V: Appendixes	Additional information about packet filters: opcode descriptions, examples, and error messages Getting Technical Support Returning products for repair

Conventions

Table 2 and Table 3 list icon and text conventions that are used throughout this guide.

Table 2 Notice Icons




Icon	Type	Description
	Information Note	Information notes call attention to important features or instructions.
	Caution	Cautions contain directions that you must follow to avoid immediate system damage or loss of data.
	Warning	Warnings contain directions that you must follow for your personal safety. Follow all instructions carefully.

Table 3 Text Conventions

Convention	Description
"Enter"	"Enter" means type something, then press the [Return] or [Enter] key.
"Syntax" vs. "Command"	<p>"Syntax" indicates that the general command syntax form is provided. You must evaluate the syntax and supply the appropriate value; for example:</p> <p>Set the date by using the following syntax:</p> <pre>mm/DD/yy hh:mm:ss xm</pre> <p>"Command" indicates that all variables in the command syntax form have been supplied and you can enter the command as shown in text; for example:</p> <p>To update the system software, enter the following command:</p> <pre>system software Update</pre>
Text represented as screen display	<p>This typeface represents text that appears on your terminal screen; for example:</p> <pre>NetLogin:</pre>
Text represented as commands	<p>This typeface represents commands that you enter; for example:</p> <pre>bridge port stpState</pre>
<i>Italic</i>	<i>Italic</i> is used to denote emphasis and buttons.
Keys	<p>When specific keys are referred to in the text, they are called out by their labels, such as "the Return key" or "the Escape key," or they may be shown as [Return] or [Esc].</p> <p>If two or more keys are to be pressed simultaneously, the keys are linked with a plus sign (+), for example:</p> <p>Press [Ctrl]+[Alt]+[Del].</p>

Switch 2200 Documentation

The following documents comprise the Switch 2200 documentation set. If you want to order a document that you do not have or order additional documents, contact your sales representative for assistance.

- *SuperStack™ II Switch 2200 Unpacking Instructions*

Describes how to unpack your Switch 2200. It also provides you with an inventory list of all the items that came with your system. (Shipped with system/Part No. 801-00312-000)

- *SuperStack™ II Switch 2200 Software Release Notes*

Provides information about the software release, including new features and bug fixes. It also provides information about any changes to the Switch 2200 documentation. (Shipped with system)

- *SuperStack™ II Switch 2200 Getting Started*
Describes all the procedures necessary for planning your configuration and for installing, cabling, powering up, and troubleshooting your Switch 2200 system. (Shipped with system/Part No. 801-00309-000)
- *SuperStack™ II Switch 2200 Operation Guide*
Provides information to help you understand system management and administration, FDDI technology, and bridging. It also describes how these concepts are implemented in the Switch 2200 system. (Shipped with system/Part No. 801-00311-000)
- *SuperStack™ II Switch 2200 Administration Console User Guide (this guide)*
Provides information about using the Administration Console to configure and manage your Switch 2200 system. (Shipped with system/Part No. 801-00310-000)
- *Command Quick Reference for the SuperStack™ II Switch 2200 Administration Console*
Contains all of the Administration Console intelligent switching commands for the Switch 2200 system. (Folded card; shipped with system/Part No. 801-00314-000)


Documentation Comments

Your suggestions are very important to us: To help make Switch documentation more useful to you, please email comments about this guide to 3Com at: **sdtechpubs_comments@3Mail.3Com.com**

Please include the following information when commenting:

- Document title
- Document part number (on back cover of document)
- Page number (if appropriate)


Example: *SuperStack™ II Switch 2200 Operation Guide*
Part No. 801-00311-000
Page 2-5 (chapter 2, page 5)




INTRODUCTION

Chapter 1 Overview of SuperStack™ II Switch 2200 Administration

Chapter 2 How to Use the Administration Console





SUPERSTACK™ II SWITCH 2200 ADMINISTRATION OVERVIEW

This chapter introduces you to SuperStack™ II Switch 2200 administration and briefly describes the system parameters that you can configure.

About Switch 2200 Administration

The Switch 2200 software is installed at the factory in flash memory on the system processor. Because this software boots from flash memory automatically when you power on your system, the system is immediately ready for use in your network. However, you might need to configure certain parameters for the system to operate effectively in your networking environment. Additionally, when managing your Switch 2200, you might want to view important MAC, port, bridge, and IP statistics. The Switch 2200 Administration Console allows you to configure your system and display these important statistics. For more complete network management, you can use an external application, such as 3Com's Transcend® Enterprise Manager.

Configuration Tasks

This section uses tables to summarize the tasks and quick commands for the SuperStack™ II Switch 2200 Administration Console.

- General System Commands (Table 1-1)
- System Management Setup Commands (Table 1-2)
- Bridging Commands (Table 1-3)
- Ethernet Commands (Table 1-4)
- FDDI Commands (Table 1-5)

These tables, which are repeated on the *Command Quick Reference* card that comes with your system, provide a brief description of most tasks, along with the Administration Console command to access the task quickly. They also tell you where to look in the documentation for additional information.

Table 1-1 General System Commands

Task	Quick Command	For Details, See. . .
Run a script of commands to set up a system Write a script of Console commands with the values you assign so that you can quickly configure one or more systems. You can run the same script on a number of systems to ensure consistent setup.	script	page 2-13
Display the system configuration Display software and hardware revisions and certain warning messages.	system display	page 4-1
Install software into flash memory Update your system software. Software is initially installed at the factory.	system softwareUpdate	<i>SuperStack™ II Switch 2200 Software Installation and Release Notes</i>
Display, set, enable, or disable a baseline for statistics Establish and use baselines for Ethernet, FDDI, and bridging statistics to evaluate recent activity in your system and on your network.	system baseline	page 5-2
Configure timeout for remote sessions Configure the system to disconnect remote sessions after a specified time interval.	system telnet	page 2-12
Control access to the Console Set passwords for levels of access (read, write, administer) and prohibit remote access during your session by locking the Console.	system password system consoleLock	page 4-2 page 3-20
Name the system Assign the system a unique name for management purposes. For example, you might name a system based on its location: <i>Switch2200-Floor2</i> .	system name	page 4-3
Set the system date and time Ensure that messages are accurately logged. The internal clock is set at the factory; change it for your time zone.	system time	page 4-3
Set screen height Adjust the console screen height for your terminal.	system screenHeight	page 2-10
Enable the [Control] keys when working in the Console Enable quick keys for the reboot (Ctrl+X) and abort (Ctrl+C) functions.	system ctlKeys	page 2-11

(continued)

Table 1-1 General System Commands (continued)

Task	Quick Command	For Details, See. . .
Save, restore, or reset nonvolatile data in the system Provide a backup for nonvolatile data, restore nonvolatile data to the system, or reset nonvolatile data to defaults.	<code>system nvData</code>	page 6-2
Reboot the system Restart the system. Disconnects rlogin and telnet sessions.	<code>system reboot</code>	page 4-4

Table 1-2 System Management Setup Commands

Task	Quick Command	For Details, See. . .
Configure the Console port baud rate Change the factory default baud rate of the Console port, which allows you to connect a VT or tty type of terminal or terminal emulator to the system using a null modem cable.	<code>system consoleSpeed</code>	page 3-2
Configure an IP address using an IP interface Communicate with the system using SNMP, rlogin, or telnet.	<code>ip interface display</code> <code>ip interface define</code> <code>ip interface modify</code> <code>ip interface remove</code>	page 3-5
Define static routes Access a menu from which you can display, define, remove, and flush static routes for transmitting traffic through the system. Static routes override routes learned through RIP.	<code>ip route</code> <code>ip route default</code>	page 3-9
Administer the ARP cache Display, remove, and flush the ARP cache (a table of known IP addresses and their corresponding MAC addresses).	<code>ip arp display</code> <code>ip arp remove</code> <code>ip arp flush</code>	page 3-11
Set RIP's operational mode Define how Routing Information Protocol (RIP) messages are processed.	<code>ip rip</code>	page 3-12
Ping an IP station or the system Find out if the system can reach an IP station or check that the system is on the network.	<code>ip ping</code>	page 3-12
Display IP statistics Display datagram statistics and current RIP operational mode.	<code>ip statistics</code>	page 3-14

(continued)

Table 1-2 System Management Setup Commands (continued)

Task	Quick Command	For Details, See. . .
Configure SNMP management Display current SNMP configurations and specify the type of authorization for SNMP management.	snmp display snmp community	page 3-15
Configure SNMP trap reporting Display SNMP trap reporting information, add or modify trap reporting destination configurations, remove trap destinations, flush all SNMP trap reporting destinations, and set up SMT event proxying.	snmp trap display snmp trap addModify snmp trap remove snmp trap flush snmp trap smtProxyTraps	page 3-16

Table 1-3 Bridging Commands

Task	Quick Command	For Details, See. . .
Display bridge information Display information about the bridge, such as statistics, bridge configurations, and spanning tree configurations.	bridge display	page 10-1
Enable or disable IP fragmentation Enable or disable the fragmenting of large FDDI packets to allow FDDI and Ethernet stations to communicate using IP.	bridge ipFragmentation	page 10-5
Enable or disable IPX snap translation Enable or disable the translation of 802.3_RAW IPX packets to FDDI_SNAP packets (when going from Ethernet to FDDI), and vice versa (when going from FDDI to Ethernet). The default is disabled.	bridge ipxSnapTranslation	page 10-5
Set the bridge address threshold Specify the reporting threshold for the total number of Ethernet addresses known to the bridge. The SNMP trap <i>addressThresholdEvent</i> is generated when the threshold is reached.	bridge addressThreshold	page 10-6
Set the bridge address aging timer Specify how often dynamically learned addresses are aged by the bridge port. Appropriately configured aging prevents packet flooding.	bridge agingTime	page 10-6

(continued)

Table 1-3 Bridging Commands (continued)

Task	Quick Command	For Details, See. . .
<p>Configure Spanning Tree Protocol (STP) parameters for a bridge</p> <p>Enable or disable STP and set the bridge priority, the maximum age of stored configuration message information, the period between the generation of messages by a root bridge, the amount of time a bridge spends in the listening and learning states, and the group address.</p>	<pre>bridge stpState bridge stpPriority bridge stpMaxAge bridge stpHelloTime bridge stpForwardDelay bridge stpGroupAddress</pre>	<p>page 10-7 to page 10-10</p>
<p>Display bridge port information</p> <p>Display information about the bridge port, including STP configurations, in a summarized or detailed format.</p>	<pre>bridge port summary bridge port detail</pre>	<p>page 11-1</p>
<p>Configure Spanning Tree Protocol (STP) parameters for a bridge port</p> <p>Enable or disable STP on a bridge port, and set the bridge port path cost and port priority.</p>	<pre>bridge port stpState bridge port stpCost bridge port stpPriority</pre>	<p>page 11-8 page 11-9 page 11-10</p>
<p>Set the multicast packet firewall threshold</p> <p>Suppress multicast storms and limit the rate at which multicast packets are propagated by the system.</p>	<pre>bridge port multicastLimit</pre>	<p>page 11-7</p>
<p>Administer bridge port addresses</p> <p>Administer the MAC address of stations connected to Ethernet and FDDI ports. This command accesses a menu from which you can list, add, remove, flush, and freeze bridge port addresses.</p>	<pre>bridge port address</pre>	<p>page 11-11</p>
<p>Use packet filters to restrict which packets are forwarded through a bridge port</p> <p>Access a menu from which you can list packet filters, display a packet filter definition, create or edit a definition, load a definition onto the system, copy a definition, and assign or unassign a definition to a port.</p>	<pre>bridge packetFilter</pre>	<p>page 12-1 and following</p>
<p>Create address and port groups to use as filtering criteria</p> <p>Access a menu from which you can specify groups (either address groups or port groups) to use in a packet filter definition. From each menu, you can list, display, create, and delete groups. You can also add and remove address and ports to and from groups.</p>	<pre>bridge packetFilter addressGroup bridge packetFilter portGroup</pre>	<p>page 13-1 and following</p>

Table 1-4 Ethernet Commands

Task	Quick Command	For Details, See. . .
Display Ethernet port information Display label, status, and statistic information on Ethernet ports in a summarized or detailed format.	ethernet summary ethernet detail	page 7-1
Label an Ethernet port Assign a unique name to an Ethernet port. Useful for port identification when managing the system.	ethernet label	page 7-8
Set the Ethernet port state Enable or disable an Ethernet port, controlling whether the port sends and receives frames.	ethernet portState	page 7-8
Configure Ethernet ports to be monitored by a network analyzer Analyze data forwarded through Ethernet ports. With roving analysis, you set up one Ethernet port for a network analyzer attachment and set up another Ethernet port (local or remote) to be monitored. Data is copied and forwarded from the port being monitored to the network analyzer.	analyzer display analyzer add analyzer remove analyzer start analyzer stop	page 9-2 to page 9-6

Table 1-5 FDDI Commands

Task	Quick Command	For Details, See. . .
Display FDDI information Display information about the system's FDDI station, paths, MAC, and ports. MAC information is available in a summarized or detailed format.	fddi station display fddi path display fddi mac summary fddi mac detail fddi port display	page 8-2 page 8-6 page 8-18 page 8-19
Set FDDI station parameters Set parameters for connection policies, the neighbor notification timer, and status reporting.	fddi station connectPolicy fddi station tNotify fddi station statusReporting	page 8-3 and page 8-5
Set FDDI path parameters Set the minimum value for the TVX timer, the minimum value for the T-Max timer, and the maximum value for the T-Req timer.	fddi path tvxLowerBound fddi path tmaxLowerBound fddi path maxTreq	page 8-7 page 8-8 page 8-9
Set FDDI MAC parameters Set the parameters for the frame error threshold and the not copied threshold, enable or disable LLC service, and set MAC paths.	fddi mac frameErrorThreshold fddi mac notCopiedThreshold fddi mac llcService fddi mac path	page 8-16 page 8-17 page 8-18 page 8-18
Set FDDI port parameters Set the parameters for the link error rate alarm threshold and the link error rate cut-off threshold, and set port paths.	fddi port lerAlarm fddi port lerCutoff fddi port path	page 8-20 page 8-21 page 8-23
Label an FDDI port Assign a unique name to an FDDI port. Useful for port identification when managing the system.	fddi port label	page 8-22

2

HOW TO USE THE ADMINISTRATION CONSOLE

This chapter familiarizes you with user access levels of the Superstack™ II Switch 2200 Administration Console and explains how to:

- Move around within the menu hierarchy to perform tasks
- Set up the interface parameters
- Access online help
- Use scripts for performing Administration Console tasks
- Exit the Administration Console

Initial User Access

As the initial user, access the system at the *administer* level and press Return at the password prompt. The first time you access the Administration Console, the password is null. Subsequent access is described in this chapter.

Levels of User Access

The Administration Console supports three password levels, allowing the network administrator to provide different levels of access for a range of Switch 2200 users. These access levels are described in Table 2-1.

Table 2-1 Password Access Levels

Access Level	For Users Who Need to...	Allows Users to...
Administer	Perform system set-up and management tasks (usually a single network administrator)	Perform system-level administration (such as setting passwords, loading new software, and so on)
Write	Perform active network management	Configure network parameters (such as setting the aging time for a bridge)
Read	Only view system parameters	Access only "display" menu items (display, summary, detail)

Each time you access the Administration Console, the system prompts you for an access level and password, as shown here:

```
Select access level (read, write, administer):
Password:
```

The passwords are stored in nonvolatile (NV) memory. You must enter the password correctly before you are allowed to continue.

The following examples show how the top-level menu structure changes based on the level of access. For information about setting passwords, see page 4-2.

Administer Access Example If you have administer access, each menu contains all options. Here is the **system** menu for users with administer access:

```
Menu options: -----
display                - Display the system configuration
softwareUpdate         - Load a new revision of system software
baseline               - Administer a statistics baseline
consoleSpeed           - Set the console serial port baud rate
telnet                 - Administer telnet sessions
password               - Set the console passwords
name                   - Set the system name
time                   - Set the date and time
screenHeight           - Set the console screen height
consoleLock            - Allow/Disallow remote access to the console
ctlKeys                - Enable/Disable Ctl-X (reboot) and Ctl-C (abort)
nvData                 - Save, restore, or reset nonvolatile data
reboot                 - Reboot the system
```

```
Type 'q' to return to the previous menu or ? for help.
```

```
-----
Select a menu option (system):
```

Write Access Example If you have write access, the **system** menu contains a subset of the complete menu, focusing on the network, as shown here:

```
Menu options: -----
display                - Display the system configuration
baseline               - Administer statistics baseline
consoleSpeed           - Set the console serial port baud rate
name                   - Set the system
screenHeight           - Set the console screen height
```

```
Type 'q' to return to the previous menu or ? for help.
```

```
-----
Select a menu option (system):
```


Read Access Example If you have read access, the **system** menu contains only the display options shown here:

```

Menu options: -----
      display                - Display the system configuration
Only the display -----
option in the   baseline    - Administer statistics baseline
baseline menu
is available

Type 'q' to return to the previous menu or ? for help.
-----
Select a menu option (system):

```

Using Menus to Perform Tasks

When you access the Administration Console, the top-level menu appears. You use the Administration Console by selecting options from this menu and from others below it. Each menu option is accompanied by a brief description. Here is the **top-level** menu:

```

                                     Option Descriptions
                                     |
Menu options: -----
      system                    - Administer system-level functions
      ethernet                  - Administer Ethernet ports
      fddi                      - Administer FDDI resources
      bridge                    - Administer bridging
      ip                        - Administer IP
      snmp                      - Administer SNMP
      analyzer                  - Administer Roving Analysis
      script                    - Run a script of console commands
      logout                    - Logout of the Administration Console
Options
(These vary with -----
level of access.)

Type ? for help.
-----
Select a menu option:

```

Administration Console Menu Structure

The following sections show the menu paths for performing tasks from the top-level menu and provide a brief description of each top-level menu option. See “Selecting Menu Options” on page 2-8 for instructions on actually using the menu system.



The following menus display the options available for users with administer access.

System Menu

From the **system** menu, you can view the system configuration, set up your system for management, configure Administration Console interface parameters, work with nonvolatile data, and reboot the system. (See Figure 2-1.)

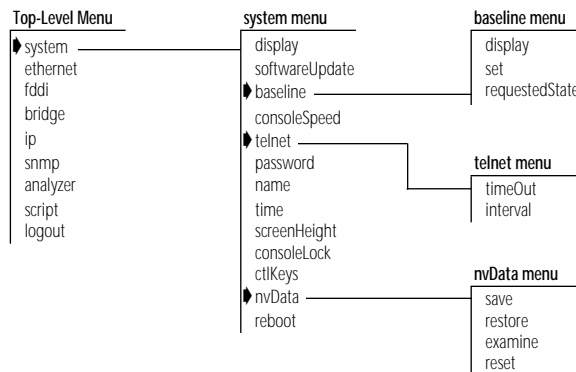


Figure 2-1 System-level Functions Menu Hierarchy for Administer Access

Ethernet Menu

From the **ethernet** menu, you can view information for and name Ethernet ports. (See Figure 2-2.) For example, to view all Ethernet port statistics, you enter **ethernet** at the top-level menu, and then **detail** at the ethernet menu.




Figure 2-2 Ethernet Menu Hierarchy for Administer Access

FDDI Menu

From the **fddi** menu, you can view information about and configure the FDDI station, paths, MAC, and ports. (See Figure 2-3.) For example, to enable the LLC service of the FDDI MAC, you enter **fddi** at the top-level menu, **mac** at the fddi menu, and then **llcService** at the mac menu.

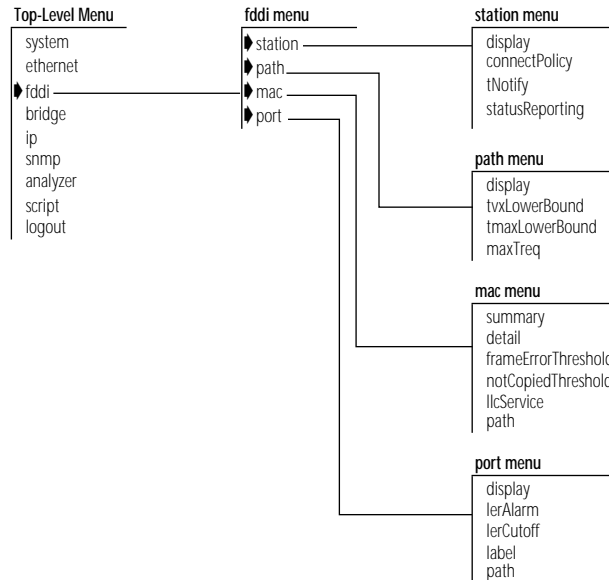


Figure 2-3 FDDI Menu Hierarchy for Administer Access

Bridge Menu

From the **bridge** menu, you can view information about and configure bridge-level parameters, including those for the Spanning Tree Protocol (STP). You can also configure the bridge at the port level and administer packet filters. (See Figure 2-4.) For example, to set the Spanning Tree state for a bridge port, you enter **bridge** at the top-level menu, **port** at the bridge menu, and **stpState** at the port menu.

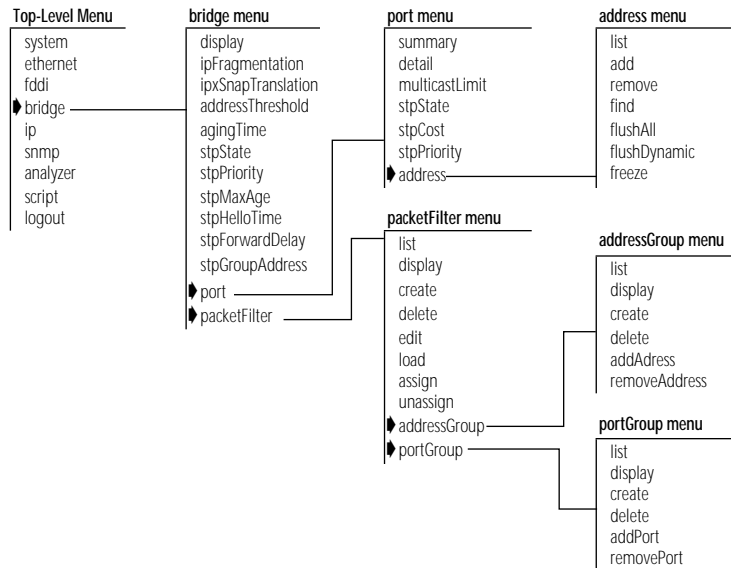


Figure 2-4 Bridging Menu Hierarchy for Administer Access

IP Menu

From the **ip** menu, you can view information about and configure Internet Protocol (IP) interfaces and routes. You can also administer the Address Resolution Protocol (ARP) and the Routing Information Protocol (RIP), and ping IP stations. (See Figure 2-5.) For example, to define a new IP interface, you enter **ip** at the top-level menu, **interface** at the ip menu, and then **define** at the interface menu.

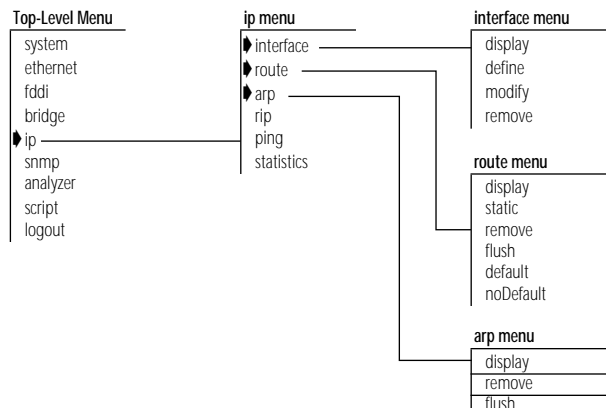


Figure 2-5 IP Menu Hierarchy for Administer Access

SNMP Menu

From the **snmp** menu, you can configure SNMP community strings and trap reporting. (See Figure 2-6.) For example, to flush all trap reporting destinations, you enter **snmp** at the top-level menu, **trap** at the snmp menu, and then **flush** at the trap menu.




Figure 2-6 SNMP Menu Hierarchy for Administer Access

Analyzer Menu

From the **analyzer** menu, you can selectively choose any Ethernet network segment attached to a Switch 2200 and monitor its activity using a network analyzer. (See Figure 2-7.) For example, to add analyzer ports, you enter **analyzer** at the top-level menu, and then **add** at the analyzer menu.




Figure 2-7 Analyzer Menu Hierarchy for Administer Access

Selecting Menu Options

You select a menu option at the selection prompt by entering its name (or enough of the name to uniquely identify it within the particular menu). For example, to access the **system** menu from the top-level menu, you enter:

```
Select a menu option: system
```

OR

```
Select a menu option: sy
```



Menu options are not case sensitive.

When you enter a menu option, you either go to the next menu in the hierarchy or you see information for the option you entered. The information is either a prompt or a screen display. If you enter the menu option incorrectly, you receive a prompt telling you that what you entered was not valid or was ambiguous. You must re-enter the command from the point at which it became incorrect. Expand a truncated command until it becomes unambiguous.

When a new menu appears, the selection prompt (with its choices in parentheses) changes to reflect your progression through the menus. For example, if you enter **system** at the top-level menu and then **baseline** at the system menu, the prompt changes at the next level:

```
Select a menu option (system/baseline):
```

Entering a command string

Once you are familiar with the menu structure, instead of working your way down the menu hierarchy to a task, you can enter a string of menu options at a selection prompt to go immediately to a task. For example, the command string for setting a baseline from the top-level menu looks like this:

```
Select a menu option: system baseline set
```

The most abbreviated version of the same command string is:

```
Select a menu option: sy b s
```

When you enter a command string, you move to the last menu level or option you entered in the command string, and information relevant to that command is displayed. It may be a menu, prompt, or screen display.

If you enter a command incorrectly, you receive a prompt telling you that what you entered was not valid or was ambiguous. You must re-enter the command from the point at which it became incorrect.

Entering Values

When you reach the level at which you perform a specific task, you are prompted for a value. The prompt usually shows all valid values (if applicable) and sometimes a suggested default value. The default might be the system default or the current user-defined value of that parameter.

The valid values are displayed in parentheses. The default value is in brackets. In this example, (disabled, enabled) are the valid values. [Enabled], shown in brackets, is the default:

```
Enter a new value (disabled,enabled) [enabled]:
```

Entering values in command strings

A command string can also contain the value of a command parameter. If you enter a value at the end of a command string, the task is completed, and you are returned to the previous menu. For example, to disable a baseline from the top-level menu, enter:

```
Select a menu option: system baseline requestedState disabled
```

Getting Out

To return to the menu one step higher in the hierarchy or to cancel an operation that you are currently performing, enter `q`, followed by [Return].

To quickly move to the top-level menu without backtracking through menus, press [Esc] (the Escape key). You immediately return to the top-level menu.

To completely leave the Administration Console, see the section “Exiting the Administration Console” on page 2-17.

Administration Console Interface Parameters

You can change two Administration Console interface parameters: the screen height and the functioning of the reboot and abort control keys.

Adjusting the Screen Height

You can change the Administration Console's screen height to increase or decrease the space available for displaying information.



The screen height setting does not affect the way the system displays menus. The screen height setting controls the way the system displays information that results from your use of the menus, such as when you request statistical summaries.

You can configure the screen height to be between 20 to 200 lines or zero (0) for infinite; the default is 24. Most terminal screens have a height of 24 lines.

Each time the screen output reaches the designated screen height, you are prompted to press a key to display more information. To receive no prompts, set the screen height to infinite (0). At this setting, however, the screen output might scroll beyond the screen, depending on your screen size.

To set the screen height:

- 1 From the top level of the Administration Console, enter:

```
system screenHeight
```

You are prompted for a screen height value.

- 2 Enter the screen height in lines (20 to 200). To receive no prompts, set the screen height to infinite (0).

Example:

```
Enter new screen height or 0 for infinite height [24]: 60
```

You are prompted about whether you want this value to be the default.

- 3 Enter **y** (yes) to use this screen height as the default for future Administration Console sessions. Enter **n** (no) if you want this screen height to be in effect only for this session.

Top-Level Menu

```

system
  ethernet
  fddi
  bridge
  ip
  snmp
  analyzer
  script
  logout
  display
  softwareUpdate
  baseline
  consoleSpeed
  telnet
  password
  name
  time
  screenHeight
  consoleLock
  ctKeys
  nvData
  reboot
  
```


Example:

```
Do you want this to be the new default screen height?
(y/n): y
```

Disabling the Reboot and Abort Keys

As shipped, the Administration Console allows you to use the [Ctrl + X] or [Ctrl + C] key combinations within the Administration Console. These key strokes allow you to reboot the system [Ctrl + X] or restart the Administration Console [Ctrl + C]. You can change this setting to disable both of these features.



CAUTION: *If you disable the control keys, only use [Ctrl + C] if instructed to by a Technical Support representative. Using [Ctrl + C] might irregularly terminate an Administration Console session.*

To enable or disable the reboot and abort control keys:


- 1 From the top level of the Administration Console, enter:

```
system ctlKeys
```

You are prompted for whether to enable or disable the functionality, as shown here:

```
Enter new value (disabled,enabled) [enabled]:
```

- 2 Enter **enabled** or **disabled** at the prompt.



Remote Access Parameters

You can reach the Administration Console remotely through a telnet or rlogin session. You can set parameters to prevent disconnections when another user remotely accesses the Administration Console, to enable the Switch 2200 to end remote sessions after a specified time period, and to specify the time interval before remote sessions are ended.

Preventing Disconnections

Because only a single shell is supported by the Administration Console, you might be disconnected from your session if someone else remotely accesses the Administration Console. A terminal connected through the Console serial port can be disconnected by a telnet or rlogin connection.

To ensure that your Administration Console session will not be pre-empted by remote access, you can lock the Administration Console. Remote access is prohibited only for that particular session.



The Administration Console is always locked when you are in the middle of a command. For example, the Administration Console is locked during a software update.

To lock the Administration Console:

Top-Level Menu

```

system
ethernet
bridge
ip
snmp
analyzer
script
logout
display
softwareUpdate
baseline
consoleSpeed
telnet
password
name
time
screenHeight
consoleLock
ctlKeys
nvData
reboot
  
```

- 1 From the top level of the Administration Console, enter:

```
system consoleLock
```

You are prompted to unlock (off) or lock (on) the Administration Console as shown here:

```
Enter new value (off,on) [on]:
```

- 2 Enter **off** to unlock the Administration Console or **on** to lock it.

Enabling Timeout of Remote Sessions

You can configure the Switch 2200 to disconnect remote sessions after a user-specified time interval of no activity. By default, the telnet timeout is disabled.

To enable or disable the telnet timeout:

Top-Level Menu

```

system
ethernet
bridge
ip
snmp
analyzer
script
logout
display
softwareUpdate
baseline
consoleSpeed
telnet
password
name
time
screenHeight
consoleLock
ctlKeys
nvData
reboot
timeOut
interval
  
```

- 1 From the top level of the Administration Console, enter:

```
system telnet timeOut
```

- 2 Enter the telnet timeout state (**off** or **on**).

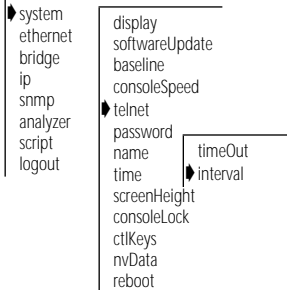
The default time interval is 30 minutes. To change this value, follow the instructions in the next section.

Setting Timeout Interval for Remote Sessions

You can set the timeout interval for remote sessions to any value from 30 minutes to 60 minutes. By default, the timeout interval is 30 minutes.

To set the telnet timeout interval:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
system telnet interval
```

- 2 Enter the telnet timeout interval (**30 minutes to 60 minutes**).

Running Scripts of Administration Console Tasks

You can use scripts to expedite and automate Administration Console tasks. Any command you enter in the Administration Console can become part of a script. You can even script your entire system setup so that you can repeat the exact setup on another Switch 2200.

You create scripts in an ASCII-based line editor, such as *EMACS* or *vi*. To run them from the Administration Console, you must access the directory where your scripts are stored. When writing scripts, you can use the number symbol (#) to identify comments in the script.

To run a script:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
script
```

You are prompted for information about where you have stored the script you want to run: host IP address, file path name, user name, and password. Press [Return] at any prompt to use the value in brackets.

- 2 Enter the host IP address of the system where the script resides.
- 3 Enter the path name.
- 4 Enter your user name.
- 5 Enter your password.
- 6 Enter the name of the script.

The task you scripted is run in the Administration Console.

The next example shows how you can script these tasks to initially configure your system:

- Setting up the Console port baud rate
- Setting the system name
- Assigning an IP address for management
- Checking the IP connection by pinging the Switch 2200
- Enabling Spanning Tree on the system
- Setting up SNMP trap reporting

```
# This script performs some start-up configurations.
#
# Set the Console serial port baud rate.
#
system consoleSpeed
300                # Console port baud rate
#
# Set the system name
#
system name
Engineering Switch2200_4
#
# Assign an IP address to the Switch 2200.
#
ip interface define
158.101.112.99     # IP address for the system
255.255.0.0        # subnet mask
158.101.255.255   # broadcast address
1                  # cost
all                # ports
#
ip interface display
#
# Validate access to management workstation
#
ip ping
158.101.112.26     # management workstation address
#
# Enable the Spanning Tree Protocol
#
bridge stpState enabled
#
# Configure my node as an SNMP trap destination
#
snmp trap add
158.101.112.26     # management workstation address
all                # turn on all traps
q                  # no more trap destinations
#
snmp trap display
#
```

Getting Help in the Administration Console

If you need assistance when using the Administration Console, it has online Help and an outlining feature, both of which can be accessed from any menu level. These features are described in this section.

Online Help

The Administration Console online Help provides an overview of the Administration Console and lets you access information about any menu option.

General online help

To get help using the Administration Console, enter `?`. The system displays general instructions for using the Administration Console.

Help for specific menu options

To get help for a specific menu option, enter `?` and the name of the option for which you want help. The system displays instructions, if available, for using that option.

For example, to get help on the **ethernet** option on the top-level menu, enter:

```
? ethernet
```

Viewing More Levels of Menu Options

The outlining feature allows you to list the menu options that fall lower than the current menu in the hierarchy. The default displays up to three levels of options.

To display the outline of available options below the current menu (up to three levels), enter **outline** (or **o**).

You can add a number to the command to modify how many levels you display. For example, to display two levels, enter:

```
outline 2
```

Exiting the Administration Console

If you are using an rlogin session to access the system, exiting will terminate the session. If you are accessing the system through the Console serial port, exiting returns you to the password prompt.


To exit from the Administration Console:

- 1 Return to the top level of the Administration Console, if you are not already there, by pressing the [ESC] key.
- 2 From the top-level menu, enter:

logout

Top-Level Menu

```
system
ethernet
fddi
bridge
ip
snmp
analyzer
script
▶ logout
```




SYSTEM-LEVEL FUNCTIONS

Chapter 3 Configuring Management Access to the System

Chapter 4 Administering Your System Environment

Chapter 5 Baselineing Statistics

Chapter 6 Saving, Restoring, and Resetting Nonvolatile Data



3

CONFIGURING MANAGEMENT ACCESS TO THE SYSTEM

This chapter describes how to configure management access to the SuperStack™ II Switch 2200 stackable switch through a serial connection or an IP interface. It also describes how to configure the Switch 2200 so that you can manage it using the Simple Network Management Protocol (SNMP).

About Management Access

You can access the Administration Console directly through the console serial port. Alternatively, from a PC or workstation, you can access the Administration Console through an Ethernet or FDDI port that has an IP interface configured for it. Once you establish an IP interface, you can also set up the system to be managed by an SNMP-based network management application, such as 3Com's Transcend® Enterprise Manager.

Using a Serial Connection

Direct access through the console serial port is often preferred because it allows you to stay attached during system reboots.



See the SuperStack™ II Switch 2200 Getting Started Guide for console port pin-outs.

Serial connections are often more readily available at a site than Ethernet connections. A Macintosh or PC attachment can use any terminal emulation program when connecting to the Console serial port. A workstation attachment under UNIX can use an emulator such as tip.

Using an IP Interface

An IP interface allows you to manage the system in-band through any Ethernet or FDDI port. Once an IP interface is configured, you can rlogin or telnet to the Administration Console using TCP/IP from a host, or you can access the SNMP agent from an external management application. The IP interface has a unique IP address.

In-band or Out-of-band?

By default, the Switch 2200 system provides in-band management through its Ethernet and FDDI ports. In-band management, management using the same network that carries regular data traffic, is often the most convenient and inexpensive way to access your system. If you are using a dedicated network for management data, then you are managing your network out-of-band.



If Spanning Tree is enabled and the port is in the blocking state, in-band management is not functional.

Setting Up the Console Serial Port

The default baud rate for the Console serial port is 9600. You might need to change the baud rate to match the port speed on your terminal.



Baud rate changes take effect immediately after you confirm the change. Adjust the baud rate of your terminal or terminal emulator appropriately to re-establish communication using the console serial port.

To set the baud rate for the Console serial port:

- 1 From the top level of the Administration Console, enter:

```
system consoleSpeed
```

- 2 Enter the baud rate for the serial port.

The system supports the following baud rates: 19200, 9600, 4800, 2400, 1200, and 300.

If you are connected to the Console serial port when you set the baud rate for that serial port, the following message is displayed:

```
Changing the baud rate may cause a loss of communication
since you are currently connected via the serial port.
Are you sure you want to change the baud rate? (y/n):
```

If you respond **y** (yes), the baud rate is changed immediately. At this time, you lose the ability to communicate on the serial port unless you adjust the baud rate of your terminal or terminal emulator (*tip*) appropriately. If you respond **n** (no), the baud rate does not change, and the previous menu is displayed.

Top-Level Menu

```

system
ethernet
fddi
bridge
ip
snmp
analyzer
script
logout
display
softwareUpdate
baseline
consoleSpeed
telnet
password
name
time
screenHeight
consoleLock
ctrlKeys
nvData
reboot

```

Setting Up an IP Interface for Management

IP is a standard networking protocol used for communications among various networking devices. To access the system using TCP/IP or to manage the system using SNMP, you must set up IP for your system as described in this section.

General Setup Process

You must first define an interface, which includes assigning an IP address to that interface, and then ping your IP management station to ensure that the connection is alive.



Assign an IP host address to every port for system management.

Then you can finalize your IP setup by ensuring that the configurations of the following are correct for your network and changing them as necessary:

- Routes (See page 3-7)
- Address Resolution Protocol (ARP) cache (See page 3-11)
- Routing Information Protocol (RIP) (See page 3-12)

You can monitor IP activity for your system by displaying the IP statistics at any time.

Administering Interfaces

You define interfaces to establish the relationship between the ports on your system and the subnets in your IP network. You can have up to 32 addresses on a single port and you can assign up to 17 ports per interface.

An IP interface has the following information associated with it:

■ IP Address

This address is specific to your network. Choose it from the range of addresses assigned to your organization. This address defines both the number of the network to which the interface is attached and the interface's host number on that network.

■ Subnet Mask

A subnet mask is a 32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, which as the subnet number, and which as the host number. Each IP address bit corresponding to a **1** in the subnet mask is in the network/subnet part of the address. Each IP address bit corresponding to a **0** is in the host part of the IP address.

- **Broadcast Address**

The system uses the IP address when it broadcasts packets to other stations on the same subnet. In particular, the system uses this address for sending RIP updates. By default, the system uses a directed broadcast (all 1s in the host field).

- **Cost**

The system uses this number, between 1 and 15, when calculating route metrics. Unless your network has special requirements, you should assign a cost of 1 to all interfaces.

- **Ports**

A single interface might contain several bridge ports. All of the ports corresponding to one interface share the same IP address, subnet mask, broadcast address, and cost. The Switch 2200 contains 17 ports: 1 FDDI and 16 Ethernet.

Be sure that the port to which your management station is attached is included in an interface.


Displaying Interfaces

You can display a table that shows all IP interfaces configured for the system, including their parameter settings.

To display IP interface information, enter the following command from the Administration Console top-level menu:

ip interface display

As shown in this example, the current configuration is displayed. It contains IP forwarding and RIP information as well as the IP interface information.



IP forwarding is enabled, RIP is active, ICMP router discovery is disabled.

Index	IP address	Subnet mask	Cost	Ports
1	158.101.1.1	255.255.255.0	1	1
2	158.101.4.1	255.255.255.0	1	2
3	158.101.6.1	255.255.255.0	1	5
4	158.101.8.1	255.255.255.0	1	8

Defining an Interface

When you define an interface, you define the interface's IP address, subnet mask, broadcast address, cost, and the collection of system ports associated with the interface.

Table 3-1 shows the recommended settings for the IP interface parameters if you are setting up the system for management.

Table 3-1 Recommended Settings for IP Management Access

Parameter	Recommended Setting
IP address	User defined
Subnet mask	User defined
Broadcast address	Directed (all 1s in the host field)
Cost	1
Ports	all



Defining an interface defines the IP broadcast domain for frames sourced from the attached segment. To avoid unintentional filtering of IP broadcasts, 3Com recommends that you include all ports. If you do not assign all ports to this interface, be sure that you include the port to which your network management station is attached.

To define an IP interface:


- 1 From the top level of the Administration Console, enter:

```
ip interface define
```

You are prompted for the interface's parameters. To use the value in brackets, press [Return] at the prompt.

- 2 Enter the IP address of the interface.

Top-Level Menu



- 3 Enter the subnet mask of the network to which the interface is to be connected.
- 4 Enter the broadcast address to be used on the interface.
- 5 Enter the cost value of the interface.
- 6 Enter the port(s) that you want to include in the interface. Separate nonconsecutive ports with commas (.). Enter a consecutive series of ports using a hyphen (-).

Example:

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter broadcast address [158.101.1.255]:
Enter cost [1]:
Ports 1-2=FDDI, 3-18=Ethernet
Select port(s) (1-18|all): 2-4,8
```



If you physically change the configuration of your system after defining IP interfaces, the ports designated for those interfaces might no longer be valid. You should reconfigure your interfaces.

Modifying an Interface

To modify an IP interface that you have already defined:


- 1 From the top level of the Administration Console, enter:

```
ip interface modify
```

You are prompted for the interface parameters. Press [Return] at the prompts for which you do not want to modify the value in parentheses.

- 2 Modify the existing interface parameters by entering a new value at the prompt.

Top-Level Menu




Removing an Interface

You might want to remove an interface if you no longer need to communicate with IP on the ports associated with that interface.

To remove an IP interface definition:

- 1 From the top level of the Administration Console, enter:
ip interface remove
- 2 Enter the index numbers of the interfaces you want to remove.



Administering Routes

Each system maintains a table of routes to other IP networks, subnets, and hosts. You can either make static entries in this table using the Administration Console or configure the system to use RIP to automatically exchange routing information.

Each routing table entry contains the following information:

- **Destination IP Address and Subnet Mask**

These elements define the address of the destination network, subnet, or host. A route matches a given IP address if the bits in the IP address that corresponds to the bits set in the route subnet mask match the route destination address. When it forwards a packet, if the system finds more than one routing table entry matching an address (for example, a route to the destination network and a route to the specific subnet within that network), it will use the most specific route (that is, the route with the most bits set in its subnet mask).

- **Routing Metric**

This metric specifies the number of networks or subnets that a packet must pass through to reach its destination. This metric is included in RIP updates to allow routers to compare routing information received from different sources.

- **Gateway IP Address**

This address tells the router how to forward packets whose destination address matches the route's IP address and subnet mask. The system forwards such packets to the indicated gateway.

- **Status**

The status of the route provides the information described in Table 3-2.

Table 3-2 Route Status

Status	Description
Direct	Route to a directly connected network
Static	Route was statically configured
Learned	Route was learned using indicated protocol
Timing out	Route was learned but is partially timed out
Timed out	Route has timed out and is no longer valid

In addition to the routes to specific destinations, the routing table can contain an additional entry, called the *default route*. The system uses the default route to forward packets that do not match any other routing table entry. You might want to use a default route in place of routes to numerous destinations that all have the same gateway IP address.

Displaying the Routing Table

You can display the routing table for the system to determine which routes are configured and if they are operating.


To display the contents of the routing table, enter the following from the top level of the Administration Console:

```
ip route display
```

In the following example, routes for the Switch 2200 are displayed. The configuration of RIP is indicated in the status display.

Destination	Subnet mask	Metric	Gateway	Status
158.101.4.0	255.255.255.0	2	158.101.2.8	Static
158.101.3.0	255.255.255.0	2	158.101.1.2	Learned(RIP)
158.101.2.0	255.255.255.	1	--	Direct
158.101.1.0	255.255.255.0	1	--	Direct
Default Route	--	5	158.101.1.2	Learned (RIP)

Top-Level Menu



Defining a Static Route

You might want to define a static route to transmit system traffic, such as system pings or SNMP response, through a consistent route. Before you define static routes, you must define at least one IP interface. (See “Defining an Interface” on page 3-5.) Static routes remain in the table until you remove them, or until you remove the corresponding interface. Static routes take precedence over dynamically learned routes to the same destination.

To define a static route:

- 1 From the top level of the Administration Console, enter:

```
ip route static
```

You are prompted for the route's parameters. To use the value in brackets, press [Return] at the prompt.

- 2 Enter the destination IP address of the route.
- 3 Enter the subnet mask of the route.
- 4 Enter the gateway IP address of the route.

A static route is defined in the following example:

```
Enter destination IP address: 158.101.4.0
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter gateway IP address: 158.101.2.8
```

Removing a Route



To remove a route:

- 1 From the top level of the Administration Console, enter:

```
ip route remove
```

- 2 Enter the destination IP address of the route.
- 3 Enter the subnet mask of the route.

The route is immediately deleted from the routing table.



Flushing a Route

Flushing deletes all learned routes from the routing table.

To flush all learned routes, enter the following from the top level of the Administration Console:

```
ip route flush
```

All learned routes are immediately deleted from the routing table.

Setting the Default Route

The system uses the default route to forward packets that do not match any other routing table entry. A system can learn a default route using RIP, or you can configure a default route statically.

If a system's routing table does not contain a default route, either statically configured or learned using RIP, then it cannot forward a packet that does not match any other routing table entry. If it cannot forward a packet for this reason, then it drops the packet and sends an ICMP "destination unreachable" message to the host that sent the packet to notify it of the problem.

To statically configure the default route:

- 1 From the top level of the Administration Console, enter:

```
ip route default
```

- 2 Enter the gateway IP address of the route.

The default route is immediately added to the routing table.


Removing the Default Route

To remove a default route, enter the following from the top level of the Administration Console:


```
ip route noDefault
```

The default route is immediately removed from the routing table.


Top-Level Menu



Top-Level Menu



Top-Level Menu



Administering the ARP Cache

The Switch 2200 uses the Address Resolution Protocol (ARP) to find the MAC addresses corresponding to the IP addresses of hosts and routers on the same subnets. An ARP cache is a table of known IP addresses and their corresponding MAC addresses.

Displaying the ARP Cache

To display the contents of the ARP cache, enter the following command from the top level of the Administration Console:

```
ip arp display
```

The contents of the ARP cache are displayed as shown in this example:

```
RIP is active.
```

IP Address	MAC Address	Interface
158.101.1.112	08-00-1e-31-a6-2	1
158.101.1.117	08-00-1e-65-21-07	1

Removing an ARP Cache Entry

You might want to remove an entry from the ARP cache if the MAC address has changed.



To remove an entry from the ARP cache:

- 1 From the top level of the Administration Console, enter:


```
ip arp remove
```

- 2 Enter the IP address you want to remove.

The address is immediately removed from the table. If necessary, the system will subsequently use ARP to find the new MAC address corresponding to that IP address.



Top-Level Menu



Flushing ARP Cache Entries

You might want to delete all entries from the ARP cache if the MAC address has changed.

To remove all entries from the ARP cache, enter the following command from the top level of the Administration Console:

```
ip arp flush
```

The ARP cache entries are immediately removed from the table.

Setting the RIP Mode

You can select a RIP mode that is appropriate for your network. RIP can operate in one of two modes:

- *Off* — The station ignores all incoming RIP packets and does not generate any RIP packets of its own.
- *Passive* — The station processes all incoming RIP packets and responds to explicit requests for routing information, but it does not broadcast periodic or triggered RIP updates.

RIP default mode

By default, RIP operates in passive mode.

To set the RIP operating mode:

- 1 From the top level of the Administration Console, enter:


```
ip rip
```

- 2 Enter the RIP mode (**off** or **passive**). To use the value in brackets, press [Return] at the prompt.

See the following example:

```
Select RIP mode (off, passive) [passive]: off
```

Top-Level Menu



Pinging an IP Station

Once you have set up your IP interface, you might want to check to see if the Switch 2200 system can communicate with other systems over the IP network. To check, you can “ping” the IP address of your management station.

Pinging uses the Internet Control Message Protocol (ICMP) echo facility to send an ICMP echo request packet to the IP station you specify. It then waits for an ICMP echo reply packet. Possible responses from pinging are:

- Alive
- No answer
- Network is unreachable. A network is unreachable when there is no route to that network.

To ping an IP station:

- 1 From the top level of the Administration Console, enter:

```
ip ping
```

- 2 Enter the IP address of the station you want to ping.

```
IP Address: 192.9.200.40
```

You could receive one of the following responses:

```
192.9.200.40 is alive
```

OR

```
no answer from 192.9.200.40
```

For a remote IP address, you can also receive the following response:

```
Network is unreachable
```

You should receive a response that the address you pinged is *Alive*. If you do not receive this response, be sure that you have defined the correct interface values.

Top-Level Menu

```
system
ethernet
fddi
bridge
ip
snmp
analyzer
script
logout
interface
route
arp
rip
ping
statistics
```

Displaying IP Statistics

The IP statistics you can view are described in Table 3-3.

Table 3-3 IP Statistics

Field	Description
inReceives	Total number of IP datagrams received, including those with errors
forwDatagrams	Number of datagrams that the IP station attempted to forward
inDelivers	Number of datagrams that the IP station delivered to local IP client protocols
outRequests	Number of datagrams that local IP client protocols passed to IP for transmission
outNoRoutes	Number of datagrams that the IP station discarded because there was no route to the destination
inHdrErrors	Number of datagrams that the IP station discarded because the IP header contained errors
inAddrErrors	Number of datagrams that the IP station discarded because of an error in the source or destination IP address

Top-Level Menu

```

system
 ethernet
  fddi
  bridge
  ip
  snmp
  analyzer
  script
  logout
  interface
  route
  arp
  rip
  ping
  statistics

```

To display IP statistics, enter the following from the top level of the Administration Console:

```
ip statistics
```

Statistics are displayed, as shown in this example:

```
IP forwarding is enabled, RIP is active, ICMP router discovery is disabled.
```

```

      inReceives      forwDatagrams      inDelivers      outRequests
             51213             49743             3227             2285

      outNoRoutes      inHdrErrors      inAddrErrors
             273             7             0

```

Setting Up SNMP on Your System

To manage the Switch 2200 from an external management application, you must configure SNMP community strings and set up trap reporting as described in this section.

You can manage the Switch 2200 using an SNMP-based external management application. This application (an SNMP manager) sends requests to the Switch 2200 system, where they are processed by the Switch SNMP agent.

The SNMP agent provides access to the collection of information about the Switch 2200. In addition, a Switch 2200 SNMP agent sends traps to an SNMP manager to report significant events. Access to system information through SNMP is controlled by community strings.

For more information about using SNMP to manage the Switch 2200, see Chapter 3: *Management Access: Protocols* in the *SuperStack™ II Switch 2200 Operation Guide*.

Displaying SNMP Settings

You can display the current Switch 2200 SNMP configurations for the community strings.


To display SNMP settings, enter the following from the top level of the Administration Console:

```
snmp display
```

The community string settings are displayed as shown here:

```
Read-only community is public
Read-write community is private
```

Top-Level Menu



Configuring Community Strings

A community string is an octet string, included in each SNMP message, that controls access to system information. The Switch 2200 SNMP agents internally maintain two community strings that you can configure:

- *Read-only* community strings with the default “public”
- *Read-write* community strings with the default “private”

When an SNMP agent receives an SNMP request, the community string in the request is compared with the community strings configured for the agent. SNMP *get*, *get-next*, and *set* requests are valid if the community string

in the request matches the agent's *read-write* community. Only the SNMP *get* and *get-next* requests are valid if the community string in the request matches the *read-only* community.

Community string length

When you set a community string, you can specify any value up to 48 characters long.

To set a community string:

- 1 From the top level of the Administration Console, enter:

```
snmp community
```

You are prompted for a read-only community value and then a read-write community value. If you do not want to change the value of a community string, press [Return] at either prompt.

- 2 At the read-only prompt, enter the community string.
- 3 At the read-write prompt, enter the community string.

Administering SNMP Trap Reporting

For network management applications, you can use the Administration Console to manually administer the trap reporting address information.

Displaying Trap Information

Displaying the trap reporting information shows you the various SNMP traps and the current configured destinations, as well as whether the proxying of remote SMT traps is enabled or disabled.

To show the configured trap reporting information, enter the following from the top level of the Administration Console:

```
snmp trap display
```

Top-Level Menu

```
system
ethernet
fddi
bridge
ip
snmp
analyzer
script
logout
  display
  community
  trap
```

Top-Level Menu

```
system
ethernet
fddi
bridge
ip
snmp
analyzer
script
logout
  display
  community
  trap
    display
    addModify
    remove
    flush
    smtProxyTraps
```

Here is an example display of the SNMP trap reporting information:

Trap Descriptions:

Trap #	Description
1	MIB II: Coldstart
2	MIB II: Authentication Failure
3	Bridge MIB: New Root
4	Bridge MIB: Topology Change
5	LANplex Systems MIB: System Overtemperature
10	LANplex Systems MIB: Address Threshold
12	LANplex Opt FDDI MIB: SMT Hold Condition
13	LANplex Opt FDDI MIB: SMT Peer Wrap Condition
14	LANplex Opt FDDI MIB: MAC Duplicate Address Condition
15	LANplex Opt FDDI MIB: MAC Frame Error Condition
16	LANplex Opt FDDI MIB: MAC Not Copied Condition
17	LANplex Opt FDDI MIB: MAC Neighbor Change
18	LANplex Opt FDDI MIB: MAC Path Change
19	LANplex Opt FDDI MIB: Port LER Condition
20	LANplex Opt FDDI MIB: Port Undesired Connection
21	LANplex Opt FDDI MIB: Port EB Error Condition
22	LANplex Opt FDDI MIB: Port Path Change

Trap Destinations Configured:

Address	Trap Numbers Enabled
158.101.112.3	1-10, 12-21


Proxying of remote SMT events is disabled

Configuring Trap Reporting

You can add new trap reporting destination configurations or you can modify an existing configuration. You can define up to ten destination addresses and the set of traps that are sent to each destination address.

To add a new trap reporting destination configuration or modify a current one:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
snmp trap addModify
```

The system prompts you for a trap destination address, that is, the IP address of the SNMP manager that will receive the traps.

- 2 Enter an IP address of the SNMP manager (destination address) .

3 Enter the trap number(s).

Separate a series of more than two trap numbers with a hyphen (-) and nonsequential trap numbers by commas. Enter **all** if you want to enable all the traps for the destination.



The trap numbers you enter allow the trap specified by that number to be sent to the destination address when the corresponding event occurs. No unlisted traps are transmitted.

This example shows a trap configuration:

```
Enter the trap destination address: 158.101.222.3
Enter the trap numbers to enable (1-5,10,12-22|all)
[1-5,10,12-22]: all
```

Address Error

If the destination address you entered is not a valid end-station or if the agent does not have a route to the destination, you receive this message:

```
Trap address invalid or unreachable
```

If you see this message, confirm the address of the end-station and confirm that it is online.

Removing Trap Destinations

When you remove a destination, no SNMP traps will be reported to that destination.

To remove a destination:

1 From the top level of the Administration Console, enter:


```
snmp trap remove
```

You are prompted for a trap destination address, that is, the IP address of the SNMP manager that will no longer receive the traps.

2 Enter the SNMP trap reporting destination address you want to remove.

The destination address is removed and you return to the previous menu.

Top-Level Menu



Flushing Trap Destinations

When flushing the SNMP trap reporting destinations, you remove all trap destination address information for the SNMP agent.

To flush all SNMP trap reporting destinations:

- 1 From the top level of the Administration Console, enter:


```
snmp trap flush
```

You receive the following prompt:

```
Are you sure? (n/y) [y]:
```

- 2 Enter **y** (yes) or **n** (no) at the prompt.

If you enter **y**, the addresses are immediately flushed. If you enter **n**, you return to the previous menu.



Setting Up SMT Event Proxying

FDDI SMT events, which occur on the FDDI ring, can be reported to stations through the Status Report Protocol. Several SNMP traps, defined in the LANplex Optional FDDI MIB, correspond to some of these events and conditions. If you want your Switch 2200 to report remote SMT events as SNMP traps, you must enable proxying of remote SMT events in that Switch 2200 system.



Local SMT events are automatically reported by the SNMP agent in a Switch 2200 system.

If you have a single Switch 2200 on your network and you have no other way to access FDDI information, then you should enable proxying of SMT events. This configuration provides access to the events occurring locally on the Switch 2200 and to those reported by other stations on the FDDI ring.

If you have multiple Switch 2200s on your FDDI network all reporting to the same SNMP management station, then you can do one of the following:

- On only one Switch 2200, 1) enable local SNMP traps as described in the “Configuring Trap Reporting” on page 3-17 and 2) enable proxying of remote SMT events. On all other Switch 2200s in your network, 1) disable proxying of remote SMT events and 2) enable only SNMP traps that are *not* SMT-related. SMT-related traps include all of those in the LANplex Optional FDDI MIB. This configuration provides access to the events

occurring locally on the one Switch 2200 and to those reported by other stations on the FDDI ring (including other Switch 2200s).

- Enable local SNMP traps and disable the proxying of remote SMT events on every Switch 2200 in your network. Local traps will be reported to the management station (which will cover all your Switch 2200s), but SMT events from systems other than Switch 2200s in your network will not be reported.

To enable or disable the proxying of remote SMT events:


- 1 From the top level of the Administration Console, enter:

```
snmp trap sntProxyTraps
```

- 2 Enter **disabled** or **enabled** at the prompt.

The proxying of remote SMT traps is disabled or enabled for the system.

Top-Level Menu



4

ADMINISTERING YOUR SYSTEM ENVIRONMENT

This chapter focuses on the administration of your SuperStack™ II Switch 2200 system environment, which involves:

- Displaying the current system configuration
- Setting system passwords
- Setting the system name
- Changing the system date and time
- Rebooting

Displaying the System Configuration

Top-Level Menu

system	display
ethernet	softwareUpdate
fdci	baseline
bridge	consoleSpeed
ip	telnet
snmp	password
analyzer	name
script	time
logout	screenHeight
	consoleLock
	ctrlKeys
	nvData
	reboot

The system configuration display provides software and hardware revisions and warning messages for certain system conditions.

To display the configuration of a Switch 2200, enter the following command from the top level of the Administration Console:

```
system display
```

Example of a Switch 2200 system configuration display:

```
Switch 2200 (rev 1.3) - System ID 0f2b00  
Intelligent Switching Software  
Version 7.1.0 - Built 7/24/96 06:26:55 PM
```

The display contains the following general system information:

- The system type (Switch 2200)
- System ID
- Software version
- Software build date and time

Warning messages

You will also see a warning message in the display, and the system bell will ring, if the system detects any of the following conditions:

- System temperature has exceeded the maximum level for normal operation
- Fan failure
- Power supply failure

Setting Passwords

The Administration Console supports three levels of password: one for browsing or viewing only (read), one for configuring network parameters (write), and one for full system administration (administer).

Initial passwords

Because the initial passwords stored in the nonvolatile memory of the system are null, just press the Return key at the password prompt.

You can only change passwords by entering the Console using the *administer* access level.

To set a password:

- 1 From the top level of the Administration Console, enter:


```
system password
```
- 2 At the prompt requesting you to enter a password access level to change, enter one of the following:


```
read
write
administer
```
- 3 At the prompt for your old password, enter the old password.
- 4 Enter the new password.

The password can have up to 32 characters and is case sensitive. To enter a null password, press [Return].
- 5 Retype the new password for verification. The system does not display the password as you type.

Example:

```
Select menu option (system): password
Password access level (read, write, administer): read
Old password:
New password:
Retype new password:
```

Top-Level Menu

```

system
ethernet
fdci
bridge
ip
snmp
analyzer
script
logout
display
softwareUpdate
baseline
consoleSpeed
telnet
password
name
time
screenHeight
consoleLock
ctlKeys
nvData
reboot

```

The administration console password has been successfully changed.

- Repeat steps 1 through 5 for each level of password you want to configure.

Setting the System Name

You should give the Switch 2200 an easily recognizable and unique name to help you manage the system. For example, you might want to name the system according to its physical location (say, SS2200 ENGLAB).

To name the system:

- From the top level of the Administration Console, enter:

system name

You are prompted for the name of the system:

Enter new string (no spaces) [Switch 2200]:

- Enter a name that is both unique on the network and meaningful to you.

The new system name appears the next time you display the system configuration.

Top-Level Menu

system	display
ethernet	softwareUpdate
fdi	baseline
bridge	consoleSpeed
ip	telnet
snmp	password
analyzer	name
script	time
logout	screenHeight
	consoleLock
	ctlKeys
	nvData
	reboot

Changing the Date and Time

The Switch 2200's internal clock is initialized at the factory. You can display and change the system's current date and time.

To change either the date or the time:

- From the top level of the Administration Console, enter:

system time

The system displays the current date and time, along with a prompt asking you if the date and time are correct. Example:

The current system time is 08/24/96 04:37:57 PM.

Is this correct? (y/n):

- Enter **y** (yes) or **n** (no) at the prompt.

If you respond **y**, you return to the main menu. If you respond **n**, the system prompts you for the correct date and time.

- Enter the correct date and time in this format: mm/dd/yy hh:mm:ss xM
Table 4-1 discusses the format variables.

Top-Level Menu

system	display
ethernet	softwareUpdate
fdi	baseline
bridge	consoleSpeed
ip	telnet
snmp	password
analyzer	name
script	time
logout	screenHeight
	consoleLock
	ctlKeys
	nvData
	reboot

Table 4-1 Date and Time Variables

Format	Description
<i>first</i> mm	month (1–12)
dd	date (1–31)
yy	last two digits of the year (00–99)
hh	hour (1–12)
<i>second</i> mm	minute (00–59)
ss	second (00–59)
xM	either AM or PM

- 4 Press [Return] when you want the system to start keeping the time that you entered.

Example:

Enter the new system time (mm/dd/yy hh:mm:ss xM):

09/30/96 10:00:00 AM

Press RETURN at the exact time:

Rebooting the System

If your system is connected to the Administration Console through an rlogin or telnet session, rebooting the system disconnects your session. To retain a connection to the Administration Console during reboots so that you can view diagnostic information, you must connect your system through the Console serial port.

To reboot the system:

- 1 From the top level of the Administration Console, enter:

system reboot

The following message appears:

Are you sure you want to reboot the system? (y/n):

- 2 Enter **y** (yes) or **n** (no).

If you enter **y**, the system reboots. If you enter **n**, you return to the previous menu.

Top-Level Menu

system	display
ethernet	softwareUpdate
fdci	baseline
bridge	consoleSpeed
ip	telnet
snmp	password
analyzer	name
script	time
logout	screenHeight
	consoleLock
	ctlKeys
	nvData
	reboot

5

BASELINING STATISTICS

This chapter describes how baselining statistics work in the SuperStack™ II Switch 2200, and how to set, display, enable, or disable a baseline statistic.

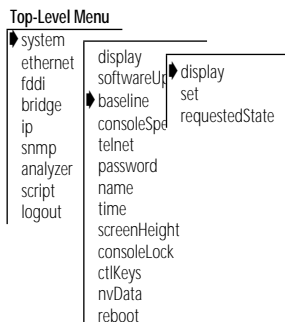
About Setting Baselines

Normally, statistics for MACs and ports start compiling at system power-up. Baselining allows you to view statistics over the period of time since a baseline was set. By viewing statistics relative to a baseline, you can more easily evaluate recent activity in your system or on your network.

Baselining is maintained across Administration Console sessions. Statistics you view after setting the baseline indicate that they are relative to the baseline. To view statistics as they relate only to the most recent power up, you must disable the baseline.

Baselining affects the statistics displayed for Ethernet ports, FDDI resources, and bridges.

Displaying the Current Baseline



You can display the current baseline to see when the baseline was last set and to determine if you need a newer baseline for viewing statistics.

To display the current baseline, enter the following commands from the top level of the Administration Console:

```
system baseline display
```

Example:

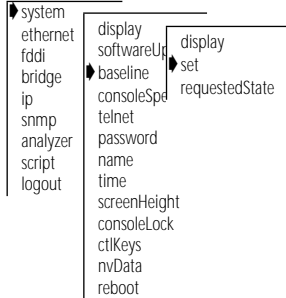
```
Baseline set at 08/07/96 10:42:52 AM is currently enabled.
```

If a baseline has not been set on the system, you see the following message:

```
Baseline has not yet been set.
```

Setting Baselines

Top-Level Menu



Setting a baseline resets the counters to zero. The accumulated totals since power up are maintained by the system. The baseline is time-stamped.

To set a baseline, enter the following commands from the top level of the Administration Console:

```
system baseline set
```

A message similar to the following appears:

```
Baseline set at 08/07/96 10:42:52 AM.
```

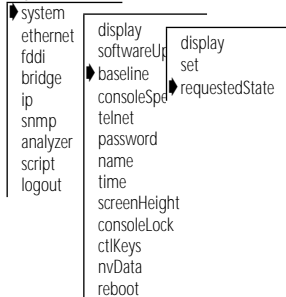
Baselining is automatically enabled when a baseline is set.

Enabling or Disabling Baselines

When you re-enable a baseline, the counters return to the values accumulated from the most recent baseline you set. Disabling a baseline returns the counters to the total accumulated values since the last power up.

To enable the current baseline:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
system baseline requestedState
```

You are prompted to enter a new baseline state, as shown here:

```
Enter new value (disabled,enabled) [enabled]:
```

- 2 Enter **disabled** or **enabled** at the prompt.

The new value is confirmed as shown here :

```
Baseline set at 08/07/96 10:42:52 AM has been disabled.
```

6

SAVING, RESTORING, AND RESETTING NONVOLATILE DATA

This chapter describes the nonvolatile (NV) data in the SuperStack™ II Switch 2200 system and how to save, restore, and reset the data.

About Working with Nonvolatile Data

If you want to transfer NV data from one system to another, save the system's NV data and restore it as appropriate. You might also want to save a certain configuration of the system for your reference and as a backup. You can also reset system data to its factory-configured values, if necessary.

During a save, the contents of NV memory are written out to a disk file. All configurable parameters are saved in nonvolatile memory, including:

- System name
- System date and time
- Passwords
- Packet filters
- Ethernet port labels
- FDDI resources settings
- Bridge and bridge port settings
- SNMP community string settings
- SNMP trap destination configurations

The file also contains the following information, which is used to resolve any inconsistencies when NV data is restored:

- Software version number
- System ID
- Date and time of creation
- Data checksums

Saving NV Data

When NV data is saved, it is written to a disk file on a host computer. The information can then be retrieved from the disk file when you use the restore command.

To save NV data:

- 1 From the top level of the Administration Console, enter:

```
system nvData save
```

You are prompted for information for saving the data. To use the value in brackets, press [Return] at the prompt. Any entry for IP address, file name, and user name becomes the new default.

- 2 Enter the IP address of the station to which you want to save the NV data.
- 3 Enter the file path name where you want to save the file.
- 4 Enter your user name on the host system.
- 5 Enter your password on the host system.
- 6 Enter a name of the file (optional).

Example:

```
Host IP Address [158.101.100.1]: 158.101.112.34
NV Data file pathname: usr/jones/systemdata
User name: Tom
Password:
Enter an optional file label: Labdata
```


If the information is incorrect or a connection could not be made with the specified host, a message similar to the following is displayed:

```
Login incorrect.
Error: Could not open ftp session
```

If a session is successfully opened, a system message notifies you of the success or failure of your save as in the following examples:

Success System NV data successfully stored in usr/jones/systemdata of host 158.101.112.34.

Failure Error - Configuration not stored.



The failure message varies depending on the problem encountered while saving the NV data.

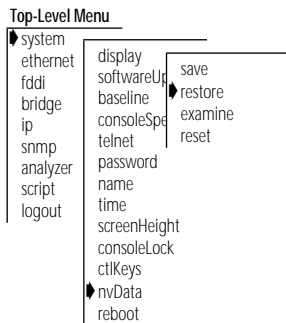
At the end of the save, you are returned to the previous menu.

Restoring NV Data

When you restore system NV data, the software presents you with a proposal for how to restore the data. This proposal is based on the restoration rules described here:

- Rule 1 Exact Match* — An exact match is one where the system IDs, module types, and module revisions (if applicable) all match between the saved configuration and the system on which you are restoring the image.
- Rule 2 System ID Mismatch* — System IDs do not match between the saved NV file and the target system. Mismatches in system IDs are allowed. Before restoring the NV data to a system with a different system ID, you should be aware of the following NV data that might cause problems when restored:
- Management IP addresses (defined in IP interface configurations) are saved as NV data and restored. Before connecting the restored system to the network, you might need to change the IP address of defined interfaces to avoid duplicate IP address problems. Modifying IP interface definitions is described on page 3-6.
 - Statically configured Ethernet addresses are saved as NV data. You must be sure not to have duplicate addresses when you restore the NV data. Listing statically configured addresses is described on page 11-11.

If none of these rules succeeds, you cannot apply the saved configuration to the system.



To restore the NV data:

- 1 From the top level of the Administration Console, enter:

```
system nvData restore
```

You are prompted for information for restoring the NV data saved to a file. Press [Return] at a prompt to use the value specified in brackets. Any entry for IP address, file name, and user name becomes the new default.

- 2 Enter the IP address of the host where the NV data file resides.
- 3 Enter the NV data file path name.
- 4 Enter your user name on the host system.
- 5 Enter your password on the host system.

If the information is incorrect, or a connection could not be made with the specified host, a message similar to the following is displayed:

```
User Tom access denied:
Error: Could not open ftp session
```

If a session is successfully opened, the system reads the header information, compares the stored configuration to the current system configuration, and proposes a method of restoration based on one of the restoration rules described on page 6-3.

You are prompted to load the proposal.

```
CAUTION - Restoring nonvolatile data may leave the system
in an inconsistent state and therefore a reboot is
necessary after each restore.
```

```
Do you wish to continue? (y/n):
```

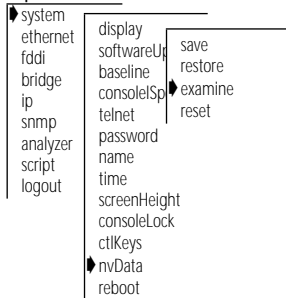
- 6 Enter **y** (yes) if you want to use the proposal. If you do not want to use the proposal, enter **n** (no).

If you enter **y**, the system NV is restored as proposed.

- 7 At the end of a restore, press [Return] to reboot the system.

Examining a Saved NV Data File

Top-Level Menu



After saving NV data to a file, you can examine the header information of that file.

To examine the file:

- 1 From the top level of the Administration Console, enter:

```
system nvData examine
```

You are prompted for information for examining a saved NV data file. Press [Return] at a prompt to use the value specified in brackets. Any entry for IP address, file name, and user name becomes the new default.

- 2 Enter the IP address of the host where the NV data file resides.
- 3 Enter the NV data file path name.
- 4 Enter your user name on the host system.
- 5 Enter your password on the host system.

If the information is incorrect, or a connection could not be made with the specified host, a message similar to the following is displayed:

```
User Tom access denied:
Error: Could not open ftp session
```

If a session is successfully opened, the system displays the header information that corresponds to the file entered. See the following example:

```
Product ID #, Product Type #
System ID 102
Saved October 8, 1994 10:24:12. Configuration version 3.
```

You are returned to the NV data menu options.

Resetting NV Data to Defaults

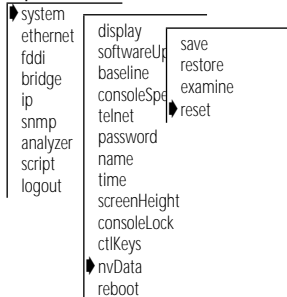
At times you may not want to restore the system NV data. Instead, you may want to reset the values to the factory defaults so that you can start configuring the system from the original settings.



CAUTION: *Resetting the NV data means that all NV memory is set back to the factory defaults. Before proceeding, ensure that you want to reset your NV data.*

To reset all the NV data on the system to the original default values:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:


```
system nvData reset
```

You see the following prompt:

```
Resetting nonvolatile data may leave the system in an
inconsistent state and therefore a reboot is necessary
after each reset.
```


```
Do you wish to continue (n,y) [y]:
```

- 2 Confirm that you want to reset NV data by entering **y** (yes) at the prompt. If you enter **y** (yes) the system will reboot. If you enter **n** (no), you are returned to the previous menu.
- 3 Reboot the system.



ETHERNET AND FDDI PARAMETERS

- Chapter 7** Administering Ethernet Ports
- Chapter 8** Administering FDDI Resources
- Chapter 9** Setting Up the System for Roving Analysis



7

ADMINISTERING ETHERNET PORTS

This chapter describes how to:

- View Ethernet port information
- Configure Ethernet port labels
- Enable or disable an Ethernet port

Displaying Ethernet Port Information

You can display either a summary of Ethernet port information or a detailed report. When you display a summary of Ethernet port information, you view its label, status, and the most pertinent statistics about general port activity and port errors. The detailed display of Ethernet port information includes the information in the summary and additional Ethernet port statistics, such as collision counters.

If you want to display Ethernet port statistics relative to a baseline, see Chapter 5 for more information.

To display information about the Ethernet ports:

- 1 From the top level of the Administration Console, enter:


```
ethernet summary
```

OR

```
ethernet detail
```

- 2 Enter the port(s) for which you want to view information.

The port information is displayed in the format you specified. The following example shows a detailed display for Ethernet ports:



port	rxFrames	rxBytes	rxFrameRate	rxByteRate
1	406430	36336795	0	0
12	242400	29275605	0	0
port	rxPeakByteRate	rxPeakFrameRate	noRxBuffers	alignmentErrs
1	90484	163	0	0
12	58438	394	0	0
port	fcsErrs	lengthErrs	rxInternalErrs	rxDiscards
1	0	0	0	0
12	0	0	0	0
port	rxUnicasts	rxMulticasts	txFrames	txBytes
1	365811	40619	1422085	234636091
12	242033	367	1256455	300242671
port	txFrameRate	txByteRate	txPeakFrameRate	txPeakByteRate
1	3	345	208	271724
12	3	345	402	321722
port	txQOverflows	excessCollision	excessDeferrals	txInternalErrs
1	0	0	0	0
12	0	0	0	0
port	carrierSenseErr	txDiscards	txUnicasts	txMulticasts
1	0	0	528268	893836
12	0	0	322389	934076
port	collisions	lateCollisions	requestedState	portState
1	0	0	enabled	on-line
12	0	0	enabled	on-line
port	portType	linkStatus	macAddress	
1	10BaseT(RJ45)	enabled	00-80-3e-0b-48-02	
12	10BaseT(RJ45)	enabled	00-80-3e-0b-48-0d	
port		portLabel	duplexMode	
1		Office113_SPARCstation5	n/a	
12		Office322_Quadra900	n/a	

An example of a summary display for Ethernet ports is shown here:

```

port                                portLabel                portState
  1                                Office113_SPARCstation5  on-line
 12                                Office322_Quadra900     on-line

port      rxFrames      txFrames      rxBytes      txBytes
  1          406876      1423733      36377226    234900612
 12          242532      1257721      29293858    300479754

port      rxErrs      txErrs      noRxBuffers  txQOverflows
  1           0           0           0             0
 12           0           0           0             0
    
```

Table 7-1 describes the information provided about an Ethernet port.

Table 7-1 Description of Fields for Ethernet Port Attributes

Field	Description
alignmentErrs	Number of frames received by this port that are not an integral number of octets in length and do not pass the FCS check
carrierSenseErr	Number of frames discarded because the carrier sense condition was lost while attempting to transmit a frame from this port
collisions	Number of collisions detected on this port
duplexMode	Current duplex mode setting. Possible values are full, half, and not applicable (n/a). Duplex mode is not applicable on the Switch 2200.
excessCollision	Number of frames that could not be transmitted on this port because the maximum allowed number of collisions was exceeded
excessDeferrals	Number of frames that could not be transmitted on this port because the maximum allowed deferral time was exceeded
fcsErrs	Number of frames received by this port that are an integral number of octets in length but do not pass the FCS check
lateCollisions	Number of times a collision was detected on this port later than 512 bit-times into the transmission of a frame
lengthErrs	Number of frames received by this port longer than 1518 bytes or shorter than 64 bytes
linkStatus	Boolean value indicating the current state of the physical link status for this port (either enabled or disabled)
macAddress	The MAC address of this port
noRxBuffers	Number of frames discarded because there was no available buffer space

(continued)

Table 7-1 Description of Fields for Ethernet Port Attributes (continued)

Field	Description
portLabel	32-character string containing a user-defined name. The maximum length of the string is 32 characters, including the null terminator.
portState	Current software operational state of this port. Possible values are on-line and off-line.
portType	Specific description of this port's type.
requestedState	Configurable parameter used to enable and disable this port. The default is enabled.
rxByteRate	Average number of bytes received per second by this port during the most recent sampling period
rxBytes	Number of bytes received by this port, including framing characters
rxDiscards	Number of received frames discarded because there was no higher layer to receive them or because the port was disabled
rxErrs	Sum of all receive errors associated with this port (summary report only)
rxFrameRate	Average number of frames received per second by this port during the most recent sampling period. Sampling periods are 1 second long and are not configurable.
rxFrames	The number of frames copied into receive buffers by this port
rxInternalErrs	Number of frames discarded because of an internal error during reception
rxMulticasts	Number of multicast frames delivered to a higher-level protocol or application by this port
rxPeakByteRate	Peak value of ethernetPortByteReceiveRate for this port since the station was last initialized
rxPeakFrameRate	Peak value of ethernetPortFrameReceiveRate for this port since the station was last initialized
rxUnicasts	Number of unicast (nonmulticast) frames delivered by this port to a higher-level protocol or application
txByteRate	Average number of bytes transmitted per second by this port during the most recent sampling period
txBytes	Number of bytes transmitted by this port, including framing characters
txDiscards	Number of transmitted frames discarded because the port was disabled
txErrs	Sum of all transmit errors associated with this port (summary report only)

(continued)

Table 7-1 Description of Fields for Ethernet Port Attributes (continued)

Field	Description
txFrameRate	Average number of frames transmitted per second by this port during the most recent sampling period. Sampling periods are 1 second long and are not configurable.
txFrames	The number of frames transmitted by this port
txInternalErrs	Number of frames discarded because of an internal error during transmission
txMulticasts	Number of multicast frames queued for transmission by a higher-level protocol or application, including those not transmitted successfully
txPeakByteRate	Peak value of ethernetPortByteTransmitRate for this port since the station was last initialized
txPeakFrameRate	Peak value of ethernetPortFrameTransmitRate for this port since the station was last initialized
txQOverflows	The number of frames lost because transmit queue was full
txUnicasts	Number of unicast (nonmulticast) frames queued for transmission by a higher-level protocol or application, including frames not transmitted successfully

*Frame Processing and
Ethernet Statistics*

All frames on the Ethernet network are received promiscuously by an Ethernet port. However, frames may be discarded for the following reasons:

- There is no buffer space available.
- The frame is in error.

Figure 7-1 shows the order in which these discard tests are made.

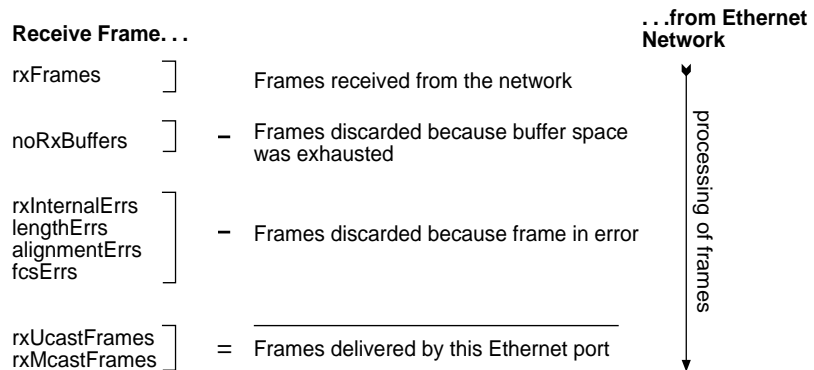


Figure 7-1 How Frame Processing Affects Ethernet Receive Frame Statistics

Frames are delivered to an Ethernet port by bridge and management applications. However, a frame may be discarded for the following reasons:

- The Ethernet port is disabled.
- There is no room on the transmit queue.
- An error occurred during frame transmission.

Figure 7-2 shows the order in which these discard tests are made.

Transmit Frame Statistics...

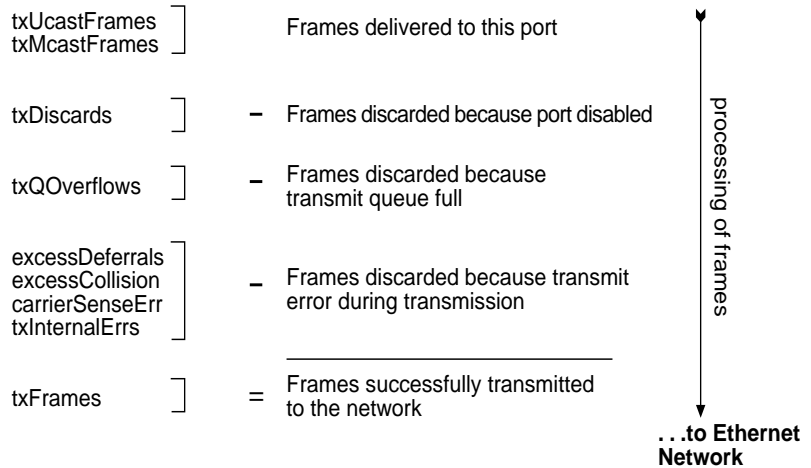


Figure 7-2 How Frame Processing Affects Ethernet Transmit Frame Statistics

Labeling a Port

Port labels serve as useful reference points and as an accurate means of identifying your ports for management. You may want to label your Ethernet ports so that you can easily identify the device specifically attached to each port (for example, LAN, workstation, or server).

To label an Ethernet port:

- 1 From the top level of the Administration Console, enter:
ethernet label
- 2 Enter the port(s) you want to label.
- 3 Enter the label of each Ethernet port.

Port labels can be a maximum of 32 characters in length. The new port label appears the next time you display information for that port.

Top-Level Menu

```
system
└─▶ ethernet
    ├── fddi
    ├── bridge
    ├── ip
    ├── snmp
    ├── analyzer
    ├── script
    └── logout
        ├── summary
        ├── detail
        └─▶ label
            └─▶ portState
```

Setting the Port State

You can enable (place online) or disable (place off-line) Ethernet ports. When an Ethernet port is enabled, frames are transmitted normally over that port. When an Ethernet port is disabled, the port does not send or receive frames.


To enable or disable an Ethernet port:

- 1 From the top level of the Administration Console, enter:
ethernet portState
- 2 Enter the number(s) of the port(s) you want to enable or disable.
- 3 Enter **enable** or **disable** for each Ethernet port.

The *portState* value (shown in the summary and detail displays) reflects online for all enabled ports displayed and off-line for all disabled ports displayed.

Top-Level Menu

```
system
└─▶ ethernet
    ├── fddi
    ├── bridge
    ├── ip
    ├── snmp
    ├── analyzer
    ├── script
    └── logout
        ├── summary
        ├── detail
        └─▶ portState
```



ADMINISTERING FDDI RESOURCES

This chapter describes how to display information about and configure the SuperStack™ II Switch 2200 system and its:

- FDDI station
- FDDI paths
- Media Access Control (MAC)
- FDDI ports



This chapter, which covers advanced FDDI topics, is intended for users familiar with the FDDI MIB. Under normal operating conditions, you do not need to change the FDDI default settings.

For more information about FDDI in the Switch 2200, see the *SuperStack™ II Switch 2200 Operation Guide*.

Administering FDDI Stations

An FDDI station is an addressable node on the network that can transmit, repeat, and receive information. A station contains only one Station Management (SMT) entity and at least one MAC or one port. Stations can be single attachment (one physical connection to the network) or dual attachment (two physical connections to the network).


You can display station information and set the following parameters:

- Connection policies
- Neighbor notification timer
- Status reporting

Displaying Station Information

When you display FDDI station information, you receive information about the station, including its configuration, status reporting, and the most pertinent statistics about general station activity and errors.

Top-Level Menu



- 1 Enter the following from the top level of the Administration Console:

```
fddi station display
```

You are prompted for a station. The Switch 2200 has only one station, which appears in brackets.

- 2 Press [Return].

See the following example of station information:

```

configuration          tNotify statusReporting  connectPolicy
      isolated                30          enabled           0x8000

      ecmState remoteDisconnect  traceMaxExp
              in              false          87500000

                                stationId
                                00-00-00-80-3e-02-95-00
  
```

Table 8-1 describes these statistics.

Table 8-1 Description of Fields for FDDI Station Attributes

Field	Description
configuration	Attachment configuration for the station or concentrator. Values can be Thru, Isolated, Wrap_A, and Wrap_B.
connectPolicy	Bit string representing the connection policies in effect on a station. How connection policies translate into bits is described in Table 8-2. This value can be user-defined.
ecmState	Current state of the ECM state machine
remoteDisconnect	Flag indicating that the station was remotely disconnected from the network as a result of receiving an fddiSMTAction with the value of <i>disconnect</i> in a Parameter Management Frame (PMF). A station requires a Connect Action to rejoin and clear the flag.
stationID	Unique identifier for an FDDI station
statusReporting	Value indicating whether <i>statusReporting</i> is enabled or disabled for the station. This attribute controls whether a station generates Status Report Frames (SRFs) to report events and conditions to network management stations. This value can be user-defined.

(continued)

Table 8-1 Description of Fields for FDDI Station Attributes (continued)

Field	Description
tNotify	Timer used in the Neighbor Notification protocol to indicate the interval of time between the generation of Neighbor Information Frames (NIF). This value can be user-defined.
traceMaxExp	Maximum propagation time for a Trace on an FDDI topology. Places a lower bound on the detection time for an unrecovering ring.

Setting the Connection Policies

The *connectPolicy* attribute is a bit string representing the connection policies in effect on a station. A connection's *type* is defined by the types of the two ports involved (A, B, M, or S) in the connection. You can set the corresponding bit for each of the connection types that you want a particular station to reject.

The Switch 2200 FDDI ports can be of type A or B. By default, all connections to the Switch 2200 FDDI ports are valid, except for M-M connections. The possible connections to reject and their corresponding bits are listed in Table 8-2.

Table 8-2 Bit to Set for Rejecting a Station Connection


This Connection Is Rejected... (Switch port - Remote port)	If This Bit Is Set	Connection Rules
A-A	0	Undesirable peer connection that creates twisted primary and secondary rings; notify station management (SMT)
A-B	1	Normal trunk ring peer connection
A-S	2	Undesirable peer connection that creates a wrapped ring; notify SMT
A-M	3	Tree connection with possible redundancy. The node may not go to Thru state in Configuration Management (CFM). In a single MAC node, Port B has precedence (with defaults) for connecting to a Port M.
B-A	4	Normal trunk ring peer connection

(continued)

Table 8-2 Bit to Set for Rejecting a Station Connection (continued)

This Connection Is Rejected... (Switch port - Remote port)	If This Bit Is Set	Connection Rules
B-B	5	Undesirable peer connection that creates twisted primary and secondary rings; notify SMT.
B-S	6	Undesirable peer connection that creates a wrapped ring; notify SMT.
B-M	7	Tree connection with possible redundancy. The node may not go to Thru state in CFM. In a single MAC node, Port B has precedence (with defaults) for connecting to a Port M.
M-A	12	Tree connection with possible redundancy
M-B	13	Tree connection with possible redundancy
M-S	14	Normal tree connection
M-M	15	Illegal connection that creates a tree of rings topology

To set the connection policies of an FDDI station:

Top-Level Menu

- 1 From the top level of the Administration Console, enter:

```
fddi station connectPolicy
```

You are prompted for a station. The Switch 2200 has one station, which appears in brackets.

- 2 Press Return.
- 3 Enter the value of the connection policy for that station.

The value is a 16-bit number with the appropriate bit(s) set for each connection type that you want to reject.

Example:


```
Select station [1]:
Station 1 - Enter new value [8000]:
```

Setting Neighbor Notification Timer

The *T-notify* attribute is a timer used in the Neighbor Notification protocol to indicate the interval of time between the generation of Neighbor Information Frames (NIF). NIF frames allow stations to discover their upstream and downstream neighbors. The T-notify value has a range of 2 to 30 seconds, with a default value of 30 seconds.

By setting the T-notify value low, your network reacts quickly to station changes, but more bandwidth is used. By setting the T-notify value high, less bandwidth is used, but your network does not react to station changes as quickly.

To set the *T-notify* timer:



- 1 From the top level of the Administration Console, enter:

```
fddi station tNotify
```

You are prompted for a station. The Switch 2200 has one station, which appears in brackets.

- 2 Press Return.
- 3 Enter the value of the *T-notify* timer for that station.

Valid values are 2–30 seconds.


See the following example:

```
Select station [1]:
Station 1 - Enter new value [30]:
```

Enabling and Disabling Status Reporting

The *statusReporting* attribute controls whether a station generates Status Report Frames (SRFs) to report events and conditions to network management stations. By default, status reporting is enabled. If you do not have an SMT management station listening to these event reports or if you use SNMP to monitor FDDI events on all FDDI end-stations, you can set this attribute to disabled so that SRFs will not be generated by the station.

To enable or disable status reporting for a station:



- 1 From the top level of the Administration Console, enter:

```
fddi station statusReporting
```

You are prompted for a station. The Switch 2200 has one station, which appears in brackets.

- 2 Press [Return].
- 3 Enter the new statusReporting value (**enabled** or **disabled**).

See the following example:

```
Select station [1]:
Station 1 - Enter new value (disabled,enabled) [enabled]:
disabled
```

Administering FDDI Paths

FDDI's dual, counter-rotating ring consists of a primary ring and a secondary ring. FDDI stations can be connected to either ring or to both rings simultaneously. Data flows downstream on the primary ring in one direction from one station to its neighboring station. The secondary ring serves as a redundant path and flows in the opposite direction. When a link failure or station failure occurs, the ring "wraps" around the location of the failure, creating a single logical ring.

You can display FDDI path information and set the time values of the following attributes:

- tvxLowerBound
- tmaxLowerBound
- maxTreq

These values are used by all MACs configured in a path.

Displaying Path Information

FDDI path information includes the time values for tvxLowerBound, tmaxLowerBound, and maxTreq, as well as values for ring latency and trace status.

To display FDDI path information:


- 1 From the top level of the Administration Console, enter:

```
fddi path display
```

You are prompted for a station and path. The Switch 2200 has one station, which appears in brackets.

- 2 When prompted for the station, press Return.

Top-Level Menu



3 Enter the path (**p** = primary, **s** = secondary).

See the following example of path information:

```

stn      path      ringLatency      traceStatus
  1      primary          16              0x0
  1      secondary       16              0x0

stn      path      tvxLowBound      tMaxLowBound      maxTReq
  1      primary      2500 us          165000 us         165000 us
  1      secondary    2500 us          165000 us         165000 us

```

Table 8-3 describes these statistics.


Table 8-3 Description of Fields for FDDI Path Attributes

Field	Description
maxTReq	Maximum time value of fddiMACT-Req that will be used by any MAC that is configured in this path. This value can be user-defined.
ringLatency	Total accumulated latency of the ring associated with this path
tmaxLowBound	Minimum time value of fddiMACT-Max that will be used by any MAC that is configured in this path. This value can be user-defined.
traceStatus	Current Trace status of the path
tvxLowBound	Minimum time value of fddiMACTvxValue that will be used by any MAC that is configured in this path. This value can be user-defined.

Setting tvxLowerBound

The *tvxLowerBound* attribute specifies the minimum time value of fddiMAC TvxFValue that will be used by any MAC that is configured onto this path. A MAC uses its valid transmission timer (TVX) to detect and recover from certain ring errors. If a valid frame has not passed through a MAC during the time indicated by fddiMACTvxValue, the MAC reinitializes the ring.

By adjusting the tvxLowerBound value, you specify how quickly the ring recovers from an error. The lower you set this value, the faster the network reacts to problems, but the ring might be reinitialized when there is no problem. The higher you set this value, the less chance of frequent reinitializations, but the network will take longer to recover from errors.



To set tvxLowerBound:

- 1 From the top level of the Administration Console, enter:

```
fddi path tvxLowerBound
```

You are prompted for a station, path, and value. The Switch 2200 has one station, which appears in brackets.

- 2 Press [Return].
- 3 Enter the path (**p** = primary, **s** = secondary).
- 4 Enter the new minimum time value.

The default is 2500 microseconds (μ s).

See the following example:

```
Select station [1]:
Select path(s) (p,s|all) [p]:
Station 1 Primary - Enter new value [2500]:
```

Setting tmaxLowerBound

The *tmaxLowerBound* attribute specifies the minimum time value of fddiMAC T-Max that will be used by any MAC that is configured onto this path. This value specifies the boundary for how high T-Req (the requested token rotation time) can be set.

To set tmaxLowerBound:

- 1 From the top level of the Administration Console, enter:

```
fddi path tmaxLowerBound
```


You are prompted for a station, path, and value. The Switch 2200 has one station, which appears in brackets.

- 2 Press Return.
- 3 Enter the path (**p** = primary, **s** = secondary).
- 4 Enter the new minimum time value.

The default is 165000 microseconds (μ s).

See the example below:

```
Select station [1]:
Select path(s) (p,s|all) [p]: s
Station 1 Primary - Enter new value [165000]:
```




Setting maxT-Req

The *maxT-Req* attribute specifies the maximum time value of fddiMACT-Req that will be used by any MAC that is configured onto this path. T-Req is the value that a MAC bids during the claim process to determine a ring's operational token rotation time, T_Opr. The lowest T-Req bid on the ring becomes T_Opr.

When T_Opr is a low value, the token rotates more quickly, so token latency is reduced. However, more of the ring's available bandwidth is used to circulate the token. Higher values of T_Opr use less bandwidth to circulate the token, but they increase token latency when the ring is saturated.

To set maxT-Req:



- 1 From the top level of the Administration Console, enter:

```
fddi path maxTreq
```

You are prompted for a station, path, and value. The Switch 2200 has one station, which appears in brackets.

- 2 Press [Return].
- 3 Enter the path (**p** = primary, **s** = secondary).
- 4 Enter the new minimum time value.

The default value is 165000 microseconds (μ s)

See the following example:

```
Select station [1]:
Select path(s) (p,s,|all) [p]:
Station 1 Primary - Enter new value [165000]:
```

Administering FDDI MACs

An FDDI MAC uses a token-passing protocol to determine which station has control of the physical medium (the ring). The primary purpose of the MAC is to deliver frames (packets) to their destination by scheduling and performing all data transfers. You can display MAC statistics and configure the following parameters:


- MAC FrameErrorThreshold
- NotCopiedThreshold
- Logical Link Control (LLC) service

Displaying MAC Information

FDDI MAC information can be viewed in a summary or in detail. When you display a summary of various FDDI MAC statistics, you receive information about the MAC, including received and transmitted frames and received and transmitted bytes. The detailed display includes the information in the summary and additional FDDI MAC statistics.

To view the FDDI MAC summary or detailed statistics:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
fddi mac summary
```

OR

```
fddi mac detail
```

You are prompted for a MAC number. The Switch 2200 has only one MAC, which appears in brackets.

- 2 Press [Return].

The following example shows the summary display of FDDI MAC information:

```

rxFrames          txFrames          rxBytes
101884            34320             22689080

txBytes          Errors          noRxBuffers
10257112         0                0

txQOverflows     upstream          downstream
0               00-80-3e-02-95-16  00-80-3e-02-95-41
  
```

The following example shows the detail display of FDDI MAC information:

```

rxFrames          rxBytes          rxFrameRate      rxByteRate
  103666          23089968             36             7582

rxPeakFrameRate  rxPeakByteRate      lostCount         lateCount
      48           10308              0              0

notCopiedCount   notCopiedThresh     notCopiedRatio    notCopiedCond
      0             6550              0             inactive

errorCount       frameErrThresh      frameErrorRatio    frameErrCond
      0             655              0             inactive

noRxBuffers      tvxExpiredCount     rxInternalErrs     rxDiscards
      0              0              0             32923

rxUnicasts       rxMulticasts        txFrames          txBytes
  34621          36158             34921          10437189

txFrameRate      txByteRate          txPeakFrameRate   txPeakByteRate
      15           4511              23             6911

txInternalErrs   txQOverflows        txDiscards         txUnicasts
      0              0              0             34861

txMulticasts     frameCount          tokenCount         ringOpCount
      94           280867          1331364113       4

currentPath      dupAddrTest         duplicateAddr      upstreamDupAddr
  primary         passed              false              false

llcAvailable     llcService          smtAddress
  true            enabled              00-80-3e-02-95-40

                  upstream                downstream
                  00-80-3e-02-95-16          00-80-3e-02-95-41

                  oldUpstream                oldDownstream
                  unknown              00-80-3e-02-95-01

downstreamType   rmtState            tMaxCapab         tvxCapab
  unknown         ring op             1342200 us        1342200 us

tReq             tNeg                tMax              tvxValue
  164986 us      164986 us          167770 us         2621 us

```

Table 8-4 describes the information provided for the FDDI MAC.

Table 8-4 Description of Fields for FDDI MAC Attributes

Field	Description
currentPath	Path on which this MAC is currently located (primary or secondary)
downstream	MAC address of this MAC's downstream neighbor
downstreamType	Indicates the PC type of this MAC's downstream neighbor
dupAddrTest	Pass or fail test for a duplicate address
duplicateAddr	Indicates whether this address is duplicated on the FDDI ring
errorCount	Number of SMT MAC errors.
Errors	The sum of errorCount, lateCount, lostCount, and tvxExpiredCount (summary report only)
frameCount	Number of frames received by this MAC
frameErrCond	Condition is active when the frameErrorRatio is greater than or equal to frameErrorThresh
frameErrorRatio	Ratio of the number lostCount plus the frameErrorCount divided by the frameCount plus lostCount
frameErrThresh	Threshold for determining when a MAC condition report will be generated
lateCount	Number of token rotation timer expirations since this MAC last received a token
llcAvailable	Indicates whether LLC frames can be sent or received on this MAC
llcService	Allows LLC frames to be sent and received on the MAC that is enabled
lostCount	Number of frames and tokens lost by this MAC during reception
noRxBuffers	Number of frames discarded because no buffer space was available
notCopiedCond	Condition is active when the notCopiedRatio is greater than or equal to notCopiedThresh
notCopiedCount	Number of frames that were addressed to this MAC but were not copied into its receive buffers
notCopiedRatio	Ratio of the notCopiedCount divided by copiedCount plus the notCopiedCount
notCopiedThresh	Threshold for determining when a MAC condition report will be generated

(continued)

Table 8-4 Description of Fields for FDDI MAC Attributes (continued)

Field	Description
oldDownstream	Previous value of the MAC address of this MAC's downstream neighbor
oldUpstream	Previous value of the MAC address of this MAC's upstream neighbor
ringOpCount	Number of times that this MAC has entered the operational state from the nonoperational state
rmtState	State of the ring management as defined in SMT
rxByteRate	Average number of bytes received per second by this MAC during the most recent sampling period
rxBytes	Number of bytes received by this MAC, including framing characters
rxDiscards	Number of good frames received by this MAC and discarded before being delivered to a higher-level protocol or application. This count does not include frames that were not received into receive buffers, such as missed frames.
rxFrameRate	Average number of frames received per second by this MAC during the most recent sampling period
rxFrames	Number of frames received by this MAC
rxInternalErrs	Number of frames discarded because of an internal hardware error during reception
rxMulticasts	Number of multicast frames delivered by this MAC to a higher-level protocol or application
rxPeakByteRate	Peak value of fddiMACByteReceiveRate for this MAC since the station was last initialized
rxPeakFrameRate	Peak value of fddiMACFrameReceiveRate for this MAC since the station was last initialized
rxUnicasts	Number of unicast (nonmulticast) frames delivered to a higher-level protocol or application by this MAC
smtAddress	Address of the MAC used for SMT frames
tMax	Maximum value of the target token rotation time
tMaxCapab	Maximum supported target token rotation time this MAC can support
tNeg	Target token rotation time negotiated during the claim process
tokenCount	Number of tokens received by this MAC
tReq	Target token rotation time requested by this MAC

(continued)

Table 8-4 Description of Fields for FDDI MAC Attributes (continued)

Field	Description
txCapab	Maximum time value of the valid transmission timer that this MAC can support
txExpiredCount	Number of times that this MAC's valid transmission timer has expired
txValue	Value of the valid transmission timer in use by this MAC
txByteRate	Average number of bytes transmitted per second by this MAC during the most recent sampling period
txBytes	Number of bytes transmitted by this MAC, including framing characters
txDiscards	Number of frames discarded because LLC Service was not enabled or the FDDI ring was not operational
txFrameRate	Average number of frames transmitted per second by this MAC during the most recent sampling period
txFrames	Number of frames transmitted by this MAC. (Note that this number does not include MAC frames.)
txInternalErrs	Number of frames discarded because of an internal hardware error during transmission
txMulticasts	Number of multicast frames queued for transmission by a higher-level protocol or application, including frames not transmitted successfully
txPeakByteRate	Peak value of fddiMACByteTransmitRate for this MAC since the station was last initialized
txPeakFrameRate	Peak value of fddiMACFrameTransmitRate for this MAC since the station was last initialized
txQOverflows	Number of frames discarded because the transmit queue was full
txUnicasts	Number of unicast frames queued for transmission by a higher-level protocol or application, including frames not transmitted successfully
upstream	MAC address of this MAC's upstream neighbor
upstreamDupAddr	Indicates whether the address upstream of this address is duplicated on the ring

*Frame Processing and
FDDI MAC Statistics*

All frames on the FDDI network are received promiscuously by an FDDI MAC. However, a frame might be discarded for the following reasons:

- There is no buffer space available.
- The frame is in error.

- LLC service is disabled.
- This is an NSA Frame and the A-bit is set.

Figure 8-1 shows the order in which these discard tests are made.

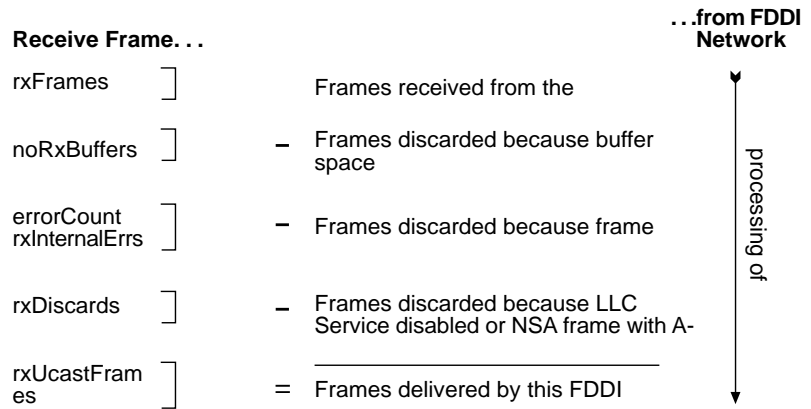


Figure 8-1 How Frame Processing Affects FDDI MAC Receive Frame Statistics

Frames are delivered to an FDDI MAC by bridges and management applications. However, a frame might be discarded for the following reasons:

- LLC Service is disabled.
- The FDDI ring is not operational.
- There is no room on the transmit queue.
- An error has occurred during frame transmission.

Figure 8-2 shows the order in which the discard tests are made.

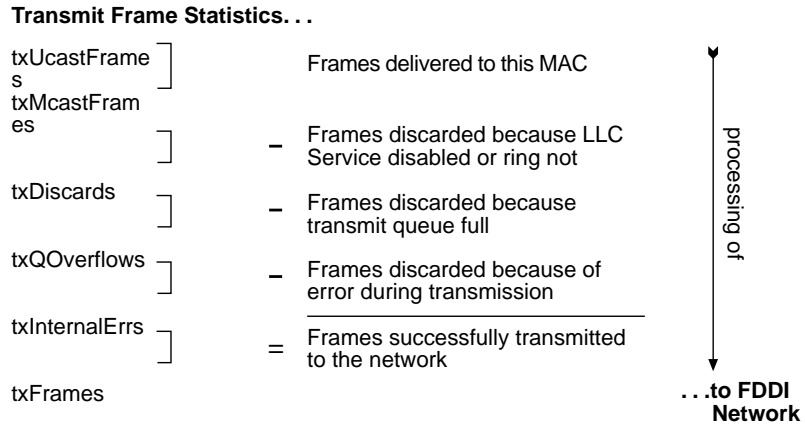



Figure 8-2 How Frame Processing Affects FDDI MAC Transmit Frame Statistics

Setting the Frame Error Threshold

The *FrameErrorThreshold* attribute determines when a MAC condition report is generated because too many frame errors have occurred. A frame error occurs when a frame becomes corrupted. A high error rate often indicates a faulty station on the FDDI ring or a dirty FDDI connector.

Station Management (SMT) monitors the ratio of frame errors to all frames transmitted within a certain period of time. The *FrameErrorThreshold* determines at what percentage the frame errors are significant enough to report to network management. The threshold value is expressed in a percentage based on 65536 (or 100%). For example, to set the threshold at 1%, the value is 655 (the system default). The lower you set the percentage, the more likely SMT will report a problem.

To set the *FrameErrorThreshold*:



- 1 From the top level of the Administration Console, enter:

```
fddi mac frameErrorThreshold
```

You are prompted for a MAC number and new value. The Switch 2200 has one MAC, which appears in brackets.

- 2 Press [Return].
- 3 Enter the new threshold value.

See the following example:

```
Select MAC [1]:
MAC 1 - Enter new value [655]:
```

Setting the Not Copied Threshold

The *NotCopiedThreshold* attribute determines when a MAC condition report is generated because too many frames could not be copied. Not-copied frames occur when there is no buffer space available in the station (which indicates that there is congestion in the station).

SMT monitors the ratio of frames not copied to all frames transmitted within a certain period of time. The *NotCopiedThreshold* determines at what percentage the frames not copied are significant enough to report to network management. The threshold value is expressed in a percentage based on 65536 (or 100%). For example, to set the threshold at 1%, the value is 655 (the system default). The lower you set the percentage, the more likely SMT will report a problem.

To set the *NotCopiedThreshold*:

- 1 From the top level of the Administration Console, enter:


```
fddi mac NotCopiedThreshold
```

You are prompted for a MAC number and new threshold value. The Switch 2200 has one MAC, which appears in brackets.

- 2 Press [Return].
- 3 Enter the new threshold value.

See the following example:

```
Select MAC [1]:
MAC 1 - Enter new value [655]:
```




Enabling and Disabling LLC Service

The Logical Link Control (LLC) service allows LLC frames to be sent and received on the MAC. LLC frames are all data frames transmitted on the network. If there is something wrong on your network, you may want to turn off data (user) traffic for a MAC by disabling LLC service. Although you have disabled data traffic from the MAC, the MAC still participates in neighbor notification and is visible to network management.

To enable or disable LLC service for the MACs in the Switch 2200:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
fddi mac llcService
```

You are prompted for a MAC number and to enable or disable LLC service. The Switch 2200 has one MAC, which appears in brackets.

- 2 Press Return.
- 3 Enter the new MAC value (**enabled** or **disabled**).


See the following example:

```
Select MAC [1]:
MAC 1 - Enter new value (disabled,enabled) [enabled]:
disabled
```

Setting the MAC Paths

The possible backplane path assignments include primary and secondary. To assign MACs to paths:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
fddi mac path
```

You are prompted for a path assignment for the MAC.

- 2 Enter the path.

Administering FDDI Ports

Within an FDDI station, the PHY and PMD entities make up a port. A port (consisting of the PHY/PMD pair that connects to the fiber media) is located at both ends of a physical connection and determines the characteristics of that connection. Each FDDI port is one of four types: A, B, M, or S. You can display port statistics and configure the following port parameters:

- lerAlarm
- lerCutoff
- port labels
- port paths

Displaying Port Information

When you display FDDI port information, you receive information about ports, including the type, path, and port label, as well as other FDDI port statistics, such as error counters.

To view FDDI port information:

- 1 From the top level of the Administration Console, enter:

```
fddi port display
```

You are prompted for a port.

- 2 Enter the port about which you want to view information. Example:

```
port                                portLabel                lemCount
  1                                Backbone1                 0
  2                                SrvrRm001                0
```

```
port    lerEstimate    lerAlarm    lerCutoff    lerCondition
  1             12             7             4            inactive
  2             12             7             4            inactive
```

```
port    lemRejectCount    lctFailCount    ebErrorCount    ebErrorCond
  1              0                0                0                inactive
  2              0                0                0                inactive
```

```
port    lineState    currentPath    connectState    pcmState
  1             qls            isolated        connecting        connect
  2             qls            isolated        connecting        connect
```

```
port    pcWithhold    myType    neighborType    pmdClass
  1             none        A          unknown        multimode
  2             none        B          unknown        multimode
```

Top-Level Menu




Table 8-5 describes the type of information provided for an FDDI port.

Table 8-5 Description of Fields for FDDI Port Attributes

Field	Description
connectState	Connect state of this port (disabled, connecting, standby, or active)
currentPath	Path on which this port is currently located
ebErrorCond	Condition is active when an elasticity buffer error has been detected during the past 2 seconds
ebErrorCount	Number of Elasticity Buffer errors that have been detected
lctFailCount	Number of consecutive times the link confidence test (LCT) has failed during connection management
lemCount	Number of link errors detected by this port
lemRejectCount	Number of times that the link error monitor rejected the link
lerAlarm	The link error rate estimate at which a link connection generates an alarm
lerCondition	Condition is active when the lerEstimate is less than or equal to lerAlarm
lerCutoff	The link error rate estimate at which a link connection is broken
lerEstimate	Average link error rate. It ranges from 10^{-4} to 10^{-15} and is reported as the absolute value of the exponent of the link error estimate
lineState	Line state of this port
myType	Type of port connector on the port
neighborType	Type of port connector at the other end of the physical connection
pcmState	Current Physical Connection Management (PCM) State defined in SMT
pcWithhold	Reason for withholding the connection
pmdClass	Type of PMD entity associated with this port
portLabel	32-character string containing a user-defined name

Setting lerAlarm


The *lerAlarm* attribute is the link error rate (LER) value at which a link connection generates an alarm. If the LER value is greater than the alarm setting, then SMT sends a Status Report Frame (SRF) to the network manager indicating a problem with a port. The *lerAlarm* value is expressed as an exponent (such as 1×10^{-10}). A healthy network has an LER exponent between 1×10^{-10} and 1×10^{-15} . You should set the *lerAlarm* below these

values so that you are only receiving alarms if your network is in poor health. The SMT Standard recommended value is 8.



The `lerAlarm` value must be higher than the `lerCutoff` value so that the network manager will be alerted to a problem before the PHY (port) is actually removed from the network.

To set `lerAlarm`:



- 1 From the top level of the Administration Console, enter:

fddi port lerAlarm

You are prompted for a port number and an estimated link error rate at which the link connection will generate an alarm.

- 2 Enter the port number.
- 3 Enter the estimated link error rate value.


Valid exponent values are -4 through -15. Even though these are negative exponents, enter the value without the negative symbol. For example, to express the value 1×10^{-8} , enter 8 as the value.

Setting `lerCutoff`

The `lerCutoff` attribute is the link error rate estimate at which a link connection is disabled. Once the `lerCutoff` value is reached, the PHY that detected a problem is disabled. The `lerCutoff` value is expressed as an exponent (such as 1×10^{-10}). A healthy network has an LER exponent between 1×10^{-10} and 1×10^{-15} . You should set the `lerCutoff` below these values so that a port will only be removed as a last resort. The SMT Standard recommended value is 7.



The `lerCutoff` value must be lower than the `lerAlarm` value so that the network manager will be alerted to a problem before the PHY (port) is actually removed from the network.



To set the *lerCutoff*:

- 1 From the top level of the Administration Console, enter:

```
fddi port lerCutoff
```

You are prompted for a port number and an estimated link error rate value at which the link connection will be broken.


- 2 Enter the port number.
- 3 Enter the estimated link error rate value.

Valid exponent values are -4 through -15. Even though these are negative exponents, enter the value without the negative symbol. For example, to express the value 1×10^{-7} , enter 7 as the value.

Setting Port Labels

Port labels serve as useful reference points and as an accurate means of identifying your ports for management. You may want to label your FDDI ports for easy identification of the devices attached to them (for example, workstation, server, FDDI backbone).

To label an FDDI port:



- 1 From the top level of the Administration Console, enter:

```
fddi port label
```

You are prompted for a port number and a label value.

- 2 Enter the port number.
- 3 Enter the label value.

Setting the Port Paths

In the Switch 2200 you can assign the A and B ports to either the primary or the secondary path.

To assign ports to paths:


- 1 From the top level of the Administration Console, enter:

```
fddi port path
```

You are prompted for a port.

- 2 Enter the port(s) you want to configure.
- 3 Select the DAS configuration **isol** or **thru** for peer mode at the prompt.
- 4 Select the DAS configuration **isol**, **wrap AB**, or **dualHome** for tree mode at the prompt.

Top-Level Menu



9

SETTING UP THE SYSTEM FOR ROVING ANALYSIS

This chapter describes how to set up the SuperStack™ II Switch 2200 system for roving analysis. With roving analysis, you can monitor Ethernet port activity either locally or remotely using a network analyzer attached to the system.

About Roving Analysis

Roving analysis is the monitoring of Ethernet port traffic for network management purposes. The Administration Console allows you to choose any Ethernet network segment attached to a Switch 2200 system and monitor its activity using a network analyzer (also called a “probe” or “sniffer”). You can monitor port activity locally (when the analyzer and the port are attached to the same Switch 2200 system) or remotely (when the analyzer and the port are on different systems).

You can monitor a port to:

- Analyze traffic loads on each segment so that you can continually optimize your network loads by moving network segments
- Troubleshoot network problems (for example, to find out why there is so much traffic on a particular segment)

When you set up an Ethernet port to analyze, port data that is switched over Ethernet is copied and forwarded to the port on which the network analyzer is attached — without disrupting the regular processing of the packets.

To enable the monitoring of ports on a Switch 2200, take these steps:

- 1 Select an Ethernet port to which you want to attach the network analyzer.
- 2 Select the Ethernet port that you want to monitor (either local or remote). If the port is remote, you must configure it from the Switch 2200 on which

the remote port is located. The remote system must be located on the same FDDI ring as the system to which the analyzer is attached.

Figure 9-1 shows the process for establishing local and remote monitoring of ports.




Figure 9-1 Roving Analysis of Local and Remote Ethernet Ports

Configuration rules

You can have a maximum of 16 network analyzers connected to a system (the maximum number of Ethernet ports on a system) and up to 8 ports monitored per system. The network analyzer cannot be located on the same segment as the port you want to monitor. In general, you will configure one analyzer port ❶ and from there monitor one Ethernet port at a time ❷.

Displaying the Roving Analysis Configuration

You can display the roving analysis configuration to see which ports are designated as analyzer ports and which ports are currently being monitored on a specific system.

When you display the roving analysis configurations for a system, you receive:

- A list of analyzer ports on the system (ports connected to a network analyzer), including the Ethernet port number and the Ethernet MAC address of the port
- A list of ports being monitored on the system, including the Ethernet port number and the Ethernet MAC address of the port to which the *analyzer* is attached

Top-Level Menu

```

system
ethernet  display
fddi      add
bridge   remove
ip        start
snmp     stop
analyzer
script
logout

```

To display the roving analysis configurations, enter the following from the top level of the Administration Console:

```
analyzer display
```

The configurations are displayed as shown in the following example:

```

Ethernet ports configured as analyzer ports:
      Ethernet Port          Address
                9          00-80-3e-0a-3b-02

Ethernet ports being monitored:
      Ethernet Port          Address
                16          00-80-3e-0a-3b-02

```

Adding an Analyzer Port

You can have as many as 16 network analyzers connected to a system (the maximum number of Ethernet ports on a system). For a more accurate analysis, attach the analyzer to a dedicated Ethernet port instead of through a repeater.

To add analyzer ports:

- 1 From the top level of the Administration Console, enter:

```
analyzer add
```

- 2 Press Return to select Ethernet as the port type.
- 3 Enter the number of the Ethernet port to which the network analyzer is attached.

The MAC address of the analyzer port is displayed. You will need this information for setting up the port you want to monitor. See the following example:

```

Select Ethernet port (1-16): 9
Analyzer port address is 00-80-3e-0a-3b-02

```

Port selection errors

If your port selection is not valid, you receive one of the following messages:

```

Error adding analyzer - monitoring already configured on
this port
Error adding analyzer - analyzer already configured on this
port

```

Top-Level Menu

```

system
ethernet  display
fddi      add
bridge   remove
ip        start
snmp     stop
analyzer
script
logout

```

Once the analyzer port is set, it is disabled from receiving or transmitting any other data. Instead, it transmits the data it receives from the monitored port to the network analyzer. If you have enabled Spanning Tree on this port, it is automatically disabled as long as the port is configured for the network analyzer. Once configured, the analyzer port also broadcasts its MAC address so that the address can be learned on remote systems.



If the port configuration changes in the system (that is, if modules are removed or rearranged), the MAC address of the analyzer port remains fixed. If the module with the analyzer port is moved to another slot, then the NVRAM is cleared.

Removing an Analyzer Port

You can change the location of your analyzer port, removing the current port you are using from the roving analysis configuration.

To remove analyzer ports:

- 1 From the top level of the Administration Console, enter:
analyzer remove
- 2 Press Return to select Ethernet as the port type.
- 3 Enter the number of the Ethernet port to which the network analyzer is attached.

The port returns to its current Spanning Tree state and functions as it did before it was set as an analyzer port.

Top-Level Menu

system	display
ethernet	add
fddi	remove
bridge	start
ip	stop
snmp	
analyzer	
script	
logout	

Starting Port Monitoring

After you have a local or remote port configured for the network analyzer, you can start monitoring port activity.



3Com recommends that you ALWAYS configure the analyzer port before configuring the monitored ports.

To start monitoring a new port:

- 1 From the top level of the Administration Console, enter:
analyzer start
- 2 Press Return to select Ethernet as the port type.
- 3 Enter the number of the Ethernet port to monitor.
- 4 Enter the MAC address of the port to which the network analyzer is attached (the port to which the data will be forwarded).



The MAC address of the analyzer port is displayed when you configure that port, and it is also available when you display the roving analysis configurations on the Switch 2200 system to which the analyzer is attached.

See the example below for starting port monitoring:

```
Select port type [Ethernet]:
Select port (1-16): 16
Enter the target analyzer port address: 00-80-3e-0a-3b-02
```

Port selection errors

If your port selection is not valid, you receive one of the following messages:

```
Error starting monitoring - analyzer already configured on
this port
Error starting monitoring - monitoring already configured
on this port
```

MAC address error

If the analyzer port is remote, its MAC address may not be learned on the local system and you receive the following error message:

```
Error starting monitoring - analyzer location unknown
```



CAUTION: *If you receive the above message, check your analyzer port configuration before proceeding. An incorrect configuration will result in frames being continuously flooded throughout your bridged network.*

Top-Level Menu

```
system
ethernet
fddi
bridge
ip
snmp
analyzer
script
logout
display
add
remove
start
stop
```

You are then prompted for an FDDI port through which the data should be forwarded, as shown below:

```
Select FDDI port (1-2): 2
```

Once you successfully configure a port to monitor, all the data received and transmitted on the port is forwarded to the selected analyzer port, as well as processed normally.

Stopping Port Monitoring

Top-Level Menu

```
system
ethernet
fddi      display
bridge   add
ip       remove
snmp    start
* analyzer stop
script
logout
```

After analyzing an Ethernet port, you can remove it from the roving analysis configuration.


To remove a port configured for monitoring:

- 1 From the top level of the Administration Console, enter:

```
analyzer stop
```


- 2 Press Return to select Ethernet as the port type.
- 3 Enter the number of the Ethernet port currently being monitored.

Port data is no longer copied and forwarded from that port to the selected analyzer port.



BRIDGING PARAMETERS

- Chapter 10** Administering the Bridge
- Chapter 11** Administering Bridge Ports
- Chapter 12** Creating and Using Packet Filters
- Chapter 13** Configuring Address and Port Groups to Use in Packet Filters



10

ADMINISTERING THE BRIDGE

This chapter describes how to view the bridge setup and how to configure the following bridge-level parameters:

- IP fragmentation
- IPX snap translation
- Address threshold
- Address aging time
- Spanning Tree Protocol (STP) parameters

For information about configuring the bridge port, see Chapter 11. For information about creating packet filters for a bridge, see Chapter 12.

Displaying Bridge Information

You can display information about the bridge. The display includes bridge statistics (such as topology change information) and configurations for the bridge and Spanning Tree.

To display bridge information, enter the following from the top level of the Administration Console:

```
bridge display
```

Information about the bridge is displayed.

Top-Level Menu

system	
ethernet	display
fdi	ipFragmentation
bridge	ipxSnapTranslation
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	stpGroupAddress
	port
	packetFilter

The following example shows a display of bridge information.

```

      stpState           timeSinceLastTopologyChange
      enabled           1 hr 28 mins 31 secs

      topologyChangeCount
                        2

      topologyChangeFlag  BridgeIdentifier
      false             8000 00803e0f2b00

      designatedRoot     stpGroupAddress     bridgeMaxAge
      0000 0000000000000  01-80-c2-00-00-00     20

      maxAge             bridgeHelloTime     helloTime
      20                 2                   2

      bridgeFwdDelay     forwardDelay        holdTime
      15                  15                 1

      rootCost           rootPort            priority
      0                   No port           0x8000

      agingTime          mode                addrTableSize
      300                 transparent      8191

      addressCount       peakAddrCount       addrThreshold
      95                  107             8000

      ipFragmentation    ipxTranslation      trFddiMode
      enabled             disabled         n/a

      SRBridgeNumber
      n/a

```

Each item in the bridge parameter list is described in Table 10-1.

Table 10-1 Bridge Attributes

Parameter	Description
addressCount	Number of addresses in the bridge address table
addrTableSize	Maximum number of addresses that will fit in the bridge address table
addrThreshold	Reporting threshold for the total number of addresses known on this bridge. When this threshold is reached, the SNMP trap addressThresholdEvent is generated. The range of valid values for setting this object is between 1 and the value reported by the addressTableSize attribute + 1.
agingTime	Time-out period in seconds (between 10 and 32267) for aging out dynamically learned forwarding information. The default value is 300 seconds (or 5 minutes).
bridgeFwdDelay	Forward delay value used when this bridge is the root bridge. This value sets the amount of time a bridge spends in the "listening" and "learning" states. The default value is 15 seconds.
bridgeHelloTime	Hello time value used when this bridge is the root bridge. This value is the time that elapses between the generation of configuration messages by a bridge that assumes itself to be the root. The default value is 2 seconds.
bridgeIdentifier	Bridge identification. It includes the bridge priority value and the MAC address of the lowest numbered port (for example: 8000 00803e003dc0).
bridgeMaxAge	Maximum age value used when this bridge is the root bridge. This value determines when the stored configuration message information is too old and is discarded. The default value is 20 seconds.
designatedRoot	Root bridge identification. It includes the root bridge's priority value and the MAC address of the lowest numbered port on that bridge (for example: 8000 00803e001520).
forwardDelay	The time a bridge spends in the "listening" and "learning" states
helloTime	The time that elapses between the generation of configuration messages by a bridge that assumes itself to be the root
holdTime	Minimum delay time between sending BPDUs (topology change Bridge Notification Protocol Data Units)
ipFragmentation	Configurable parameter that controls whether IP fragmentation is enabled or disabled. The default value is enabled.
ipxTranslation	Configurable parameter that controls whether IPX snap translation is enabled or disabled

(continued)

Table 10-1 Bridge Attributes (continued)

Parameter	Description
maxAge	The maximum age value at which the stored configuration message information is judged too old and discarded. This value is determined by the root bridge.
mode	Operational mode of the bridge. Valid value is <i>transparent</i> for IEEE 802.1d Transparent bridging.
peakAddrCount	Peak value of addressCount
priority	Configurable value appended as the most significant portion of a bridge identifier
rootCost	Cost of the best path to the root from the root port of the bridge (for example, one determining factor of cost is the speed of the network interface — the faster the speed, the smaller the cost)
rootPort	Port with the best path from the bridge to the root bridge
stpGroupAddress	Address that bridge listens to when receiving STP information
stpState	Configurable parameter that provides the state of the bridge (that is, whether Spanning Tree is <i>enabled</i> or <i>disabled</i> for that bridge). The default value is <i>disabled</i> .
timeSinceLast- TopologyChange	Value (in hours, minutes, and seconds) indicating how long since Spanning Tree Protocol last reconfigured the network topology
topologyChange- Flag	Indicates whether a topology change is currently occurring on the bridge (<i>true</i>). A value of <i>false</i> means that no topology change is occurring.
topologyChange- Count	Number of times that Spanning Tree Protocol has reconfigured the network topology

Enabling and Disabling IP Fragmentation

When IP fragmentation is enabled, large FDDI packets are “fragmented” into smaller packets. IP fragmentation allows FDDI and Ethernet stations connected to the Switch 2200 to communicate using IP even if the FDDI stations are transmitting packets that would typically be too large to bridge.

Default value The default value is *enabled*.

To enable or disable IP fragmentation for a bridge:

- 1 From the top level of the Administration Console, enter:

bridge ipfragmentation

- 2 To enable IP fragmentation on a bridge, enter:

enabled

To disable IP fragmentation on a bridge, enter:

disabled

Top-Level Menu

system	display
ethernet	ipFragmentation
bridge	ipxSnapTranslation
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	stpGroupAddress
	port
	packetFilter

Enabling and Disabling IPX Snap Translation

When IPX snap translation is enabled, any 802.3_RAW IPX packets being forwarded from Ethernet to FDDI will be translated to FDDI_SNAP. Likewise, SNAP IPX packets being forwarded from FDDI to Ethernet will be translated to 802.3_RAW packets. When IPX snap translation is disabled, standard (IEEE 802.1H) bridging from 802.3_RAW packets to FDDI_RAW packets is implemented.

Default value The default value is *enabled*.

To enable or disable IPX snap translation for a bridge:

- 1 From the top level of the Administration Console, enter:

bridge ipxSnapTranslation

- 2 To enable IPX snap translation on a bridge, enter:

enabled

To disable IPX snap translation on a bridge, enter:

disabled

Top-Level Menu

system	display
ethernet	ipFragmentaion
bridge	ipxSnapTranslation
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	stpGroupAddress
	port
	packetFilter

Setting the Address Threshold

Address threshold values

The address threshold for a bridge is the reporting threshold for the total number of Ethernet addresses known to the system. When this threshold is reached, the SNMP trap *addressThresholdEvent* is generated.

The range of valid values for this parameter is between 1 and the address table size + 1. Setting the address threshold to one greater than the address table size disables the generation of *addressThresholdEvents* because the limit will never be reached. The default value is 8000.

To set the address threshold:

- 1 From the top level of the Administration Console, enter:
bridge addressThreshold
- 2 Enter the value of the threshold.

Top-Level Menu

```

system
ethernet
└─▶ bridge
    ip
    snmp
    analyzer
    script
    logout
    display
    ipFragmentation
    ipxSnapTranslation
    └─▶ addressThreshold
        agingTime
        stpState
        stpPriority
        stpMaxAge
        stpHelloTime
        stpForwardDelay
        port
        packetFilter
  
```

Setting the Aging Time

Aging time values

The bridge aging time is the maximum period (in seconds) for aging out dynamically learned forwarding information. This parameter allows you to configure the system to age addresses in a timely manner, without increasing packet flooding.

The values can range from 10 to 32,267 seconds. The default value is 300 seconds, which is 5 minutes.

To set the bridge aging time:

- 1 From the top level of the Administration Console, enter:
bridge agingTime
- 2 Enter the aging time value.

Top-Level Menu

```

system
ethernet
└─▶ bridge
    ip
    snmp
    analyzer
    script
    logout
    display
    ipFragmentation
    ipxSnapTranslation
    addressThreshold
    └─▶ agingTime
        stpState
        stpPriority
        stpMaxAge
        stpHelloTime
        stpForwardDelay
        stpGroupAddress
        port
        packetFilter
  
```

Administering STP Bridge Parameters

You can enable or disable Spanning Tree Protocol in the system and set the following STP bridge parameters: priority, maximum age, hello time, and forward delay. For more information about how the Spanning Tree parameters interact at the bridge level to create a loopless network, see Chapter 5: *Transparent Bridging* in the *SuperStack™ II Switch 2200 Operation Guide*.

Enabling and Disabling STP on a Bridge

When Spanning Tree Protocol is disabled, the bridge does not participate in the Spanning Tree algorithm.

The default value is *disabled*.

To enable or disable Spanning Tree Protocol:

- 1 From the top level of the Administration Console, enter:
bridge stpState
- 2 Enter **enabled** or **disabled** at the prompt.

Top-Level Menu

system	
ethernet	display
bridge	ipFragmentation
ip	ipxSnapTranslation
snmp	addressThreshold
analyzer	agingTime
script	stpState
logout	stpPriority
	stpMaxAge
	stpHelloTime
	stpForwardDelay
	stpGroupAddress
	port
	packetFilter

Setting the Bridge Priority

The bridge priority influences the choice of the root bridge and the designated bridge. The *lower* the bridge's priority number, the more likely that the bridge will be chosen as the root bridge or a designated bridge.

Bridge priority values

The bridge priority value is appended as the most significant portion of a bridge identifier (for example: 8000 00803e003dca0). It is a 2-octet value.

Top-Level Menu

system	display
ethernet	ipFragmentation
bridge	ipxSnapTranslation
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	stpGroupAddress
	port
	packetFilter

To configure the STP bridge priority:

- 1 From the top level of the Administration Console, enter:

```
bridge stpPriority
```

- 2 Enter the priority value at the prompt.

If your configuration was successful, you return to the previous menu. If the configuration was not successful, you are notified that your changes failed, and you can try to re-enter those changes.

Setting the Bridge Maximum Age

The bridge maximum age determines when the stored configuration message information is judged too old and discarded from the bridge's memory.

When the Spanning Tree Protocol is configured properly, the maximum age value should ideally never be reached. If the value is too small, then the Spanning Tree Protocol may reconfigure too often, causing temporary loss of connectivity in the network. If the value is too large, the network will take longer than necessary to adjust to a new Spanning Tree configuration after a topology change such as the restarting of a bridge.

*Maximum Age
recommended value*

A conservative value is to assume a delay variance of 2 seconds per hop. The recommended value is 20 seconds.

To configure the bridge max age:

- 1 From the top level of the Administration Console, enter:

```
bridge stpMaxAge
```

- 2 Enter the bridge max age value.

If your configuration was successful, you return to the previous menu. If the configuration was not successful, you are notified that your changes failed, and you can try to re-enter those changes.

Top-Level Menu

system	display
ethernet	ipFragmentation
bridge	ipxSnapTranslation
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	stpGroupAddress
	port
	packetFilter

Setting the Bridge Hello Time

Hello time is the period between the generation of configuration messages by a root bridge. If the probability of losing configuration messages is high, shortening the time makes the protocol more robust. However, lengthening the time lowers the overhead of the algorithm.

*Hello time
recommended value*

The recommended time is 2 seconds.

To configure the bridge hello time:

Top-Level Menu

```

system
ethernet
bridge
ip
snmp
analyzer
script
logout
display
ipFragmentation
ipxSnapTranslation
addressThreshold
agingTime
stpState
stpPriority
stpMaxAge
stpHelloTime
stpForwardDelay
stpGroupAddress
port
packetFilter

```

- 1 From the top level of the Administration Console, enter:

```
bridge stpHelloTime
```

- 2 Enter the bridge hello time value.

If your configuration was successful, you return to the previous menu. If the configuration was not successful, you are notified that your changes failed, and you can try to re-enter those changes.

Setting the Bridge Forward Delay

The forward delay value specifies the amount of time a bridge spends in the “listening” and “learning” states. This value temporarily prevents a bridge from starting to forward data packets to and from a link until news of a topology change has spread to all parts of a bridged network. This delay gives all links that need to be turned off in the new topology time to turn off before new links are turned on.

*Forward delay
recommended value*

The recommended value is 15 seconds.

To configure the forward delay value:

Top-Level Menu

```

system
ethernet
bridge
ip
snmp
analyzer
script
logout
display
ipFragmentation
ipxSnapTranslation
addressThreshold
agingTime
stpState
stpPriority
stpMaxAge
stpHelloTime
stpForwardDelay
stpGroupAddress
port
packetFilter

```

- 1 From the top level of the Administration Console, enter:

```
bridge stpForwardDelay
```

- 2 Enter the forward delay value.

If your configuration was successful, you return to the previous menu. If the configuration was not successful, you are notified that your changes failed, and you can try to re-enter those changes.

Setting the STP Group Address

The STP group address is a single address that bridges listen to when receiving STP information. Each bridge on the network sends STP packets to the group address. Every bridge on the network receives STP packets sent to the group address, regardless of which bridge sent the packets.

Because there is no industry standard on what the group address should be, products from different vendors may respond to different group addresses. If STP does not seem to be working in a mixed-vendor environment, other vendors' products might have different group addresses. In this case, you need to set the STP group address.

To set the STP group address:

- 1 From the top level of the Administration Console, enter:

bridge stpGroupAddress

You are prompted for the new address.

- 2 Enter the group address.

For IBM Spanning Tree Protocol, the group address must be
CO:00:00:00:01:00

Top-Level Menu

system	display
ethernet	ipFragmentation
bridge	ipxSnapTranslation
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	stpGroupAddress
	port
	packetFilter

11

ADMINISTERING BRIDGE PORTS

This chapter describes how to view bridge port information and configure the following:

- Multicast packet threshold
- Spanning Tree Protocol (STP) parameters
- Bridge port addresses

Displaying Bridge Port Information

Bridge port information includes the STP configurations for the bridge port. You can display this information in both summary and detail formats.

To display bridge information:

- 1 From the top level of the Administration Console, enter:

```
bridge port summary
```

OR


```
bridge port detail
```

You are prompted for the port type.

- 2 Enter **Ethernet**, **FDDI**, or **all**.

You are prompted for port number(s).

- 3 Enter the number(s) of the port(s) or **all** to view port parameters for all ports on the bridge.



The following example shows a bridge port summary display.

```

port                rxFrames    rxDiscards    txFrames
Ethernet 1          411180       0              1353766
Ethernet 12         243559       0              1184225

port                portId       stp            state          fwdTransitions
Ethernet 1          0x8003      enabled        forwarding     1
Ethernet 12         0x800e      enabled        forwarding     1

```

The following example shows a bridge port detail display.

```

port                rxFrames    rxBlockedDiscs  rxSameSegDiscs
Ethernet 1          412404      0                0
Ethernet 12         243932      0                0

port                rxErrorDiscs  rxMcastLimit    rxMcastExcDiscs
Ethernet 1          0            0                0
Ethernet 12         0            0                0

port                rxMcastExceeds  rxSecurityDiscs  rxOtherDiscs
Ethernet 1          0            0                0
Ethernet 12         0            0                0

port                rxAllFilters    rxMcastFilters    rxForwardUcasts
Ethernet 1          0              0                  0
Ethernet 12         0              0                  0

port                rxFloodUcasts  rxForwardMcasts  txBlockedDiscs
Ethernet 1          1499           79983             0
Ethernet 12         0              0                 0

port                txMtuExcDiscs  txAllFilters     txMcastFilters
Ethernet 1          0              0                0
Ethernet 12         0              0                0

port                txFrames                portId            stp
Ethernet 1          1357939                0x8003            enabled
Ethernet 12         1187369                0x800e            enabled

port                state          fwdTransitions    priority
Ethernet 1          forwarding     1                  0x80
Ethernet 12         forwarding     1                  0x80

port                pathCost        designatedCost    designatedPort
Ethernet 1          100             10                0x08003
Ethernet 12         100             10                0x0800e

port                SRRingNumber    designatedRoot    designatedBridge
Ethernet 1          n/a             7fff 00803e028e02 8000 00803e0b4800
Ethernet 12         n/a             7fff 00803e028e02 8000 00803e0b4800

```

Table 11-1 describes the type of information provided for the bridge port.

Table 11-1 Bridge Port Attributes

Parameter	Description
designatedBridge	Identification of the designated bridge of the LAN to which the port is attached
designatedCost	Cost through this port to get to the root bridge. The designated cost of the root port is the same as the cost received in incoming BPDUs from the designated bridge for that LAN.
designatedPort	Identification of the designated port on the designated bridge
designatedRoot	Identification of the bridge designated as root
fwdTransitions	Number of times the port has entered forwarding state. This value is useful for checking the stability of a bridged topology. The more transitions in and out of the forwarding state, the more unstable is the topology.
pathCost	Cost to be added to the total path cost when this port is the root port
port	Either Ethernet or FDDI (maximum count: 1 = FDDI and 2–17 = Ethernet)
portId	Identification of the port, which includes the port priority and the port number (for example: 8002)
priority	First factor to determine if a port is to be the designated port when more than one bridge port is attached to the same LAN. If all ports in a bridge have the same priority, then the port number is used as the determining factor.
rxAllFilters	Number of frames discarded because of a user-defined packet filter on the receive all path of this bridge port
rxBlockedDiscs	Number of frames discarded by this port because the receiving bridge port was not in the forwarding state
rxDiscards	Total number of received frames discarded (summary report only)
rxErrorDiscs	Number of frames discarded by this port because of internal bridge system errors (such as hardware and software address table discrepancies)
rxFloodUcasts	Number of unicast frames received on this port that were flooded to one or more ports
rxForwardMcasts	Number of multicast frames received on this bridge port
rxForwardUcasts	Number of unicast frames received on this bridge port

(continued)

Table 11-1 Bridge Port Attributes (continued)

Parameter	Description
rxFrames	Number of frames that have been received by this port from its segment. A frame received on the interface corresponding to this port is only counted by this object if the frame is for a protocol being processed by the local bridging function, including bridge management frames.
rxMcastExcDiscs	Number of multicast frames discarded when rxMcastLimit is exceeded
rxMcastExceeds	Amount of time rxMcastLimit has been exceeded
rxMcastFilters	Number of frames discarded because of a user-defined packet filter on the receive multicast path of this port
rxMcastLimit	Configurable parameter that limits the rate of multicast frames forwarded on a bridge port
rxOtherDiscs	Number of frames discarded by this port because they contained either invalid (group) source addresses or source addresses belonging to this bridge (indicates network loops)
rxSameSegDiscs	Number of frames discarded by this port because the destination address is known on the same network segment as the source address (that is, the frame does not need to be bridged)
rxSecurityDiscs	Number of frames discarded by this port because they contained source addresses that were statically configured on another bridge port (that is, a statically configured station, which is not allowed to move, appears to have moved)

(continued)

Table 11-1 Bridge Port Attributes (continued)

Parameter	Description
state	<p>Spanning Tree state (blocking, listening, learning, forwarding, disabled) in which the port is currently operating:</p> <p><i>Blocking:</i> The bridge continues to run the Spanning Tree algorithm on that port, but the bridge does not receive data packets from the port, learn locations of station addresses from it, or forward packets onto it.</p> <p><i>Listening:</i> The bridge continues running the Spanning Tree algorithm and transmitting configuration messages on the port, but it discards data packets received on that port and does not transmit data packets forwarded to that port.</p> <p><i>Learning:</i> Similar to listening, but the bridge receives data packets on that port to learn the location of some of the stations located on that port.</p> <p><i>Forwarding:</i> The bridge receives packets on that port and forwards or does not forward them depending on address comparisons with the bridge's source address list.</p> <p><i>Disabled:</i> The port has been disabled by management.</p>
stp	Whether the port is <i>enabled</i> or <i>disabled</i> for the Spanning Tree Protocol
txAllFilters	Number of frames discarded because of a user-defined packet filter on the transmit all path of this bridge port
txBlockedDiscs	Number of frames discarded by this port because the transmitting bridge port was not in the forwarding state
txFrames	Number of frames that have been transmitted by this port to its segment. A frame transmitted on the interface corresponding to this port is only counted by this object if the frame is for a protocol being processed by the local bridging function, including bridge management frames.
txMcastFilters	Number of frames discarded because of a user-defined packet filter on the transmit multicast path of this port
txMtuExcDiscs	Number of frames discarded by this port due to an excessive size

*Frame Processing and
Bridge Port Statistics*

All frames received on a physical (Ethernet or FDDI) interface and not explicitly directed to the Switch 2200 are delivered to the corresponding bridge port. A frame is then either forwarded to another bridge port or discarded. A frame might be discarded for the following reasons:

- The destination station is on the same segment as the source station.
- The receive bridge port is blocked.
- There is some problem with the frame.
- A user-defined packet filter indicated that the frame should not be forwarded.

Figure 11-1 shows the order in which the discard decisions are made.

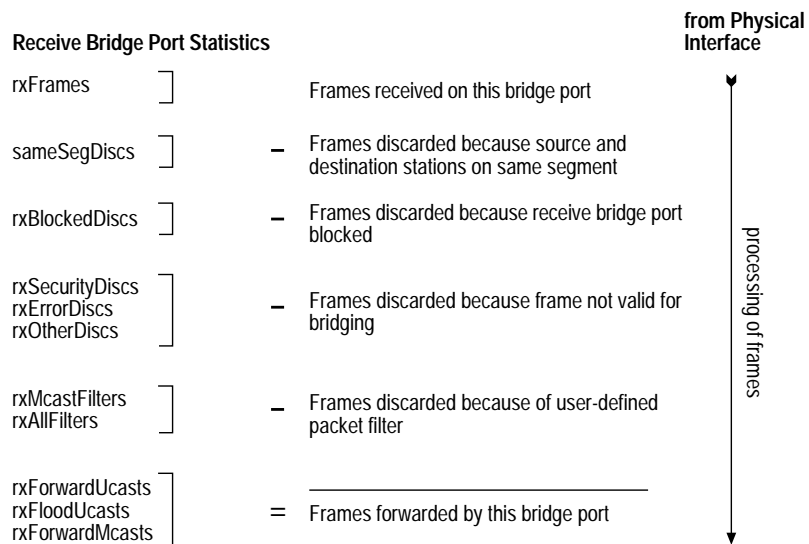


Figure 11-1 How Frame Processing Affects Receive Bridge Port Statistics

A frame forwarded to a bridge port is transmitted onto a physical interface unless it is discarded. A frame might be discarded for the following reasons:

- The transmit bridge port is blocked.
- The frame is too large for the corresponding physical interface.
- A user-defined packet filter indicated that the frame should not be forwarded.

Figure 11-2 shows the order in which the discard decisions are made.

Transmit Bridge Port Statistics

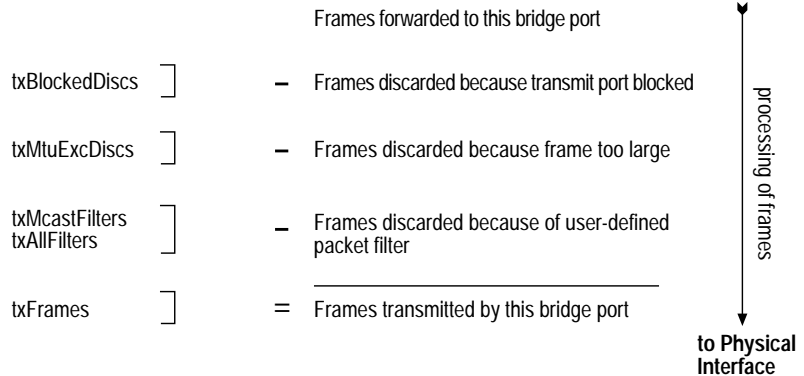


Figure 11-2 How Frame Processing Affects Transmit Bridge Port Statistics

Setting the Multicast Limit

You can assign a multicast packet firewall threshold to a bridge port on the Switch 2200 to limit the forwarding rate of multicast traffic originating on the Ethernet segment connected to the port. For more information about the multicast packet firewall, see Chapter 7: *Bridging Extensions* in the *SuperStack™ II Switch 2200 Operation Guide*.

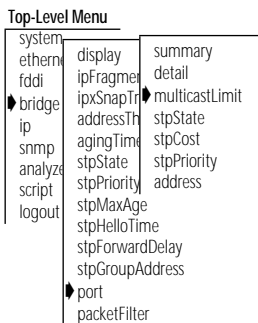
Default value The default is zero (0), which means that no threshold is set.

To set the multicast limit:

- 1 From the top level of the Administration Console, enter:

```
bridge port multicastLimit
```

 You are prompted for port type.
- 2 Enter **Ethernet**, **FDDI**, or **all**.
 You are prompted for port number(s).
- 3 Enter the number(s) of the port(s) or **all** to set the threshold for all ports on the bridge.
 You are prompted for a new value for each port you specified.



- 4 Enter the new multicast threshold value for the port(s).

See the example below:

```
Ethernet port 4 - Enter new value [0]: 400
Ethernet port 5 - Enter new value [0]: 400
```

Administering STP Bridge Port Parameters

You can enable or disable the Spanning Tree Protocol for one or more ports on the system. This only affects the operation of the port if the Spanning Tree Protocol is enabled. You can also set the following STP port parameters: path cost and priority. For more information about how Spanning Tree parameters interact at the bridge-port level, see Chapter 5: *Transparent Bridging in the SuperStack™ II Switch 2200 Operation Guide*.

Enabling and Disabling STP on a Port


You can enable and disable the Spanning Tree Protocol for any port in the system. When STP is disabled for a port but enabled for the entire bridge, a port does not forward frames or participate in the Spanning Tree algorithm. (See page 10-7 for instructions on enabling STP for the entire bridge.) When STP is disabled for a port as well as for the entire bridge, the port will continue to forward frames.

Default value

By default the Spanning Tree state value on a port is the same as the Spanning Tree state value set for the bridge.

To enable or disable STP on a port:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
bridge port stpState
```

You are prompted for the port type.

- 2 Enter **Ethernet**, **FDDI**, or **all**.

You are prompted for the port number(s).

- 3 Enter the number(s) of the port(s) or **all** to enable or disable all ports for the Spanning Tree Protocol.

You are prompted for a new value for each port you specified.

- 4 Enter **enabled** or **disabled** at the prompts.

The following example shows values being set for more than one port:

```
Ethernet port 4 - Enter new value (disabled,enabled)
[enabled]: disabled
Ethernet port 5 - Enter new value (disabled,enabled)
[enabled]: disabled
```

Setting the Port Path Cost

You can set the path cost for a bridge port. The path cost is the cost to be added to the root cost field in a configuration message received on this port. This value is used to determine the path cost to the root through this port. You can set this value individually on each port.

Path cost value

A larger path cost value makes the LAN reached through the port more likely to be low in the Spanning Tree topology. The lower the LAN is in the topology, the less through traffic it will carry. For this reason, you might want to assign a large path cost to a LAN with a lower bandwidth or one on which you want to minimize traffic.


To configure the path cost:

- 1 From the top level of the Administration Console, enter:
bridge port stpCost
You are prompted for the port type.
- 2 Enter **Ethernet**, **FDDI**, or **all**.
You are prompted for the port number(s).
- 3 Enter the number(s) of the port(s) or **all** to configure path cost for all ports on each bridge.
You are prompted for the path cost for each port you specified.
- 4 Enter the path cost for the port(s).

The following example shows values being set for more than one port:

```
FDDI port 1 - Enter new value [100]: 50
Ethernet port 3 - Enter new value [100]: 200
Ethernet port 4 - Enter new value [100]: 200
```

If your configuration was successful, you return to the previous menu. If the configuration was not successful, you are notified that your changes failed, and you can try to re-enter those changes.



Setting the Port Priority

The STP port priority influences the choice of port when the bridge has two ports connected to the same LAN, creating a loop. The port with the lowest port priority will be the one used by the Spanning Tree Protocol.

Port priority value Port priority is a 1-octet value.

To configure the port priority:

- 1 From the top level of the Administration Console, enter:

```
bridge port stpPriority
```

You are prompted for the port type.

- 2 Enter **Ethernet**, **FDDI**, or **all**.

You are prompted for the port number(s).

- 3 Enter the number(s) of the port(s) or **all** to configure the port priority for all ports on each bridge.

You are prompted for the port priority for each port you specified.

- 4 Enter the port priority for the port(s).

The following example shows values being set for more than one port:

```
Ethernet port 3 - Enter new value [0x80]: 1
Ethernet port 4 - Enter new value [0x80]: 500
```

If your configuration was successful, you return to the previous menu. If the configuration was not successful, you are notified that your changes failed, and you can try to re-enter those changes.

Top-Level Menu

system		
ethernet	display	summary
fddi	ipFragm	detail
bridge	ipxSnap	multicastLimit
ip	address	stpState
snmp	agingTim	stpCost
analyze	stpState	stpPriority
script	stpPrior	address
logout	stpMaxAge	
	stpHelloTime	
	stpForwardDelay	
	stpGroupAddress	
	port	
	packetFilter	

Administering Port Addresses


You can administer the MAC addresses of stations connected to Ethernet and FDDI ports on the Switch 2200.

Listing Addresses

You can display MAC addresses currently associated with the selected ports. Each address type (static or dynamic), assigned port, and age are also listed.

To list currently defined MAC addresses:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
bridge port address list
```

You are prompted for the port type.

- 2 Enter **Ethernet**, **FDDI**, or **all**.

You are prompted for the port number(s).

- 3 Enter the number(s) of the port(s) or **all** to display all MAC addresses for the ports you selected.

An example of an address list follows.

Addresses for Ethernet port 1:

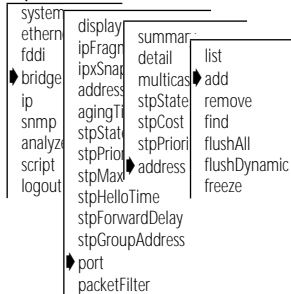
Ethernet address	Type	Age (secs.)
08-00-20-1d-67-e2	Dynamic	219
00-80-3e-02-68-00	Dynamic	219
00-20-af-29-7b-74	Dynamic	219
08-00-02-05-91-c1	Dynamic	219
00-80-3e-02-6d-00	Dynamic	219
00-80-3e-08-5f-00	Dynamic	219
00-80-3e-00-3d-00	Dynamic	219

Adding New Addresses

When you assign new MAC addresses to the selected ports, these addresses are added as statically configured addresses. A statically configured address is never aged and can never be learned on a different Ethernet port.

To add a MAC address:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
bridge port address add
```

You are prompted for the port type.

- 2 Enter **Ethernet** or **FDDI**.

You are prompted for the port number.

- 3 Enter the number of the port.

You are prompted for one or more addresses to add.

- 4 Add each MAC address, pressing [Return] after each entry.


When you finish adding addresses, enter **q** to return to the previous menu.

Removing Addresses

You can remove individual MAC addresses from selected ports.

To remove an address:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
bridge port address remove
```

You are prompted for the port type.

- 2 Enter **Ethernet** or **FDDI**.

You are prompted for the port number.

- 3 Enter the number of the port.

You are prompted for one or more addresses to remove.

- 4 Enter addresses to remove, pressing [Return] after each entry.


Once you have entered all of the addresses to be removed, enter **q** to return to the previous menu.

Flushing All Addresses

You can flush all static and dynamic MAC addresses from the selected port(s). Static MAC addresses are those that you specified using the *add* menu option. Dynamic MAC addresses are those that were automatically learned by the bridge.

To flush *all* addresses:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
bridge port address flushAll
```

You are prompted for the port type.

- 2 Enter **Ethernet, FDDI, or all**.

You are prompted for the port number(s).

- 3 Enter the number(s) of the port(s) or **all**.


All addresses are flushed from the ports you specified.

Flushing Dynamic Addresses

You can flush all dynamic (automatically learned) addresses from the selected port(s).

To flush dynamic addresses:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:

```
bridge port address flushDynamic
```

You are prompted for the port type.

- 2 Enter **Ethernet, FDDI, or all**.

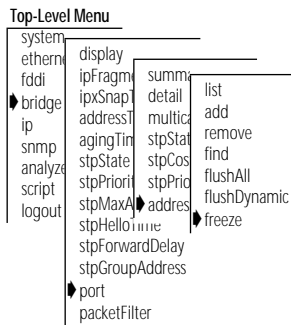
You are prompted for the port number(s).

- 3 Enter the number(s) of the port(s) or **all**.

The addresses are flushed from the address table.

Freezing Dynamic Addresses

You can convert all the dynamic addresses associated with the selected port(s) into static addresses. This conversion is called “freezing” the addresses. Freezing dynamic addresses is a way to improve your network security.



To freeze all dynamic addresses:

- 1 From the top level of the Administration Console, enter:

bridge port address freeze

You are prompted for the port type.

- 2 Enter **Ethernet**, **FDDI**, or **all**.

You are prompted for the port number(s).

- 3 Enter the number(s) of the port(s) or **all**.

The dynamic addresses become static.

12

CREATING AND USING PACKET FILTERS

This chapter describes how to create and edit packet filters using the packet filter language. This chapter also provides instructions for how to:

- List, display, and delete currently defined filters
- Load packet filter definitions created in an ASCII-based editor onto the Switch 2200 system
- Assign filters to ports on the system

About Packet Filtering

Independently configurable packet filtering is provided for the various packet processing paths on each Ethernet port of a Switch 2200. The packet processing paths are defined in Table 12-1.

Table 12-1 Packet Processing Paths

Path	Description
Transmit all	All frames that are transmitted to the segment connected to the port
Transmit multicast	All multicast (including broadcast) frames that are transmitted to the segment connected to the port
Receive all	All frames that are received by the port from the segment connected to the port
Receive multicast	All multicast (including broadcast) frames that are received by the port from the segment connected to the port

When you create a packet filter, you can assign it to the transmit or the receive path of each port, or to both paths.



For additional detailed explanations of packet filter concepts, see Chapter 7: User-defined Packet Filtering in the SuperStack™ II Switch 2200 Operation Guide.

Listing Packet Filters

Top-Level Menu

```

system
ethernet
fddi
bridge
ip
snmp
analyze
script
logout
  display
  ipFragme
  ipxSnapT
  addressT
  agingTim
  stpState
  stpPriority
  stpMaxAg
  stpHelloT
  stpForwal
  stpGroupAddress
  port
  packetFilter
  list
  display
  create
  delete
  edit
  load
  assign
  unassign
  addressGroup
  portGroup

```

When you list the packet filters for the system, the filter identification, filter name (if any), and filter assignments are displayed.

To list the currently defined packet filters, enter the following from the top level of the Administration Console:

```
bridge packetFilter list
```

The listing of packet filters is displayed. An example of the output follows:

```

Ethernet Packet Filters
  Packet Filter 1 - Receive OUI 08-00-1E
    Port 4, Transmit Multicast
    Port 3, Transmit Multicast
    Port 3, Receive Multicast
    Port 5, Receive Multicast
  Packet Filter 2 - Type > 900 or Multicast
    Port 6, Receive All
    Port 8, Transmit All
    Port 8, Receive All
  Packet Filter 3 - Forward IP packets only
  No port assignments

```

In this example, there are two packet filters on the system. The first packet filter has a filter id of 1 and a user-defined name of "Receive OUI 08-00-1E." This filter is loaded onto ports 4, 3, and 5. On port 3, the filter is assigned to both the *transmit multicast* and *receive multicast* paths.

The second filter (filter id 2, user name "Type > 900 or Multicast") is assigned to ports 6 and 8. The filter is assigned to both the *receive all* and *transmit all* paths of port 8.

Displaying Packet Filters

When displaying the contents of a single packet filter, you select the packet filter using the filter id (which you can obtain by listing the packet filters as described in the previous section). The packet filter instructions are displayed; however, any comments in the original packet filter definition file are not displayed because they are not saved with the packet filter.

To display the contents of a packet filter:

- 1 From the top level of the Administration Console, enter:


```
bridge packetFilter display
```

You are prompted for the number of the packet filter you want to display.

- 2 Enter the packet filter id number.

The contents of the packet filter are displayed. An example of the output generated by this command is shown next. The packet filter id and name are displayed, followed by a listing of the packet filter instructions.

```
Select packet filter to be displayed [1-n]: 2
Packet filter 2 - Type > 900 or Multicast
  name "Type > 900 or Multicast"
  pushLiteral.w                0x900
  pushField.w                  12
  gt
  reject
  pushField.b                   0
  pushLiteral.b                0x01
  and
  not
```



Creating Packet Filters

You create custom packet filters by writing a *packet filter definition*. Each packet-processing path on a port may have a unique packet filter definition or may share a definition with other ports. Packet filter definitions are written in the *packet filter language*. This language allows you to construct complex logical expressions.

After writing a packet filter definition, you load it into a Switch 2200 and the corresponding port assignments are preserved in the nonvolatile memory (NVRAM) of the system. This ensures that the packet filter configuration for each system is saved across system reboots and power failures.

Concepts for Writing a Filter

Before writing a packet filter, you should understand these basic concepts:

- How the packet filter language works
- The basic elements of a packet filter
- How to implement sequential tests in a packet filter
- The pre-processed and run-time storage requirements

How the Packet Filter Language Works

You define packet filters using a simple, *stack-oriented* language. Stack-oriented means that the language uses a LIFO (last in, first out) queue when the packet filter is running. The program places values (called operands) on the stack and tests them with various logical expressions (called operators), such as *and*, *or*, *equal*, and *not equal* (see Table 12-3 and Table 12-4). These expressions typically test the values of various fields in the received packet, which include MAC addresses, type fields, IP addresses, and Service Access Points (SAPs).

A program in the packet filter language consists of a series of one or more instructions that results in the top of the stack containing a byte value after execution of the last instruction in the program. This byte value determines whether to forward or discard the packet.

In this stack-oriented language, instructions:

- *push* operands onto the stack
- *pop* the operands from the stack for comparison purposes
- *push* the results back onto the stack

Therefore, with the exception of the push instructions, instructions (such as logical operators) locate their operands implicitly and do not require additional operand specifiers in the instruction stream.

Opcodes are the variables used to identify the type of operands and operators you are specifying in the packet filter instructions.

Table 12-2 describes the instructions and stacks of a packet filter.

Table 12-2 Packet Filter Instructions and Stacks — Descriptions and Guidelines

Element	Descriptions and Guidelines
Instructions	<p>Each instruction in a packet filter definition must be on a separate line in the packet filter definition file.</p>
<i>Instruction format</i>	<p>An instruction consists of an opcode followed by explicit operands and a comment. Although comments are optional, it is recommended that you use them throughout the packet filter for easier administration of the filters. The opcode includes an explicit operand size specification.</p> <p>The general syntax of an instruction is:</p> <pre data-bbox="468 630 1172 656"><opcode>[.<size>] [<operand>...] [# <comment>]</pre> <p>For example:</p> <pre data-bbox="468 716 1300 743">pushliteral.l 0xffffffff00 #load the type field mask</pre> <p>Use any combination of uppercase and lowercase letters for the opcode and size.</p> <p>The contents of a line following the first # (outside a quoted string) are ignored.</p>
<i>Operand sizes</i>	<p>The following operand sizes are supported:</p> <ul style="list-style-type: none"> ■ 1 byte = .b ■ 2 bytes = .w ■ 4 bytes = .l ■ 6 bytes = .a (Included primarily for use with 48-bit, IEEE, globally assigned MAC addresses)
<i>Maximum length</i>	<p>The maximum length for a filter definition is 4096 bytes.</p>
Stack	<p>The packet filter language uses a stack to store the operands that will be used by an instruction and the results of the instruction.</p> <p>Operands are popped from the stack as required by the instructions. An instruction using two or more operands takes the first operand from the top of the stack, with subsequent operands taken in order from succeeding levels of the stack.</p> <p>The stack is a maximum of 64 bytes long, with space in the stack allocated in multiples of 4 bytes. This rule provides for a maximum of 16 operands on the stack.</p> <p>An address size operand (.a) consumes 8 bytes on the stack, decreasing the maximum number of operands on the stack.</p>

Basic Elements of a Packet Filter

Before creating a packet filter, you must decide which part of the packet you want to filter. You can filter Ethernet packets by the destination address, source address, type/length, or some part of the data. You can filter FDDI packets by the destination address, source address, or some part of the data. A packet filter operates on these fields to make filtering decisions. Ethernet and FDDI packet fields are shown in Figure 12-1.




Figure 12-1 Ethernet and FDDI Packet Fields

The Ethernet and FDDI packet fields in Figure 12-1 are used as *operands* in the packet filter. The two simplest operands are described in Table 12-3.

Table 12-3 Packet Filter Operands

Operand	Description	Opcode
packet field	A field in the packet that can reside at any offset. The size of the field can be 1, 2, 4, or 6 bytes. Typically, you only specify a 6-byte field when you want the filter to examine a 48-bit address.	pushField
constant	A literal value to which you are comparing a packet field. As with a field, a constant can be 1, 2, 4, or 6 bytes long.	pushLiteral

The *operators* that you specify in the packet filter allow the filter to make a logical decision about whether the packet should be forwarded or discarded. These operators are described in Table 12-4.

Table 12-4 Packet Filter Operators

Operator	Result	Opcode
equal	true if operand 1 = operand 2	eq
not equal	true if operand 1 \neq operand 2	ne
less than	true if operand 1 < operand 2	lt
less than or equal	true if operand 1 \leq operand 2	le
greater than	true if operand 1 > operand 2	gt
greater than or equal	true if operand 1 \geq operand 2	ge
and	operand 1 bit-wise AND operand 2	and
or	operand 1 bit-wise OR operand 2	or
exclusive or	operand 1 bit-wise XOR operand 2	xor
not	true if operand 1 = false	not
shift left	operand 1 SHIFT LEFT operand 2	shifl
shift right	operand 1 SHIFT RIGHT operand 2	shiftr



*The operators **and**, **or**, and **exclusive or** are bit-wise operators. Each bit of the operands is logically compared to produce the resulting bit.*

Implementing Sequential Tests in a Packet Filter

Filter language expressions are normally evaluated to completion — a packet is accepted if the value remaining on the top of the stack is non-zero. Frequently, however, a single test is insufficient to filter packets effectively. When more tests are warranted, you want to accept a packet that either:

- Satisfies at least one criterion specified in two or more tests (that is, ORs the results of the tests), or
- Satisfies all criteria specified in two or more tests (ANDs the results of the tests)

The *accept* and *reject* instructions are used to implement sequential tests, as shown in Figure 12-2. When using *accept* or *reject*, construct the packet filter so that the tests more likely to be satisfied are performed *before* tests that are less likely to be satisfied.

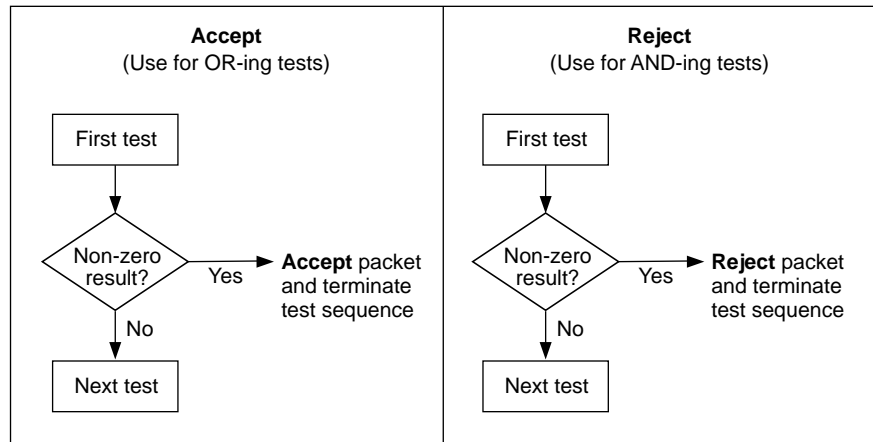


Figure 12-2 Accept and Reject Instructions

The following example shows the use of both accept and reject in a packet filter. This packet filter was created for a network running both Phase I and Phase II AppleTalk.™ The goal of the filter is to eliminate the AppleTalk traffic.

```

Name          "Filter AppleTalk datagrams"
pushField.w   12                # Get the type field.
pushTop       # Make a copy.
pushLiteral   0x809b           # EtherTalk Phase I type
eq            # Test if the packet type is
              # equal to the AppleTalk type.
reject        # Reject the packet and end.
              # Otherwise
pushLiteral.w 0x5dc            # Largest 802.3 packet size
lt            # If this value is less than the
              # value in the packet's
              # type/length field, then this
              # is an Ethernet frame, so
accept        # accept the packet if it is not
              # 802.3, otherwise...
pushField.a   16                # get the SNAP OUI and Ethertype
pushLiteral.a 0x03080007809b   # value to compare.
ne            # If not equal then forward the
              # packet, otherwise drop it

```

Preprocessed and Run-time Storage

A packet filter program is stored in a preprocessed format to minimize the space required by the packet filter definition. When assigned to a port, the packet filter is converted from the stored format to a run-time format to optimize the performance of the filter. Each SuperStack™ II Switch 2200 system is limited to a maximum of 16 packet filter programs.

Preprocessed packet filters Each system provides a maximum of 2048 bytes of nonvolatile storage for *preprocessed* packet filter programs. In the preprocessed stored format:

- A single packet filter program is limited to 254 bytes.
- Each instruction in the packet filter program requires 1 byte for the opcode and size, plus additional bytes for any explicit operands.
- System overhead is 22 bytes, plus a per-packet-filter overhead of 13 bytes. For example, assume a packet filter program requires 200 bytes for storing the instructions in the program. If this packet filter is the only one loaded, the nonvolatile memory required is 22 bytes (for system overhead) plus 13 bytes (for packet filter overhead) plus 200 bytes (for the program itself) — a total of 235 bytes.

Run-time storage of packet filters

For *run-time* storage of packet filter programs, each Switch 2200 system provides a maximum of 8192 bytes. There is no explicit system or per-packet-filter overhead; however, performance considerations can result in unused areas of the run-time storage.

The run-time format is approximately eight times the size of the stored format. Thus a 200-byte packet filter program in stored format expands to approximately 1600 bytes in the run-time format. A single packet filter program cannot exceed 2048 bytes in the run-time format.

Procedure for Writing a Filter

The following steps show the process of writing a packet filter. Detailed examples are provided in the section “Examples of Creating Filters” on page 12-11.

You write the instructions for the packet filter using the following syntax:

```
<opcode>[.<size>]  [<operand>...] [# <comment>]
```

The opcode descriptions are in the section Appendix A: *Packet Filter Opcodes, Examples, and Syntax Errors*. The description of the supported operand sizes can be found in Table 12-2. The operand value is determined by what you are testing (for example, an address or a length).



Implicit operands for an instruction must be of the size expected by the instruction. Any mismatch in implicit operand size results in an error “operand size mismatch” when you load the program into the system.

When writing a packet filter, be sure that you use comments (preceded by #) to describe each step in the filter. This habit will help you to revise filters in the future and enable others to understand and use the filters you create.

To write a packet filter:

- 1 Assign a unique, descriptive name to the filter using the `NAME` opcode.
- 2 Specify what to test. For example, use the `PUSHFIELD` opcode to select a field in the packet.
- 3 Specify what to compare to the value in step 2. For example, use the `PUSHLITERAL` opcode to select a constant value.

- 4 Apply a logic operation to the values in steps 2 and 3. The operator you use depends on what comparison you want to make.

Variations on these four basic steps of writing packet filters include:

- Use `pushTop` for each additional comparison you intend to make with the `pushField` value. This opcode makes a duplicate of the `pushField` value and places it on top of the original `pushField` on the stack. The `pushtop` instruction makes a copy of the field more efficiently than if you use a second `pushfield` instruction.
- Use `accept` or `reject` with `and` and `or` operators when you have sequential tests and you would like the filter to accept or reject a packet before the entire expression has been evaluated. Using `accept` and `reject` can significantly improve the performance of certain types of filters. See the section “Implementing Sequential Tests in a Packet Filter” on page 12-8 for more information.
- Use `pushSAGM`, `pushDAGM`, `pushSPGM`, or `pushDPGM` for filtering by address or port groups. See Chapter 13 for more information.

Examples of Creating Filters

The following example shows a complex packet filter built from three simple packet filters. Each of the shorter, simpler packet filters can be used on its own to accomplish its own task. Combined, these filters create a solution for a larger filtering problem.

Filtering Problem

Your network contains market data feed servers that receive time-critical financial data needed for trading floor applications. At the center of the trading floor networks is a Switch 2200 system that is being used to switch Ethernet traffic and to concentrate the market data feed servers onto the FDDI departmental backbone.

The difficulty is that the market data feed servers transmit data to users with broadcast packets that are forwarded to all stations on all segments attached to the Switch 2200 system. Not all of the segments attached to the Switch 2200 system have stations that require these broadcast updates. In order to optimize the performance of these Ethernet segments, you need to filter the broadcasts.

Packet Filter Solution

The solution described here is to create a highly sophisticated packet filter that prevents only the broadcast packets from the market data servers from being forwarded onto the segments that are not part of an active trading floor.

Before writing the packet filter, it is important to understand the functions that the filter must provide. The broadcast packets that are transmitted by the servers are based on either TCP/IP or XNS protocol. In both cases, the broadcast packets have socket values that are greater than 0x076c and less than 0x0898. The socket value is located 24 bytes into the packet in IP datagrams and 30 bytes into the packet in XNS datagrams.

You can use this information to create pseudocode that simplifies the process of writing the actual filter. It helps to write the pseudocode in outline form, as shown here:

- 1 Determine if the packet has a broadcast address. (Use the packet filter path assignment.)
- 2 Determine if the packet is an XNS datagram.
- 3 Check socket values and discard the packet if:
 - a The socket value is greater than or equal to 0x76c
AND
 - b The socket value is less than 0x898
- 4 Determine if the packet is an IP datagram.
- 5 Check socket values and discard the packet if:
 - a The socket value is greater than or equal to 0x76c
AND
 - b The socket value is less than 0x898
- 6 End the filter.

The pseudocode translates into the following packet filter:

```

Name      "IP XNS ticker bcast filter"
          # Assign this filter in the multicast path
          # of a port only--this is very important
          #
          # XNS FILTERING SECTION
          #
pushField.w      12      # get the type field of the packet and
pushLiteral.w    0x0600  # place it on top of the stack.
eq               # put the type value for XNS on top of
                 # the stack.
pushLiteral.w    0x76c   # if the two values on the top of the
ge               # stack are equal, then return a non-zero
                 # value.
pushField.w      30      # put the lowest socket value on top of
                 # the stack.
pushField.w      30      # put the value of the socket from the
ge               # packet on top of the stack.
                 # compare if the value of the socket is
pushLiteral.w    0x0898  # greater than or equal to lower bound.
                 # put the highest socket value on top of
pushField.w      30      # the stack.
lt               # put the value of the socket from the
                 # packet on top of the stack.
and              # compare if the value of the socket is
                 # less than the upper bound
and              # "and" together with "ge" and "lt" test
                 # to determine if the socket value is
                 # "within" the range. If it is, a "one"
                 # will be placed on the stack.
and              # compare if XNS & in range
                 #
                 # IP FILTERING SECTION
                 #
pushField.w      12      # get the type field of the packet and
pushLiteral.w    0x0800  # place it on top of the stack.
eq               # put the type value for IP on top of
                 # the stack.
eq               # if the two values on the top of the
                 # stack are equal, then return a non-zero
                 # value.
pushLiteral.w    0x76c   # put the lowest socket value on top of
                 # the stack (1900).
pushField.w      24      # put the value of the socket from the
ge               # packet on top of the stack.
                 # compare if the value of the socket is
pushLiteral.w    0x0898  # greater than or equal to lower bound.
                 # put the highest socket value on top of
pushField.w      24      # the stack (2200).
lt               # put the value of the socket from the
                 # packet on top of the stack.
and              # compare if the value of the socket is
                 # less than the upper bound
and              # "and" together with "ge" and "lt".
                 # test to determine if the socket value is
                 # "within" the range. If it is in range, a
                 # "one" will be placed on the stack.
and              # compare if IP and in range.
or               # determine if the type field is either
                 # XNS or IP.
not              # discard if (IP & in range) and (XNS & in
                 # range).

```

The rest of this section concentrates on the parts of the filter, showing you how to translate the pseudocode's requirements into filter language. The large filter on page 12-13 is broken down into subsets to show how you can create small filters that perform one or two tasks, and then combine them for more sophisticated filtering. Table 12-5 shows how the purpose of each pseudocode step is accomplished in the small series of packet filters.

Table 12-5 Pseudocode Requirements Mapped to the Packet Filter

Step	Accomplished Through...
1	The path to which you assign the packet filter. For administrative purposes, this path is specified in the first two comment lines in the filter definition. The filter must be assigned to a multicast path to filter packets with broadcast addresses.
2	Packet Filter One — Forwarding XNS packets
3	Packet Filter Two — Checking for specified socket range
4 & 5	Combining a Subset of Filters — Forwarding IP packets within specified socket range

Packet Filter One. This filter is designed to forward XNS packets. These steps show how to create this filter.

1 Name the filter:

```
"Forward only XNS packets"
```

It is important to distinguish the function of each filter when it is loaded onto a Switch 2200 that has more than one filter stored in memory. Naming is also useful for archiving filters on an ftp server so that the filters can be saved and loaded on one or more Switch 2200 systems.

2 Enter executable instruction #1:

```
pushField.w 12 # get the type field of the packet and  
# place it on top of the stack
```

3 Enter executable instruction #2:

```
pushLiteral.w 0x0600 # put the type value for XNS on top  
# of the stack
```

4 Enter executable instruction #3:

```
eq # if the two values on the top of the stack are equal,  
   # then return a non-zero value
```

Packet Filter Two. This filter is designed to accept packets within the socket range of 0x76c and 0x898. These steps show how to create this filter:

1 Name the filter:

```
"Socket range filter"
```

2 Enter executable instruction #1:

```
pushLiteral.w 0x76c # put the lowest socket value on top  
                  # of the stack
```

3 Enter executable instruction #2:

```
pushField.w 30 # put the value of the socket from the  
              # packet on top of the stack
```

4 Enter executable instruction #3:

```
ge # compare if the value of the socket is greater than  
   # or equal to the lower bound
```

5 Enter executable instruction #4:

```
pushLiteral.w 0x0898 # put the highest socket value on  
                   # top of the stack
```

6 Enter executable instruction #5:

```
pushField.w 30 # put the value of the socket from the  
              # packet on top of the stack
```

7 Enter executable instruction #6:

```
lt # compare if the value of the socket is less than the  
   # upper bound
```

8 Enter executable instruction #7:

```
and # "and" together with "ge" and "lt" test to determine  
    # if the socket value is "within" the range. If it is,  
    # a "one" will be placed on the stack.
```


Combining a Subset of the Filters. The next filter accepts IP packets with a socket range of 0x76c (1900) and 0x898 (2200). The filter combines packet filters one and two, modifying them for IP. These steps show how to create this filter.

1 Name the filter:

```
"Only IP pkts w/in socket range"
```

2 Perform steps 2 through 4 as described in "Packet Filter One" on page 12-14, except give the pushLiteral instruction (in step 3) a value of 0x0800 for IP.

3 Perform steps 2 through 8 as described in "Packet Filter Two" on page 12-15, except the socket value for IP (in step 3) is located 24 bytes into the packet (instead of 30 as for XNS).

4 Add an *and* statement to compare the results of step 2 with the results of step 3:

```
and # compare if IP and in range
```

This combination looks like this:

```
Name      "Only IP pkts w/in socket range"
pushField.w    12      # get the type field of the packet and
                  # place it on top of the stack
pushLiteral.w  0x0800  # put the type value for IP on top of
                  # the stack
eq            # if the two values on the top of the
                  # stack are equal, then return a non-zero
                  # value
pushLiteral.w  0x76c   # put the lowest socket value on top of
                  # the stack (1900)
pushField.w    24      # put the value of the socket from the
                  # packet on top of the stack
ge            # compare if the value of the socket is
                  # greater than or equal to the lower bound
pushLiteral.w  0x0898  # put the highest socket value on top of
                  # the stack (2200)
pushField.w    24      # put the value of the socket from the
                  # packet on top of the stack
lt            # compare if the value of the socket is
                  # less than the upper bound
and           # "and" together with "ge" and "lt" test
                  # to determine if the socket value is
                  # "within" the range. If it is in range, a
                  # "one" will be placed on the stack.
and           # compare if IP and in range
```

Combining All the Filters. Together, the four packet filters work to perform the solution to the problem: filtering the broadcast packets from the market data servers. These steps show how to create this filter:

- 1 Name the filter:

```
"Discard XNS & IP pkts w/in socket range"
```

- 2 Perform steps 2 through 4 as described in "Packet Filter One" on page 12-14.
- 3 Perform steps 2 through 8 as described in "Packet Filter Two" on page 12-15.
- 4 Add an *and* statement to compare the results of step 2 and the results of step 3:

```
and # compare if XNS & in range
```

- 5 Perform steps 2 through 4 as described in "Combining a Subset of the Filters" on page 12-16.

- 6 Add an *or* statement:

```
or # determine if the type field is either XNS or IP
```

- 7 Add a *not* statement to discard any matching packets:

```
not # discard if (IP & in range) & (XNS & in range)
```

The complete packet filter that discards IP and XNS packets that are within the specified range is shown on page 12-13.

Tools for Writing a Filter

You can create a new packet filter using either an ASCII-based text editor (such as *EMACS* or *vi*) or the line editor built into the Administration Console. Using an ASCII-based text editor allows you to create multiple copies of the packet filter definition, which you can then copy onto one or more Switch 2200 systems from a networked workstation. This method also allows you to archive copies of filter definitions.

Using the Built-in Line Editor


The built-in text editor provides a minimal set of editing functions that you can use to edit a packet filter definition one line at a time. A single line is limited to no more than 79 characters. The number of lines is limited only by available memory.



The maximum length of a packet filter definition is 4096 bytes.

The editor assumes a terminal capability no higher than a glass tty (that is, it does not assume an addressable screen). You can place any ASCII printable character into the editing buffer at the cursor position. If a character exceeds the maximum line length, the character is discarded and a bell sounds. The editor initially operates in *insert* mode. The commands supported by the editor are summarized in Table 12-6.

To use the built-in line editor to create a packet filter definition:



- 1 From the top level of the Administration Console, enter:

```
bridge packetFilter create
```

The packet filter line editor appears.

- 2 Enter the definition for the packet filter. See the command in Table 12-6.
- 3 Save the packet filter by pressing Ctrl+W.

The syntax of the filter definition is checked. If any errors are detected, the errors are displayed and the editor is re-entered at the line containing the first error. After correcting the errors, attempt to save the packet filter again.

After you have corrected all errors and successfully saved the packet filter, it is converted to internal form and stored on the system.

Table 12-6 Packet Filter Editor Commands

Command	Keys	Description
List buffer	Ctrl+l	Displays each of the lines in the editing buffer and then redisplayes the line currently being edited
Next Line	Ctrl+n	Moves cursor to next line; positions cursor at start of line
Previous Line	Ctrl+p	Moves cursor to previous line; positions cursor at start of line
Start of Line	Ctrl+a	Moves cursor within a line to the start of the present line
End of Line	Ctrl+e	Moves cursor within a line to the end of the present line
Left 1 Character	Ctrl+b	Moves cursor <i>left</i> one character within a line
Right 1 Character	Ctrl+f	Moves cursor <i>right</i> one character within a line
Insert Line	Enter	Inserts a new line. The new line becomes the current line, with the cursor positioned at the start. If the cursor is positioned over the first character on a line when you press Enter, a blank new line is inserted prior to the current line. Otherwise, the current line is split at the cursor position, with the current line retaining the characters before the cursor, followed by the new line containing the remainder of the characters.
Delete Previous Character	Ctrl+h	Deletes a single character preceding the cursor and shifts the remainder of the line <i>left</i> one position
Delete Current Character	Ctrl+d	Deletes a single character under the cursor and shifts the remainder of the line <i>left</i> one position
Delete Line	Ctrl+k	Deletes the remainder of the line from the current cursor position. If the cursor is positioned over the first character, all of the characters on the line are deleted, but the line is retained. A second Delete Line command removes the line from the edit buffer.
Insert/Overstrike Toggle	Ctrl+o	Toggles between the insert mode and overstrike mode
Write Changes	Ctrl+w	Writes (saves) the current contents of the edit buffer into the packet filter definition. No syntax checking of the definition is performed at this point other than to verify that the length of the source is within the maximum limits. If the source is too long, the message <code>Error: Edit buffer exceeds maximum length</code> is displayed. The contents of the edit buffer are unaffected; however, the packet filter definition contains only those lines that fit entirely within the length limitation.
Exit Editor	ESC	Allows you to leave the editor. You receive a warning if the edit buffer has not been successfully written since the last modification. You can either discard the changes or return to the editor. Note that only those changes made since the last Write Changes command are discarded.

Using an External Text Editor

To use an ASCII-based editor to create a packet filter:

- 1 Create the definition in a text file.
- 2 From a networked workstation, ftp the file to the Switch 2200 on which you want to load the filter.
- 3 Load the filter as described in "Loading Packet Filters" on page 12-22.

Deleting Packet Filters

Deleting a packet filter removes the filter from the Switch 2200 system.

To delete a packet filter:

- 1 From the top level of the Administration Console, enter:
bridge packetFilter delete
- 2 Enter the id of the filter to delete. To find the id of the filter, list the filters as described in "Listing Packet Filters" on page 12-2.
You are prompted to confirm the deletion.
- 3 Enter **y** (yes) to delete or **n** (no) to return to the previous menu.

Top-Level Menu

```


system
ethernet
fddi
bridge
ip
snmp
analyze
script
logout
  display
  ipFragme
  ipxSnapT
  addressT
  agingTim
  stpState
  stpPriority
  stpMaxAc
  stpHelloT
  stpForwa
  stpGroupAddress
  port
  packetFilter
  list
  display
  create
  delete
  edit
  load
  assign
  unassign
  addressGroup
  portGroup

```

Editing, Checking and Saving Packet Filters

You can use the Switch 2200 system line editor to edit packet filters. Once you save the packet filter, it is checked for syntax errors. The Switch 2200 system software will not allow you to assign the packet filter to a port until the filter is error-free.

You can also edit a packet filter using an ASCII-based text editor such as *EMACS* or *vi*. You can then use ftp to send the filter text to the Switch 2200 system from a networked workstation.



To edit a packet filter using the Switch 2200 system line editor:

- 1 From the top level of the Administration Console, enter:
bridge packetFilter edit
- 2 Enter the packet filter id number.
Specifying a filter id loads that filter into the edit buffer.
- 3 Edit the filter. For more information, see the section “Using the Built-in Line Editor” on page 12-17.
- 4 Press [Esc] to exit the line editor.
- 5 At the `Edit buffer has been changed. Quit anyway?` prompt, enter **y** (yes) to end the editing session or **n** (no) to return to editing.
- 6 You have three choices of what to do next:
 - To overwrite the existing filter with the contents of the edit buffer, enter **y** at the `Replace existing filter?` prompt.
 - To store the definition as a new filter, enter **n** at the `Replace existing filter?` prompt and **y** at the `Store as new filter?` prompt. The packet filter is assigned a number.
 - To exit from the editor without saving changes, enter **n** at both prompts.

Correcting errors in a packet filter

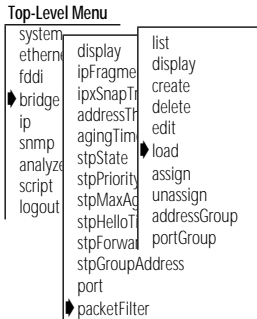
When you save a packet filter edited with the built-in text editor, the system checks the syntax of the filter definition. If any errors are detected, the errors are displayed and the editor is re-entered at the line containing the first error. After correcting the errors, you must exit the editor and attempt to save the packet filter again.

After you have corrected all errors and saved the packet filter, it is converted to internal form and updated on the system.

Loading Packet Filters

When you create packet filters using an external text editor, you must load the filters onto the system from the network host on which you created them. Once loaded, the packet filter definition is converted into the internal format that is used by the packet filter code in the system.

To load a packet filter:



- 1 From the top level of the Administration Console, enter:

```
bridge packetFilter load
```

You are prompted for a host IP address, file path name, user name, and password. To use the value in brackets, press [Return] at any prompt.

- 2 Enter the host IP address.
- 3 Enter the path name.
- 4 Enter your user name.
- 5 Enter your password.

The packet filter is loaded onto the Switch 2200.

Any syntax errors in the packet filter definition are reported to you at this time. See Appendix A: *Packet Filter Opcodes, Examples, and Syntax Errors* for a description of these errors. If errors are detected, you are offered the option of editing the filter definition or terminating the load.

The load might fail if the system has insufficient nonvolatile RAM to store the filter. In this case, an error message tells you that the system did not accept the load.

Assigning Packet Filters to Ports


To assign a packet filter to one or more ports, the packet filter must reside on the system. Each path (transmit all, transmit multicast, receive all, and receive multicast) of a port can have only one packet filter assigned to it; however, you can assign a single packet filter to multiple paths and ports.

Packet filter path assignments

Placing a filter on the transmit path confines the packet to the segment it originated from if it does not meet the forwarding criteria. Placing a filter on the receive path prohibits a packet from accessing certain segments unless

it meets the forwarding criteria. A packet that does not meet the forwarding criteria defined in the filter is discarded.

To assign a packet filter:



- 1 From the top level of the Administration Console, enter:
bridge packetFilter assign
- 2 Enter the id number of the packet filter to be assigned. To get the id of the packet filter, you can list all packet filters as described in “Listing Packet Filters” on page 12-2.
- 3 Enter the port type (**Ethernet, FDDI, all**).
- 4 Enter the port(s) to assign the filter.
- 5 Enter the path(s) you want to place the filter (**txA, txM, rxA, rxM, all**).

In this example, the assignment is to the transmit all (txA) path and the receive all (rxA) path on port 1.

```

Select filter [1-n]: 1
Select port type(s) (Ethernet,FDDI|all) [Ethernet,FDDI]:
FDDI
Select port(s) (1-16|all) [1-16]: 1
Select path(s) (txA,txM,rxA,rxM|all): txA,rxA

```

The ports are limited to those that have at least one path unassigned, while the paths are limited to those that are unassigned. Because you can specify multiple selections at each level, you can assign a wildcard that attempts to assign the filter to the set indicated by the ports and paths taken in combination.



One or more assignments might fail because of a previous assignment.

Unassigning Packet Filters from Ports

Top-Level Menu

```

system
ethernet
fddi
bridge
ip
snmp
analyze
script
logout
  display
  ipFragme
  ipxSnapTr
  addressTr
  agingTim
  stpState
  stpPriority
  stpMaxAg
  stpHelloT
  stpForwar
  stpGroupAddress
  port
  packetFilter
  list
  display
  create
  delete
  edit
  load
  assign
  unassign
  addressGroup
  portGroup

```

To unassign a packet filter from one or more ports, the packet filter must have been previously assigned to at least one port.

To unassign a packet filter:

- 1 From the top level of the Administration Console, enter:
bridge packetFilter unassign
- 2 Enter the id number of the packet filter to unassign.
- 3 Enter the port type (**Ethernet, FDDI, all**).
- 4 Enter the port number(s) of the packet filter to unassign.
- 5 Enter the path(s) of the packet filter to unassign.

An example of unassigning a packet filter is shown next. In this example, the unassignment is from the transmit all (txA) paths on port 1.

```

Select filter [1-n]: 1
Select port type(s) (Ethernet,FDDI|all) [Ethernet,FDDI]: FDDI
Select port(s) (1-16|all) [1-16]: 1
Select path(s) (txA,txM,rxA,rxM|all) [txA,rxA]: txA

```

Because you can specify multiple selections at each level, you can assign a wildcard that attempts to unassign the filter from the set indicated by the ports and paths taken in combination.



One or more of the unassignments might fail if the filter is not assigned.

13

CONFIGURING ADDRESS AND PORT GROUPS TO USE IN PACKET FILTERS

This chapter describes how to use address and port groups as filtering criteria in a packet filter, and how to administer address and port groups.

Using Groups in Packet Filters

You can use address groups (a list of MAC addresses) and port groups (a list of Switch 2200 Ethernet and FDDI ports) as filtering criteria in a packet filter.



For more information about address and port group concepts, see Chapter 7: User-defined Packet Filtering in the SuperStack™ II Switch 2200 Operation Guide.

A packet filter uses a group to make filtering decisions by accessing the group's source group mask and destination group mask. You reference these group masks using the opcodes SAGM (source address group mask), DAGM (destination address group mask), SPGM (source port group mask), and DPGM (destination port group mask). Here are some examples of using address and port groups in packet filters.

Address group packet filter example

In this example, the filter only forwards packets among stations that are within the same address group.

```
Name      "Accept Same Source and Destination"
pushSAGM           # Get source address group mask
pushDAGM           # Get destination address
                  # group mask
and                # Compare if source address and
                  # destination address are common
                  # members of an address group (result
                  # is either zero or non-zero)
pushLiteral.1     0      # Put a zero on the stack
ne                 # If not equal, returns a "one" to
                  # stack, resulting in packet
                  # forwarded
```

Port group packet filter example

In this example, packets are not forwarded to ports in groups 3 and 8.

```
Name      "Discard Groups 3 and 8"
pushSPGM           # Get source port group mask
pushLiteral.1     0x0084 # Select bits 3 and 8
and                # If port group bits 3 & 8 are common
                  # with SPGM, then non-zero value is
                  # pushed onto stack
pushLiteral.1     0      # Push zero
eq                # Only if SPGM is not in port groups
                  # corresponding to bits 3 & 8, then
                  # packet is forwarded
```

In the Administration Console you can:

- List the groups
- Display specific information about a group
- Create a new group
- Delete a group
- Copy a group from one module to another (address groups only)
- Add and remove addresses and ports to or from a group

Listing Groups

You can list the address and port groups currently defined for your Switch 2200 system. The group id, group name (if any), and group mask are displayed.

Top-Level Menu

```
system
ethernet
fdi
bridge
ip
snmp
analyze
script
logout
  display
  ipFragm
  ipxSnap
  address
  agingTim
  stpState
  stpPrior
  stpMaxA
  stpHeld
  stpForw
  stpGrou
  port
  packetFilter
  list
  display
  create
  delete
  edit
  load
  assign
  unassign
  addressGroup
  portGroup
```

- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter the following command:

```
bridge packetFilter portGroup
```

- 2 To list the currently defined groups, enter this command:

```
list
```

The listing of address or port groups is displayed, as shown in the next example.

Address group example

In this example, three address groups are defined in the system. The first address group has an id of 1 and the name *Accounting*. This group uses an address group mask of 1 (the bit set in the mask) .

```
Address Groups
Address Group 1 - Accounting
    Address group mask - bit 1
Address Group 2 - Development
    Address group mask - bit 6
Address Group 3 - Sales
    Address group mask - bit 3
```

Port group example

In this example of listing port groups, two port groups are defined in the system. The first port group has an id of 1 and the name *Sales*. This group uses a port group mask of 7 (the bit set in the mask).

```
Port Groups
Port Group 1 - Sales
    Port group mask - bit 7
Port Group 2 - Manufacturing
    Port group mask - bit 16
```

Displaying Groups

The display of an address or port group shows the group id, the name of the group, and all the addresses or ports included in that group.

To display address or port groups:

- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter the following command:

```
bridge packetFilter portGroup
```

- 2 Enter this command:


```
display
```

- 3 Enter the id number of the address or port group you want to display.

The address or port group you selected is displayed.

Address group example

In this example, address group 2 is displayed. The address group id and the name (if any) are displayed, followed by Ethernet addresses that are



members of the group. The name of the address group in this example is *Development*, and the group has five members.

```
Select address group to be displayed [1-n]: 2
```

```
Address Group 2 - Development
05-39-24-56-ab-ee      08-29-34-fd-32-14      08-29-34-dd-ee-01
09-34-56-32-12-e3     00-14-32-54-fd-4e
```

Port group example

In this example, port group 2 is displayed. The port group id and the name (if any) are displayed, followed by the ports that are members of the group. The name of the port group in this example is *Manufacturing* and the group has three members.

```
Select port group to be displayed [1-n]: 2
```

```
Port Group 2 - Manufacturing
Ethernet port 1  Ethernet port 5      FDDI port 1
```

Creating New Groups

When you create a new address or port group, an unused address or port group must be available. A port group is limited to the number of ports on the system.

Top-Level Menu

```
system
ethernet
fdi
bridge
ip
snmp
analyze
script
logout
display
ipFragm
ipxSnap
address
agingTim
stpState
stpPrior
stpMax
stpHello
stpForw
stpGroup
port
packetFilter
list
display
create
delete
edit
load
assign
unassign
addressGroup
portGroup
```

- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter the following command:

```
bridge packetFilter portGroup
```

- 2 Enter this command:

```
create
```

- 3 For address groups, enter the address group mask.

For port groups, enter the port group mask.

- 4 Enter the address or port group name.

- 5 Enter the addresses or ports to add to the new group. Type **q** after entering all the addresses or ports.

Enter the addresses in MAC format as:

```
xx-xx-xx-xx-xx-xx
```

Enter the ports in this syntax:

```
< Ethernet | E | FDDI | F > [port] < port number >
```

As you enter each address or port, the system attempts to add it to the group. If the address or port you enter is already a member of the group, the system displays a message, as shown next, and the address or port is ignored.

```
Warning: Selected address was already a member of the
address group.
```

For an address group, if the system fails to accept the additional address, the address is not added to the group and an error message is displayed as follows:

```
Error: No room in group for an additional address.
```

When this message occurs, the specified address is ignored and creation of the address group stops. All addresses entered up to the last address are added to the group and the group is loaded on the system.

If you enter an invalid port name, the port is not added to the group, and you receive one of the following error messages:

```
Error: No port type specified for the port.
```

```
Error: No port number specified for the port.
```

```
The correct format is < Ethernet | E | FDDI | F > [port] <
port number >
```

```
Specified port number is invalid.
```

```
Valid FDDI port for this group is 1 or 2.
```

*Address group
example*

In this example, a new address group is created and loaded on the system. The address group mask for the group is 5 and the name of the group is *Marketing*. Two Ethernet addresses are entered and assigned to the group.

```
Select a bit in the address group mask [3-8, 14-32]: 5
```

```
Enter the address group name: Marketing
```

```
Enter the addresses for the group - type q to return to the menu:
```

```
Address: 08-32-45-fe-76-d3
```

```
Address: 08-32-45-e3-32-21
```

```
Address: q
```

```
Address Group 4 - Marketing - has been loaded
```

Port group example In this example, a new port group is created and loaded on the system. The bit in the port group mask for the group is 12 and the name of the group is *Education*. One port is entered and assigned to the group.

```
Select a bit in the port group mask [3-8, 14-32]: 12
Enter the port group name: Education
Enter the ports for the group - type q to return to the menu:
Port: Ethernet 2
Port: q
Port Group 6 - Education - has been loaded
```

Deleting Groups

When you delete address or port groups from the system, those groups are no longer available for use in packet filters.



If you want to use a group later but want to delete it now, first save it to an ASCII file.

To delete an address or port group:

Top-Level Menu

```
system
ethernet
fddi
bridge
ip
snmp
analyze
script
logout
  display
  ipFragm
  ipxSnap
  address
  agingTim
  stpState
  stpPrior
  stpMaxA
  stpHeld
  stpForw
  stpGrou
  port
  packetFilter
    list
    display
    create
    delete
    edit
    load
    assign
    unassign
    addressGroup
    portGroup
```

- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter the following command:

```
bridge packetFilter portGroup
```

- 2 Enter this command:

```
delete
```

You are prompted for the ID of the address group or port group that you want to delete.

- 3 Enter the ID number of the group you want to delete.

Adding Addresses and Ports to Groups

When adding addresses or ports to an existing group, you can either enter the addresses or ports at the prompts or import them from a file. At least one address group or port group must exist before you can add addresses or ports. (See “Creating New Groups” on page 13-4.) An address may be in multiple address groups.

Address group size


An address group for the Switch 2200 system supports a maximum of 8192 addresses in both 802.1d Bridging mode and Express switching mode. When you load an address group, the addresses that are not currently in the table are added. Therefore, the actual number of entries that you can add to an address group is limited by the address table size.

Port group size

The maximum number of ports a port group can contain is 17, which is the maximum number of ports on a Switch 2200.

For clarity, only one menu (the address group menu) is displayed here.

To add addresses or ports to an existing group:



- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter the following command:

```
bridge packetFilter portGroup
```

- 2 To add an *address* to a group, enter:

```
addAddress
```

OR, to add a *port* to a group, enter:

```
addPort
```

- 3 Enter the number of the group to modify.
- 4 Enter the addresses or ports to add to the group. Enter **q** after entering all the addresses or ports.

Enter the addresses in MAC format as:

```
xx-xx-xx-xx-xx-xx
```


Enter the ports in this syntax:

```
< Ethernet | E | FDDI | F > [port] < port number >
```

As you enter each address or port, the system attempts to add it to the group.

If the address or port you enter is already a member of the group, a message is displayed, as shown next, and the address or port is ignored.

```
Warning: Selected address was already a member of the
address group
```

OR

```
Error: Port grp - no error for the current software
```

For address groups, if the system fails to accept the additional address, the address is not added to the group and an error message is displayed as follows:

```
Error: No room in group for additional address.
```

The point at which the system runs out of room for additional addresses depends on:

- The number of addresses currently in the address table.
- The number of unique addresses configured across all address groups on the system. (Each statically configured address and each unique address assigned to one or more address groups consumes one address storage location.)

For port groups, entering an invalid port specification results in error messages, similar to those described on page 15-5.

*Address group
example*

In the example, two additional addresses are added to the *Development* address group.

```
Select address group to be modified [1-4]: 2
```

```
Adding addresses to group 2 - Development
```

```
Enter the addresses to be added - type q to return to the menu:
```

```
Address: 08-21-42-62-98-ab
```

```
Address: 08-37-21-65-78-c4
```

```
Address: q
```

Port group example

This example shows a port successfully added to the *Manufacturing* port group.

Select port group to be modified [1-4]: **2**

Adding ports to group 2 - Manufacturing

Enter the ports to be added - type **q** to return to the menu:

Port: **Ethernet 3**

Port: **q**

Removing Addresses or Ports from a Group

When removing addresses or ports from a group, you can enter the addresses or groups at the prompts. At least one group must exist to remove an address or port.

To remove addresses or ports from an address group:

- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter the following command:

```
bridge packetFilter portGroup
```

- 2 To remove an *address* from a group, enter:

```
removeAddress
```

OR, to remove a *port* from a group, enter:

```
removePort
```

- 3 Enter the number of the group to modify.
- 4 Enter the addresses or ports to remove from the new group. Type **q** after entering all the addresses or ports.


Enter the addresses in MAC format as:

```
xx-xx-xx-xx-xx-xx
```

Enter the ports in the syntax:

```
< Ethernet | E | FDDI | F > [port] < port number >
```

Top-Level Menu



As you enter addresses and ports, the system attempts to remove them from the group.

If the address or port is not found in the group, a warning message is displayed, as shown here:

```
Warning: Specified address was not a member of the
address group.
```

OR

```
Warning: Specified port was not a member of the port
group.
```

When this message occurs, the specified address or port is ignored and you are prompted for the next one to be removed.

*Address group
example*

In this example, two Ethernet addresses are removed from the *Marketing* address group.

```
Select address group to be modified [1-4]: 4
Removing addresses from group 4 - Marketing
Enter the addresses to be removed - type q to return to the menu:
Address: 08-37-21-65-78-c4
Address: 08-42-21-84-78-f1
Address: q
```

Port group example

In this example, an Ethernet and an FDDI port are removed from the *Education* port group.

```
Select port group to be modified [1-4]: 4
Removing ports from group 4 - Education
Enter the ports to be removed - type q to return to the menu:
Port: FDDI 1
Port: Ethernet 4
Port: q
```


Loading Groups

There is no explicit menu item to load address and port groups that are defined in a file on a remote host. However, you can “load” groups by creating a script on a remote host (which includes your address or port group) and then running that script.

The following example shows a script that builds an address group:

```
bridge packetFilter addressGroup create
08-37-21-65-78-c4
08-32-18-55-40-a0
08-22-12-65-78-05
08-18-23-00-82-00
08-52-12-65-5f-22
08-25-43-41-6e-09
08-00-65-23-00-ee
08-5a-42-77-8a-01
08-22-13-66-00-2a
08-8e-54-11-78-3b
08-77-12-65-78-8c
q
```


When you run the script, the address group is automatically created and stored on the system. For more information on running scripts, see *Chapter 2: How to Use the Administration Console*, on page 2-13.




APPENDICES

Appendix A Packet Filter Opcodes, Examples, and Sytax Errors

Appendix B Technical Support





PACKET FILTER OPCODES, EXAMPLES, AND SYNTAX ERRORS

This appendix:

- Describes the specific opcodes you can use when creating a packet filter
- Provides numerous examples of commonly used packet filters
- Describes the possible syntax errors you might receive when loading a packet filter



For information on creating and using packet filters, see Chapter 12.

Opcodes

Opcodes are instructions used in packet filter definitions. The available opcodes are described in this section:

name "<name>"

Description:

Assigns a user-defined <name> to the packet filter. The name may be any sequence of ASCII characters other than quotation marks. The name is limited to 32 characters. Only a single name statement can be included in a packet filter program.

Storage Needed:

2 + n bytes of packet filter storage where n is the length of the <name>

pushField.size <offset>

Description:

Pushes a field from the target packet onto the stack. Packet data starting at <offset> is copied onto the stack. The most significant byte of the field is the byte at the specified offset. The number of bytes pushed is determined by the size field of the instruction. The pushField instruction provides direct access to any 1, 2, 4, or 6 byte field contained within the first 65535 bytes of the target packet.

Certain implementations of the packet filter language further limit the maximum offset, based on the packet lengths supported by the underlying network. Ethernet-based packet filters are limited to accessing fields in the first 1518 bytes of the target packet.

Specify the offset as either an octal, decimal, or hexadecimal number.

- Precede an octal number by a "0".
- Precede a hexadecimal number by either "0x" or "0X".
- Use either upper or lower case letters for the hexadecimal digits "a" through "f".

Storage Needed:

3 bytes

pushLiteral.size <value>

Description:

Pushes a literal constant <value> onto the stack. The most significant byte of the <value> is the first byte of the literal. Bytes are copied directly from the instruction stream onto the stack. The number of bytes pushed is determined by the size field of the instruction.

Specify the value as either an octal, decimal, or hexadecimal number.

- Precede an octal number by a "0".
- Precede a hexadecimal number by either "0x" or "0X".
- Use either upper or lower case letters for the hexadecimal digits "a" through "f".

Storage Needed:

1 (.b), 2 (.w), 4 (.l), or 6 (.a) bytes—depending on the size of the operand

pushTop

Description:

Pushes the current top of the stack onto the stack (that is, it reads the top of the stack and pushes the value onto the stack). The size of the push is determined by the size of the contents of the stack.

Storage Needed:

1 byte

pushSAGM

Description:

Pushes the source address group mask (SAGM) onto the top of the stack. The SAGM is a bitmap representing the groups to which the source address of a packet belongs. This instruction pushes 4 bytes onto the stack.

Each address group is represented by a single bit in the SAGM.

Multicast addresses (including broadcast addresses) are in all groups.

Storage Needed:

1 byte

pushDAGM

Description:

Pushes the destination address group mask (DAGM) onto the top of the stack. The DAGM is a bitmap representing the groups to which the destination address of a packet belongs. This instruction pushes 4 bytes onto the stack.

Each address group is represented by a single bit in the DAGM.

Multicast addresses (including broadcast addresses) are in all groups.

Storage Needed:

1 byte

pushSPGM

Description:

Pushes the source port group mask (SPGM) onto the top of the stack. The SPGM is a bitmap representing the groups to which the source port of a packet belongs. This instruction pushes 4 bytes on to the stack.

Each port group mask is represented by a single bit in the SPGM bitmap. Port group masks are assigned to the bitmap in sequence, starting with port group mask 1 as the least significant bit through port group mask 32 as the most significant bit.

Storage Needed:

1 byte

pushDPGM

Description:

Pushes the destination port group mask (DPGM) onto the top of the stack. The DPGM is a bitmap representing the groups to which the destination port of a packet belongs. This instruction pushes 4 bytes on to the stack.

Each port group mask is represented by a single bit in the DPGM bitmap. Port group masks are assigned to the bitmap in sequence, starting with port group mask 1 as the least significant bit through port group mask 32 as the most significant bit.

Storage Needed:

1 byte

eq (equal)

Description:

Pops two values from the stack and compares them. If they are equal, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

ne (not equal)

Description:

Pops two values from the stack and compares them. If they are not equal, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

lt (less than)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is less than the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

le (less than or equal to)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is less than or equal to the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

gt (greater than)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is greater than the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

ge (greater than or equal to)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is greater than or equal to the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

and (bit-wise AND)

Description:

Pops two values from the stack and pushes the bit-wise *AND* of these values back onto the stack. The size of the operands and the result are determined by the contents of the stack.

Storage Needed:

1 byte

or (bit-wise OR)

Description:

Pops two values from the stack and pushes the bit-wise *OR* of these values back onto the stack. The size of the operands and the result are determined by the contents of the stack.

Storage Needed:

1 byte

xor (bit-wise exclusive-OR)

Description:

Pops two values from the stack and pushes the bit-wise *exclusive-OR* of these values back onto the stack. The size of the operands and the result are determined by the contents of the stack.

Storage Needed:

1 byte

not

Description:

A byte is popped from the stack; if it is non-zero, a zero byte is pushed back onto the stack. Otherwise, a non-zero byte is pushed back onto the stack.

Storage Needed:

1 byte

accept

Description:

Conditionally accepts the packet being examined. A byte is popped from the stack. If it is non-zero, the packet is accepted, and evaluation of the filter ends immediately; otherwise, filter evaluation continues with the next instruction.

Storage Needed:

1 byte

reject

Description:

Conditionally rejects the packet being examined. A byte is popped from the stack. If it is non-zero, the packet is rejected and evaluation of the filter ends immediately; otherwise, filter evaluation continues with the next instruction.

Storage Needed:

1 byte

shifl (shift left)

Description:

Pops two values from the stack and shifts the first operand left by the number of bits specified by the second operand. Bits shifted out of the left side of the operand are discarded and zeros are shifted in from the right. The resulting value is pushed back onto the stack. The size of the first operand and the size of the result are determined by the contents of the top of the stack. The second operand is always 1 byte and only the low 5 bits of the byte are used as the shift count.

Storage Needed:

1 byte

shiftr (shift right)

Description:

Pops two values from the stack and shifts the first operand right by the number of bits specified by the second operand. Bits shifted out of the right side of the operand are discarded and zeros are shifted in from the left. The resulting value is pushed back onto the stack. The size of the first operand and the size of the result are determined by the contents of the top of the stack. The second operand is always 1 byte and only the low 5 bits of the byte are used as the shift count.

Storage Needed:

1 byte

Packet Filter Examples

The following examples of using the packet filter language start with basic packet filter concepts.

Destination Address Filter

This filter operates on the destination address field of a frame. It allows packets to be forwarded that are destined for stations with an Organizationally Unique Identifier (OUI) of 08-00-02. To customize this filter to another OUI value, change the literal value loaded in the last **pushLiteral.l** instruction. Note that the OUI must be padded with an additional 00 to fill out the literal to 4 bytes.

```
name                "Forward to 08-00-02"
pushField.l         0          # Get first 4 bytes of
                        # destination address
pushLiteral.l       0xffffffff # Set up mask to isolate first
                        # 3 bytes
and
pushLiteral.l       0x08000200 # Top of stack now has OUI
                        # Load OUI value
eq                  # Check for match
```

Source Address Filter

This filter operates on the source address field of a frame. It allows packets to be forwarded that are from stations with an OUI of 08-00-02. To customize this filter to another OUI value, change the literal value loaded in the last **pushLiteral.l** instruction. Note that the OUI must be padded with an additional 00 to fill out the literal to 4 bytes.

```
name                "Forward from 08-00-02"
pushField.l         6          # Get first 4 bytes of source
                        # address
pushLiteral.l       0xffffffff # Set up mask to isolate first
                        # 3 bytes
and
pushLiteral.l       0x08000200 # Top of stack now has OUI
                        # Load OUI value
eq                  # Check for match
```

Length Filter

This filter operates on the length field of a frame. It allows packets to be forwarded that are less than 400 bytes in length. To customize this filter to another length value, change the literal value loaded in the **pushLiteral.w** instruction.

```
name                "Forward < 400"
pushField.w         12         # Get length field
pushLiteral.w       400        # Load length limit
lt                  # Check for frame length < limit
```

Type Filter This filter operates on the type field of a frame. It allows packets to be forwarded that are IP frames. To customize this filter to another type value, change the literal value loaded in the **pushLiteral.w** instruction.

```

name                "Forward IP frames"
pushField.w         12                # Get type field
pushLiteral.w       0x0800           # Load IP type value
eq                  # Check for match

```

Ethernet Type IPX and Multicast Filter This filter *rejects* frames that have either a Novell IPX Ethernet type field (8134 hex) or a multicast destination address.

```

name                "Type > 900 or Multicast"
pushField.w         12                # Get type field
pushLiteral.w       0x900            # Push type value to test
# against
gt                  # Is type field > 900 (hex)?
reject              # If yes: reject frame (done)
pushLiteral.b       0x01            # Multicast bit is low-order bit
pushField.b         0                # Get 1st byte of destination
and                 # Isolate multicast bit
not                 # Top of stack 1 to accept,
# 0 to reject

```

Multiple Destination Address Filter This filter operates on the destination address field of a frame. It allows packets to be forwarded that are destined for one of four different stations. To customize this filter to other destination stations, change the literal values.

```

name                "Forward to four stations"
pushField.a         0                # Get destination address
pushTop             # Make 3 copies of address
pushTop             #
pushTop             #
pushLiteral.a       0x367002010203 # Load allowed destination
# address
eq                  # Check for match
accept             # Forward if valid address
pushLiteral.a       0x468462236526 # Load allowed destination
# address
eq                  # Check for match
accept             # Forward if valid address
pushLiteral.a       0x347872927352 # Load allowed destination
# address
eq                  # Check for match
accept             # Forward if valid address
pushLiteral.a       0x080239572897 # Load allowed destination
# address
eq                  # Check for match

```


Address Group Filter This filter accepts only frames whose source and destination address are in the same group.

```

name                "Forward Same Source and Destination"
pushSAGM            # Get source address group mask
pushDAGM            # Get destination address group
                    # mask
and                 # Compare if source and
destination         # groups are common members of
                    # an address group (result is
                    # either zero or non-zero)
                    # address group masks
pushLiteral.1      0 # Put a zero on the stack
ne                  # If not equal, returns a "one"
                    # to stack, resulting in packet
                    # forwarded

```

Port Group Filter This filter discards all frames sourced from a port in either group three or eight.

```

name                "Discard Port Groups 3 and 8"
pushSPGM            # Get source port group mask
pushLiteral.1      0x0084 # Select bits 3 and 8
and                 # If port group bits 3 & 8 are
                    # common with SPGM, then
                    # non-zero value is
                    # pushed onto stack
pushLiteral.1      0 # Push zero
eq                  # Only if SPGM is not in port
                    # groups corresponding to bits
                    # 3 & 8, then packet is
                    # forwarded

```

Common Syntax Errors

When a packet filter definition is loaded, the definition is checked for syntax errors. The syntax errors and their causes are listed in Table A-1.


Table A-1 Possible Syntax Errors When Loading Packet Filters

Syntax Error	Description
Opcode not found	An opcode was expected on the line and was not found. The opcode must be one of those described in "Opcodes" on page A-1 and must include the size, if any. The opcode and size must be separated by a single "." with no intervening spaces. Any mix of uppercase and lowercase characters is permitted.
Unknown opcode	An opcode was expected on the line and was not found. The opcode must be one of those described in "Opcodes" on page A-1 and must include the size, if any. The opcode and size must be separated by a single "." with no intervening spaces. Any mix of upper and lower case characters is permitted.
Operands are not the same size	The opcode requires two operands of the same size. The top two operands currently on the stack are of different sizes.
Stack underflow	The opcode requires one or more operands. An insufficient number of operands are currently on the stack.
Stack overflow	The opcode pushes an operand on the stack. The stack does not have sufficient room for the operand.
No result found on top of stack	The program must end with a byte operand on the top of the stack. After the last instruction in the program is executed, the stack is either empty or contains an operand other than a byte.
Extra characters on line	The source line contains extraneous characters that are not part of the instruction and are not preceded by a comment character (#).
Expected a byte operand	The opcode requires a byte operand as one of its parameters. The operand is of a size other than a byte.
Offset not found	The opcode requires an offset to be specified. None was found on the line.
Literal not found	The opcode requires a literal value to be specified. None was found on the line.
String not found	The opcode requires a quoted string to be specified. None was found on the line.

(continued)

Table A-1 Possible Syntax Errors When Loading Packet Filters (continued)

Syntax Error	Description
Invalid characters in number	<p>The number specified as an offset or literal is improperly formatted. Possible causes are 1) lack of white space setting off the number, and 2) invalid characters in the number.</p> <p>Note: The radix of the number is determined by the first 1 or 2 characters of the number.</p> <ul style="list-style-type: none"> ■ A number with a leading "0x" or "0X" is treated as hexadecimal. ■ A number with a leading 0 is treated as octal. ■ All other numbers are treated as decimal.
Number is too large	<p>The number specified as an offset or literal is too large. An offset is limited to 1518 minus the size of the operand. For example, the offset for pushField.b can be no more than 1517 and the offset for pushField.w no more than 1516. A literal value is limited to the number of bytes in the operand size (1, 2, 4, or 6).</p>
Missing open quote on string	<p>The string specified does not have a starting quotation mark (").</p>
String is too long	<p>The string specified is too long. Strings are limited to 32 characters exclusive of the opening and closing quotation marks.</p>
Missing close quote on string	<p>The string specified does not have an ending quotation mark (").</p>
Multiple name statements in program	<p>More than one name statement was found in the program. Only a single name statement is allowed.</p>
Program too large	<p>The program exceeds the maximum size allowed. The causes of this error include a source definition exceeding 4096 bytes, a stored format exceeding 254 bytes, or a run-time format exceeding 2048 bytes. All of these boundary conditions are checked when the filter is loaded. See Table 13-2 for more information on packet filter sizes.</p>
Too many errors - compilation aborted	<p>The program contains an excessive number of errors. No further syntax errors will be reported. The program stops compiling when this condition occurs.</p>



TECHNICAL SUPPORT

3Com provides easy access to technical support information through a variety of services. This appendix describes these services.

Online Technical Services

3Com offers worldwide product support seven days a week, 24 hours a day, through the following online systems:

- 3Com Bulletin Board Service (3ComBBS)
- World Wide Web site
- 3ComForum on CompuServe®
- 3ComFactsSM automated fax service

3Com Bulletin Board Service

3ComBBS contains patches, software, and drivers for all 3Com products, as well as technical articles. This service is available via modem or ISDN seven days a week, 24 hours a day.

Access by Modem

To reach the service by modem, set your modem to 8 data bits, no parity, and 1 stop bit. Call the telephone number nearest you:

Country	Data Rate	Telephone Number
Australia	up to 14400 bps	(61) (2) 9955 2073
France	up to 14400 bps	(33) (1) 69 86 69 54
Germany	up to 9600 bps	(49) (89) 627 32 188 or (49) (89) 627 32 189
Hong Kong	up to 14400 bps	(852) 2537 5608
Italy (fee required)	up to 14400 bps	(39) (2) 273 00680
Japan	up to 14400 bps	(81) (3) 3345 7266
Singapore	up to 14400 bps	(65) 534 5693
Taiwan	up to 14400 bps	(886) (2) 377 5840
U.K.	up to 28800 bps	(44) (1442) 278278
U.S.	up to 28800 bps	(1) (408) 980 8204

Access by ISDN

ISDN users can dial in to 3ComBBS using a digital modem for fast access up to 56 Kbps. To access 3ComBBS using ISDN, dial the following number:

(408) 654 2703

World Wide Web Site

Access the latest networking information on 3Com's World Wide Web site by entering our URL into your Internet browser:

<http://www.3Com.com/>

This service features news and information about 3Com products, customer service and support, 3Com's latest news releases, selected articles from 3TECH™ journal (3Com's award-winning technical journal), and more.

3ComForum on CompuServe®

3ComForum is a CompuServe service containing patches, software, drivers, and technical articles about all 3Com products, as well as a messaging section for peer support. To use 3ComForum, you need a CompuServe account.

To use 3ComForum:

- 1 Log on to CompuServe.
- 2 Enter **go threecom**
- 3 Press [Return] to see the 3ComForum main menu.

3ComFactsSM Automated Fax Service

3Com Corporation's interactive fax service, 3ComFacts, provides data sheets, technical articles, diagrams, and troubleshooting instructions on 3Com products 24 hours a day, seven days a week.

Call 3ComFacts using your touch-tone telephone and international access numbers:

Country	Telephone Number
Hong Kong	(852) 2537 5610
U.K.	(44) (1442) 278279
U.S.	(1) (408) 727 7021

Local access numbers are available within the following countries:

Country	Telephone Number	Country	Telephone Number
Australia	800 123853	Netherlands	06 0228049
Belgium	0800 71279	Norway	800 11062
Denmark	800 17319	Portugal	0505 442607
Finland	98 001 4444	Russia (Moscow only)	956 0815
France	05 90 81 58	Spain	900 964445
Germany	0130 8180 63	Sweden	020 792954
Italy	1678 99085	U.K.	0800 626403

Support from Your Network Supplier

If additional assistance is required, contact your network supplier. Many suppliers are authorized 3Com service partners who are qualified to provide a variety of services, including network planning, installation, hardware maintenance, application training, and support services.

When you contact your network supplier for assistance, have the following information ready:

- Diagnostic error messages
- A list of system hardware and software, including revision levels
- Details about recent configuration changes, if applicable

If you are unable to contact your network supplier, see the following section on how to contact 3Com.

Support from 3Com

If you are unable to receive support from your network supplier, technical support contracts are available from 3Com.

In the U.S. and Canada, call **(800) 876-3266** for customer service.

If you are outside the U.S. and Canada, contact your local 3Com sales office to find your authorized service provider:

Country	Telephone Number	Country	Telephone Number
Australia (Sydney)	(61) (2) 9937 5000	Japan	(81) (3) 3345 7251
(Melbourne)	(61) (3) 9866 8022	Mexico	(525) 531 0591
Belgium*	0800 71429	Netherlands*	06 0227788
Brazil	(55) (11) 546 0869	Norway*	800 13376
Canada	(905) 882 9964	Singapore	(65) 538 9368
Denmark*	800 17309	South Africa	(27) (11) 803 7404
Finland*	0800 113153	Spain*	900 983125
France*	05 917959	Sweden*	120 795482
Germany*	0130 821502	Taiwan	(886) (2) 577 4352
Hong Kong	(852) 2501 1111	United Arab Emirates	(971) (4) 349049
Ireland*	1 800 553117	U.K.*	0800 966197
Italy*	1678 79489	U.S.	(1) (408) 492 1790

* These numbers are toll-free.

Returning Products for Repair

A product sent directly to 3Com for repair must first be assigned a Return Materials Authorization (RMA) number. A product sent to 3Com without an RMA number will be returned to the sender unopened, at the sender's expense.

To obtain an RMA number, call or fax:

Country	Telephone Number	Fax Number
U.S. and Canada	(800) 876 3266, option 2	(408) 764 7120
Europe	31 30 60 29900, option 5	(44) (1442) 275822
Outside Europe, U.S., and Canada	(1) (408) 492 1790	(1) (408) 764 7290

INDEX

Numerics

3Com Bulletin Board Service (3ComBBS) B-1
3Com sales offices B-4
3ComFacts B-3
3ComForum B-2

A

abort

- at prompts 2-9
- enabling CTL+C 2-11

accept opcode 12-8, A-7

access levels 2-1

address

- adding static 11-12
- aging time 10-6
- filters A-9
- flushing 11-13
- for SNMP trap reporting 3-17
- freezing 11-13
- in routing table 3-7
- IP 3-3
- IP to MAC, translating 3-11
- maximum number in group 13-7
- removing static 11-12

address group

- adding addresses 13-7 to 13-9
- as filtering criteria 13-1
- copying 13-7
- creating 13-4
- deleting 13-6
- displaying contents 13-3
- listing 13-2
- removing addresses 13-9
- used in packet filter 13-1

Address Resolution Protocol. *See* ARP

address threshold

- values 10-6

addressThresholdEvent 10-6

administer access example 2-2

Administration Console

- command strings 2-8
- Control keys 2-11

- entering values 2-9

- exiting 2-17

- initial user access 2-1

- interface parameters 2-10, 2-11

- locking 2-12

- menu descriptions 2-3 to 2-7

- menu hierarchy, moving up 2-9

- menu options, selecting 2-8

- password access 2-1, 4-2

- restart 2-11

- screen height, setting 2-10

- scripts 2-13

- top-level menu 2-3

aging time

- setting for bridge 10-6

- values 10-6

analyzer

- connecting 9-3

- MAC address display 9-3

- removing port 9-4

- setting up monitored port 9-5

and (bit-wise AND) opcode A-6

AppleTalk

- packet filter 12-9

ARP

- See also* ARP cache

- defined 3-11

ARP cache

- flushing 3-12

- removing entry 3-11

ASCII-based editor

- and scripts 2-13

- for packet filters 12-17

B

backup

- saving NV data 6-2

baseline

- displaying current 5-2

- enabling and disabling 5-2

- reasons for 5-1

- setting 5-2

baud rate
 console serial port 3-2
 bell, warning 4-1
 blocking state 11-5
 bridge
 See also packet filter
 address threshold, setting 10-6
 aging time, setting 10-6
 designated 11-3
 IP fragmentation, enabling 10-5
 IPX Snap Translation, enabling 10-5
 menus 2-5
 Spanning Tree
 bridge priority, setting 10-7
 enabling 10-7
 forward delay, setting 10-9
 hello time, setting 10-9
 maximum age, setting 10-8
 statistics, displaying 10-1
 bridge port
 MAC addresses
 adding 11-12
 flushing 11-13
 freezing 11-13
 listing 11-11
 removing 11-12
 multicast limit, setting 11-7
 Spanning Tree
 enabling 11-8
 path cost, setting 11-9
 port priority, setting 11-10
 states defined 11-5
 statistics, displaying 11-1
 bridging
 commands, full 10-1 to 10-10
 commands, quick 1-4
 mode defined 10-4
 broadcast address 3-4
 bulletin board service B-1

C

code. *See* scripting *and* packet filter
 commands
 and entering values 2-9
 quick 1-1
 using 2-8
 community strings
 setting 3-16
 values 3-15
 CompuServe® B-2

connectPolicy
 configuring 8-4
 Console serial port
 for management 3-1
 reasons for disconnecting 2-11
 rebooting the system 4-4
 setting baud rate 3-2
 Control keys
 enabling 2-11
 conventions
 notice icons 3
 cost
 See also metric
 of IP interface 3-4
 Spanning Tree settings 10-4, 11-3, 11-9
 CTL+C (abort) 2-11
 CTL+X (reboot) 2-11

D

datagrams, statistics 3-14
 date
 formats 4-4
 setting system 4-3
 default route
 defined 3-8
 displayed 3-8
 removing 3-10
 setting 3-10
 destination address
 for SNMP trap reporting 3-18
 destination address group mask (DAGM) 13-1
 destination IP address
 in routing table 3-7
 destination port group mask (DPGM) 13-1
 direct, route status 3-8
 documentation
 for the Switch 2200 system 4

E

editor
 for packet filters 12-17
 for scripts 2-13
 EMACS editor 2-13, 12-17
 eq opcode A-4
 Ethernet
 analyzing segments 9-1 to 9-6
 commands, quick 1-5
 fragmenting packets 10-5
 menus 2-4
 packet fields 12-6

- portState 7-8
- station MAC addresses 11-11

Ethernet address

- and restoring NV data 6-3
- for the monitored port 9-5

Ethernet port

- analyzer attached 9-3
- displaying information 7-1
- label 7-4
- labeling 7-8
- setting state (on-line or off-line) 7-8
- static MAC addresses 11-12
- statistics 7-3

F

fan, warning 4-2

fax service B-3

FDDI

- commands, quick 1-7
- fragmenting packets 10-5
- management 8-1
- menus 2-5
- packet fields 12-6
- port label 8-20
- rings 8-6
- station MAC addresses 11-11
- wrapped ring 8-6

FDDI MAC

- condition report 8-16
- defined 8-9
- FrameErrorThreshold, setting 8-16
- LLC Service, enabling 8-18
- NotCopiedThreshold, setting 8-17
- statistics, displaying 8-10

FDDI path

- defined 8-6
- maxT-Req, setting 8-9
- statistics, displaying 8-6
- tmaxLowerBound, setting 8-8
- tvxLowerBound, setting 8-7

FDDI port

- and roving analysis 9-6
- defined 8-19
- labeling 8-22
- lerAlarm, setting 8-20
- lerCutoff, setting 8-21
- statistics, displaying 8-19

FDDI station

- and SMT 8-1
- and SRFs 8-2, 8-5
- connection policies, setting 8-4

- defined 8-1
- statistics, displaying 8-2
- status reporting, enabling 8-5
- T-notify, setting 8-5

filter id 12-2

flushing

- ARP cache 3-12
- learned routes 3-10
- MAC addresses 11-13
- SNMP trap addresses 3-19

forward delay 10-9

forwarding state 11-5

FrameErrorThreshold

- defined 8-16
- setting 8-16

freezing addresses 11-13

ftp

- IP address 3-1, 3-3

G

gateway

- IP address 3-8

ge opcode A-6

group address

- Spanning Tree, setting 10-10

group. *See* address group or port group

gt opcode A-6

H

hello time 10-9

Help

- Administration Console 2-16
- topical 2-16

I

ICMP

- and ping 3-13
- echo (request and reply) 3-13

in-band management 3-2

instructions

- opcodes 12-5, A-1
- operands 12-5, 12-7
- operators 12-7

interface

- Administration Console parameters 2-10, 2-11
- defining 3-5
- displaying 3-4
- parts of 3-3
- removing definition 3-7

Internet Control Message Protocol. *See* ICMP

IP

- address translation 3-11
- ARP cache 3-11
- interface 3-3
- management access 3-1
- menus 2-6
- pinging 3-12
- RIP mode 3-12
- route table 3-8
- routes 3-7
- statistics, displaying 3-14

IP address

- and restoring NV data 6-3
- configuring 3-5
- for IP interface 3-3
- in routing table 3-7

IP fragmentation

- enabling 10-5

IP interface

- address 3-3
- broadcast address 3-4
- cost 3-4
- defining 3-3, 3-5
- displaying 3-4
- removing definition 3-7
- subnet mask 3-3

IP packets filter 12-12, 12-16

IP route

- default 3-8, 3-10
- defining static 3-9
- destination address 3-7
- gateway IP address 3-8
- metric 3-7
- removing from table 3-9, 3-10
- status 3-8
- subnet mask 3-7

IPX Snap Translation

- enabling 10-5

L

le opcode A-5

learned, route status 3-8

learning state 10-9, 11-5

LER (Link Error Rate)

- alarm value 8-20
- cutoff value 8-21

lerAlarm

- and lerCutoff value 8-21
- defined 8-20
- setting 8-21

lerCutoff

- and lerAlarm value 8-21
- defined 8-21

Link Error Rate. *See* LER

listening state 10-9, 11-5

LLC

- enabling 8-18
- service description 8-18

Logical Link Control. *See* LLC

lt opcode A-5

M

MAC (Media Access Control) address

- adding 11-12
- and ARP 3-11
- configuring 11-11
- displaying 11-11
- dynamic to static 11-13
- flushing 11-13
- removing static 11-12
- roving analysis configuration 9-2

MAC (Media Access Control). *See* FDDI MAC management

- and naming the system 4-3
- configuring system access 3-1 to 3-13
- FDDI 8-1
- in-band 3-2
- IP interface 3-1, 3-3
- out-of-band 3-2
- port labels 7-8, 8-22
- setup, quick commands 1-3
- SNMP community strings 3-15
- system name 4-3
- Transcend® Enterprise Manager 1-1

maximum age 10-8

maxT-Req

- defined 8-9
- setting 8-9

menu

- analyzer (roving analysis) 2-7
- and command strings 2-8
- bridge 2-5
- ethernet 2-4
- fddi 2-5
- IP 2-6
- moving up hierarchy 2-9
- selecting options 2-8
- SNMP 2-7
- system 2-4

metric

- in routing table 3-7

- multicast frames
 - and packet filters 12-1
- multicast limit
 - configuring 11-7
 - defined 11-7

N

- name opcode A-1
- naming the Switch 2200 4-3
- ne opcode A-5
- neighbor notification
 - and LLC Service 8-18
- network monitoring. *See* roving analysis *and* analyzer
- network supplier support B-3
- network troubleshooting 9-1
- not opcode A-7
- NotCopiedThreshold
 - defined 8-17
 - setting 8-17
- Novell
 - in packet filter A-10
- NV data
 - and packet filters 12-3
 - backup 6-1
 - contents saved 6-1
 - examining a saved file 6-5
 - file information 6-1
 - resetting 6-6
 - saving 6-2
 - transferring 6-1

O

- off-line
 - port state 7-8
- on-line
 - port state 7-8
- on-line Help 2-16
- on-line technical services B-1
- opcode
 - and packet filter language 12-4
 - and writing packet filters 12-10
 - descriptions A-1 to A-8
- operand 12-5
 - and opcodes 12-7
 - sizes supported 12-5
- operator
 - and opcodes 12-7
 - purpose 12-7
- or opcode A-7

- OUI
 - in packet filter A-11
- out-of-band management 3-2

P

- packet
 - Ethernet type 12-6
 - FDDI type 12-6
 - fields for operands 12-7
- packet filter
 - See also* address group *and* port group
 - address group example 13-1
 - assigning to ports 12-22
 - basic elements 12-6
 - concepts 12-4 to 12-11
 - correcting errors 12-21
 - creating 12-3 to 12-17
 - definitions 12-3
 - deleting 12-20
 - displaying contents 12-3
 - editing 12-20
- editor
 - commands 12-19
 - description 12-17
 - using 12-18
- examples 12-11 to 12-17, A-9 to A-12
- external editor 12-20
- filter id 12-2
- filtering criteria, groups 13-1
- instructions 12-5
- language description 12-3, 12-4
- listing 12-2
- loading 12-22
- opcodes A-1
- operands 12-5
- port group example 13-2
- pre-processed storage 12-9
- procedure for writing 12-10
- processing paths 12-1, 12-22
- pseudocode 12-12
- run-time storage 12-10
- sequential tests 12-8
- stack 12-5
- storage space 12-9
- syntax errors A-13, A-14
- unassigning from ports 12-24

- password
- configuring 4-2
- initial system access 2-1
- levels of user access 2-1

path cost
 defined 11-9
 setting 11-9

path. *See* FDDI path *and* backplane paths

PHY
 and FDDI ports 8-19

ping
 IP station 3-12

PMD
 and FDDI ports 8-19

port
 See also FDDI port
 bridging priority 11-10
 for analyzer 9-3
 including in IP interface 3-4
 label 8-20
 maximum number in group 13-7
 path cost 11-9
 speed, setting 3-2
 state, setting 7-8
 types 8-19

port group
 adding ports 13-7 to 13-9
 as filtering criteria 13-1
 copying 13-7
 creating 13-4
 deleting 13-6
 displaying contents 13-3
 listing 13-2
 removing ports 13-9
 used in packet filter 13-2

power supply warning 4-2

probe. *See* roving analysis *and* analyzer

pushDAGM opcode 13-1, A-3

pushDPGM opcode 13-1, A-4

pushField.opcode A-2

pushLiteral.opcode A-2

pushSAGM opcode 13-1, A-3

pushSPGM opcode 13-1, A-4

pushTop opcode A-3

R

read access example 2-3

reboot
 enabling CTL+X 2-11
 resetting the system 4-4

reboot system 2-11

receive all
 packet processing path 12-1

receive multicast
 packet processing path 12-1

reject opcode 12-8, A-8

remote sessions
 enabling timeout 2-12
 setting timeout interval 2-13

restart, Administration Console 2-11

returning products for repair B-4

RIP
 and broadcast address 3-4
 default mode 3-12
 displaying state 3-4
 Off mode 3-12
 Passive mode 3-12
 setting mode 3-12

rlogin
 and exiting the Console 2-17
 and rebooting the system 4-4

route
 See also routing table
 default 3-8
 defining static 3-9
 destination IP address 3-7
 flushing learned routes 3-10
 gateway IP address 3-8
 metric 3-7
 removing default 3-10
 removing from table 3-9, 3-10
 status 3-8
 subnet mask 3-7

Routing Information Protocol. *See* RIP

routing table
 contents 3-7
 default route, setting 3-10
 display routes 3-8
 flushing learned routes 3-10
 removing default route 3-10
 removing route 3-10
 removing routes 3-9

roving analysis
 adding analyzer port 9-3
 and Spanning Tree 9-4
 configuration rules 9-2
 configuration, displaying 9-2, 9-3
 configuring 9-2
 defined 9-1
 menu 2-7
 process overview 9-1
 removing analyzer port 9-4
 starting port monitoring 9-5
 stopping port monitoring 9-6

S

- SAGM (source address group mask) 13-1
- screen height
 - adjusting 2-10
- scripts for the Administration Console
 - examples 2-15
 - running 2-13
- serial port (console)
 - for management 3-1
 - rebooting the system 4-4
 - setting baud rate 3-2
- Service Access Points (SAPs)
 - and packet filters 12-4
- shiftr opcode A-8
- shiftr opcode A-8
- SMT (Station Management)
 - and FDDI stations 8-1
 - lerAlarm value 8-21
 - lerCutoff value 8-21
- SMT event
 - enabling proxying 3-20
 - proxying defined 3-19
- Sniffing. *See* roving analysis *and* analyzer
- SNMP
 - See also* trap *and* community strings
 - community strings
 - setting 3-16
 - values 3-15
 - displaying configurations 3-15
 - menus 2-7
 - proxying remote SMT events 3-20
 - trap reporting
 - and SMT event proxying 3-19
 - configuring destinations 3-17
 - descriptions of traps 3-17
 - displaying configuration 3-16
 - flushing addresses 3-19
- SNMP agent
 - accessing through IP 3-1
 - defined 3-15
- SNMP trap
 - Address Threshold 3-17
 - addressThresholdEvent 10-6
 - Authentication Failure 3-17
 - Coldstart 3-17
 - MAC Duplicate Address Condition 3-17
 - MAC Frame Error Condition 3-17
 - MAC Neighbor Change 3-17
 - MAC Not Copied Condition 3-17
 - MAC Path Change 3-17
 - New Root 3-17
 - Port EB Error Condition 3-17
 - Port LER Condition 3-17
 - Port Path Change 3-17
 - Port Undesired Connection 3-17
 - SMT Hold Condition 3-17
 - SMT Peer Wrap Condition 3-17
 - System Overtemperature 3-17
 - Topology Change 3-17
- socket values filter 12-12, 12-15
- software
 - backup NV data 6-1, 6-2
 - build date and time 4-1
 - from factory 1-1
 - version number 4-1
- source address group mask (SAGM) 13-1
- source port group mask (SPGM) 13-1
- Spanning Tree Protocol. *See* STP
- SPGM (source port group mask) 13-1
- SRF (Status Report Frames)
 - and FDDI stations 8-2, 8-5
 - and lerAlarm 8-20
- stack 12-5
- static route status 3-8
- station. *See* FDDI station
- Station Management. *See* SMT
- statistics
 - baselining 5-1
 - Ethernet ports 7-3
 - FDDI MAC 8-10, 8-11
 - FDDI path 8-6
 - FDDI station 8-2
 - IP 3-14
- Status Report Frames. *See* SRF
- status reporting
 - configuring 8-5
 - defined 8-5
- STP (Spanning Tree Protocol)
 - bridge priority, setting 10-7
 - designated bridge 11-3
 - designated cost 11-3
 - designated port 11-3
 - designated root 11-3
 - enabling on bridge 10-7
 - enabling on bridge port 11-8
 - forward delay, setting 10-9
 - group address, setting 10-10
 - hello time, setting 10-9
 - maximum age, setting 10-8
 - port priority 11-10
 - states 11-5
- subnet mask
 - for IP address 3-3
 - in routing table 3-7

Switch 2200

- administration overview 1-1
- and network monitoring 9-1
- bell warning 4-1
- documentation 4
- fan warning 4-2
- naming 4-3
- NV data restoration 6-3
- ports and IP interfaces 3-6
- power supply warning 4-2
- quick commands 1-1
- rebooting 4-4
- resetting to system defaults 6-6
- system backup 6-2
- system configuration, displaying 4-1
- system date and time 4-3
- temperature warning 4-2
- user access levels 2-1
- warning messages 4-1

system configuration

- displaying 4-1

system menus 2-4

T

T_Opr 8-9

technical support B-1

telnet

- enabling timeout 2-12
- rebooting the system 4-4
- setting timeout interval 2-13

temperature, warning 4-2

terminal emulation

- and the serial port 3-1

text editor, built-in 12-17

time

- formats 4-4
- setting system 4-3

timing out route status 3-8

tmaxLowerBound

- defined 8-8
- setting 8-8

T-notify

- configuring 8-5
- defined 8-5

token

- and FDDI MAC 8-9

transmit all

- packet processing path 12-1

transmit multicast

- packet processing path 12-1

trap reporting

- configuring destinations 3-17
- descriptions of traps 3-17
- flushing addresses 3-19

T-Req 8-9

tvxLowerBound

- defined 8-7
- setting 8-8

U

UNIX

- and terminal emulation with Switch 2200 3-1

V

vi editor 2-13, 12-17

W

warning messages for system 4-1

wrapped ring 8-6

write access example 2-2

X

XNS

- in packet filter 12-12, 12-14, A-11

xor opcode A-7