




SUPERSTACK™ II SWITCH 2200 OPERATION GUIDE



Part No. 801-00311-000
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 Unpublished rights reserved under the copyright laws of the United States.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, EtherDisk, EtherLink, EtherLink II, LANplex, LinkBuilder, NETBuilder, NETBuilder II, NetFacts, SmartAgent, TokenDisk, TokenLink, and Transcend are registered trademarks of 3Com Corporation. 3TECH, FDDLlink, NetProbe, SuperStack, and Star-Tek are trademarks of 3Com Corporation. 3ComFacts is a service mark of 3Com Corporation.

AIX and IBM 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. 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, Trish Crawford, Lynne Gelfand, Michael Jenness, Patricia Johnson, Michael Taillon, and Iain Young. Edited by Bonnie Jo Collins.

CONTENTS

ABOUT THIS GUIDE

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

PART I MANAGEMENT AND ADMINISTRATION

1 MANAGEMENT AND ADMINISTRATION OVERVIEW

- About the Switch 2200 System 1-1
- User Interfaces and the Switch 2200 System 1-1

2 USER ACCESS: WHAT YOU SEE

- About the User Interfaces to the Switch 2200 2-1
- Switch 2200 Administration Console 2-2
- External Network Management Applications 2-4

3 MANAGEMENT ACCESS: PROTOCOLS

- About Switch 2200 Protocols 3-1
- Virtual Terminal Protocols 3-3
- SNMP 3-4
 - SNMP Agent 3-4
 - SNMP MIBs 3-5
 - SNMP Traps 3-6
 - Access Control 3-6
- SMT 3-8
- SNMP and SMT Proxy Agents 3-8

4 PHYSICAL ACCESS: PORTS AND CABLING

- In-band and Out-of-band Management 4-2
- Management Access 4-2
 - Console Serial Port 4-2
 - Ethernet and FDDI Ports 4-3

PART II BRIDGING

5 TRANSPARENT BRIDGING

- About Transparent Bridging 5-1
- What Makes a Bridge 802.1d Compliant? 5-1
- How a Bridge Learns Addresses 5-2
- How a Bridge Ages Addresses 5-3
- Packet Forwarding 5-4
- Spanning Tree and the Bridged Network 5-6
 - Packet Looping in a Bridged Network 5-6
 - The Spanning Tree Algorithm 5-7
 - How the Spanning Tree Algorithm Works 5-8
 - How Spanning Tree Is Calculated for the Network 5-13
 - Spanning Tree Port States 5-16
 - Reconfiguring the Bridged Network Topology 5-18
- Bridging References 5-18

6 USER-DEFINED PACKET FILTERING

- About User-defined Packet Filtering 6-1
 - Designing a Packet Filter 6-1
 - Assigning Packet Filters to Paths 6-2
 - Packet Filter Examples 6-3
 - Example 1: Isolating IP Segments 6-3
 - Example 2: Filtering AppleTalk Phase II Packets 6-5
- Using Address Groups and Port Groups in a Packet Filter 6-7
 - What Is an Address Group? 6-8
 - What Is a Port Group? 6-9
 - Referencing Address Groups and Port Groups from a Packet Filter 6-10
 - Example: Using Address Groups in a Packet Filter 6-10
- Globally Administering Packet Filters 6-15

7	BRIDGING EXTENSIONS
	Multicast Packet Firewalls 7-1
	IP Fragmentation 7-2
	Reduced Packet Flooding 7-3
	Enhanced Network Security 7-4

PART III FDDI TECHNOLOGY

8	FDDI OVERVIEW AND IMPLEMENTATION
	About FDDI 8-1
	Ports 8-3
	MACs 8-4
	MAC Services 8-4
	MAC Operation 8-4
	Paths 8-5
	Nodes and Attachments 8-5
	Nodes 8-6
	Attachments 8-6
	Node Types 8-7
	Station Management 8-9
	SMT Operation 8-9
	The FDDI MIB 8-9
	Frame-based Protocols 8-10
	FDDI and the Switch 2200 System 8-11

9	FDDI NETWORKS
	About FDDI Networks 9-1
	FDDI Network Topologies 9-2
	Physical Topology: The Ring of Trees 9-2
	Logical Topology: The Dual Ring 9-4
	FDDI Connection Rules 9-4
	Dual Homing 9-6

PART IV APPENDIXES

A SNMP MIB SUPPORT

SNMP MIBs A-1

SNMP MIB Compilers A-3

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

 3ComFacts Automated Fax Service B-3

Support from Your Network Supplier B-3

Support from 3Com B-4

Returning Products for Repair B-4

OPERATION GLOSSARY

INDEX

ABOUT THIS GUIDE

Introduction

The *SuperStack™ II Switch 2200 Operation Guide* provides all the information you need to understand how your Switch 2200 system works in both FDDI and Ethernet networking environments.

Audience 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 used on interconnected LANs.



If the information in the release notes shipped with your product differs from the information in this guide, follow the release notes.

How to Use This Guide

The following table shows where to find specific information.

Table 1 Locating Information in This Guide

If you are looking for information on...	Turn to...
An overview of Switch 2200 system operation	Chapter 1
Access interfaces for the user/network administrator	Chapter 2
Management protocols	Chapter 3
Cabling the Switch 2200 for management access	Chapter 4
Transparent bridging issues	Chapter 5
Spanning Tree information	Chapter 5
User-defined packet filters	Chapter 6
Bridging extensions	Chapter 7
FDDI technology and its implementation in the Switch 2200	Chapter 8
FDDI networks	Chapter 9
SNMP MIB Support	Appendix A
3Com Technical Support	Appendix B
General definitions	Glossary

Conventions

Table 2 and Table 3 list 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" vs. "Type"	The word "enter" means type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says "type."
"Syntax" vs. "Command"	<p>The word "syntax" indicates that the general form of a command syntax is provided. You must evaluate the syntax and supply the appropriate port, path, value, address, or string. Example:</p> <p style="padding-left: 40px;">The following syntax specifies the time and date:</p> <p style="padding-left: 40px;">mm/dd/yy hh:mm:ss</p> <p>The word "command" indicates that all variables in the command have been supplied and you must enter the command as shown. Example:</p> <p style="padding-left: 40px;">The following command enables Spanning Tree:</p> <p style="padding-left: 40px;">bridge stpState enabled</p>
Text for <code>screen display</code>	This typeface represents displays that appear on your terminal screen. Example:
Text for commands	This typeface represents commands that you enter. Example:
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 you must press two or more keys simultaneously, the keys are linked with a plus sign (+). Example:</p> <p>Press [Ctrl]+[Alt]+[Del].</p>
<i>Italics</i>	<i>Italics</i> are used to denote <i>emphasis</i> or new terms where they are defined.

Switch 2200 Documentation

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

- *SuperStack™ II Switch 2200 Unpacking Instructions*

Describes how to unpack your Switch 2200 system. It also provides you with an inventory list of all the items shipped with your system. (Shipped with your system)
- *SuperStack™ II Switch 2200 Software Installation and Release Notes*

Provides information about the software release, including new features, installation, procedures, and bug fixes. It also describes any changes to the Switch 2200 system's documentation. (Shipped with your system)
- *SuperStack™ II Switch 2200 Getting Started*

Describes all the procedures necessary for planning your configuration and for installing, cabling, powering up, configuring management access, and troubleshooting your Switch 2200 system. (Shipped with your system/Part No. 801-00309-000)
- *SuperStack™ II Switch 2200 Operation Guide (This guide)*

Helps 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 your system/Part No. 801-00311-000)
- *SuperStack™ II Switch 2200 Administration Console User Guide*

Provides information about using the Administration Console embedded system software to configure and manage your Switch 2200 system. (Shipped with your system/Part No. 801-00310-000)
- *SuperStack™ II Switch 2200 Administration Console Command Quick Reference*

Contains Administration Console intelligent switching commands for the Switch 2200 system. (Folding card; shipped with your system/Part No. 801-00314-000)


Documentation Comments

Your suggestions are very important to us. To help make the documentation more useful to you, please send comments about this document in e-mail to 3Com at: **sdtechpubs_comments@3Mail.3Com.com**

Please include the following information when commenting:

- Document title
- Document part number (listed on the back cover and the title page)
- Page number (if appropriate)

*Example: SuperStack™ II Switch 2200 Administration Console User Guide
Part No. 801-00310-000
Page 2-5 (chapter 2, page 5)*




MANAGEMENT AND ADMINISTRATION

Chapter 1 Management and Administration Overview

Chapter 2 User Access: What You See

Chapter 3 Management Access: Protocols

Chapter 4 Physical Access: Ports and Cabling



1

MANAGEMENT AND ADMINISTRATION OVERVIEW

This chapter introduces you to how your SuperStack™ II Switch 2200 system is managed and administered.

About the Switch 2200 System

The SuperStack™ II Switch 2200 system combines high-port-density Ethernet switching and Ethernet-to-FDDI bridging in an integrated system. You can configure much of this functionality to meet your specific networking needs. You can use two different mechanisms to configure your Switch 2200 system: the Administration Console or Transcend® Enterprise Manager or another SNMP-based network management application.

The Administration Console is a character-oriented, menu-driven user interface for performing complete system-level and module administration. You operate this console from a terminal or through a virtual terminal protocol, such as telnet or rlogin.

For more complete network management, use an external SNMP-based application such as Transcend® Enterprise Manager.

User Interfaces and the Switch 2200 System

Figure 1-1 shows how the user interfaces interact with the Switch 2200 system. The illustration also indicates where to find related information in this guide.

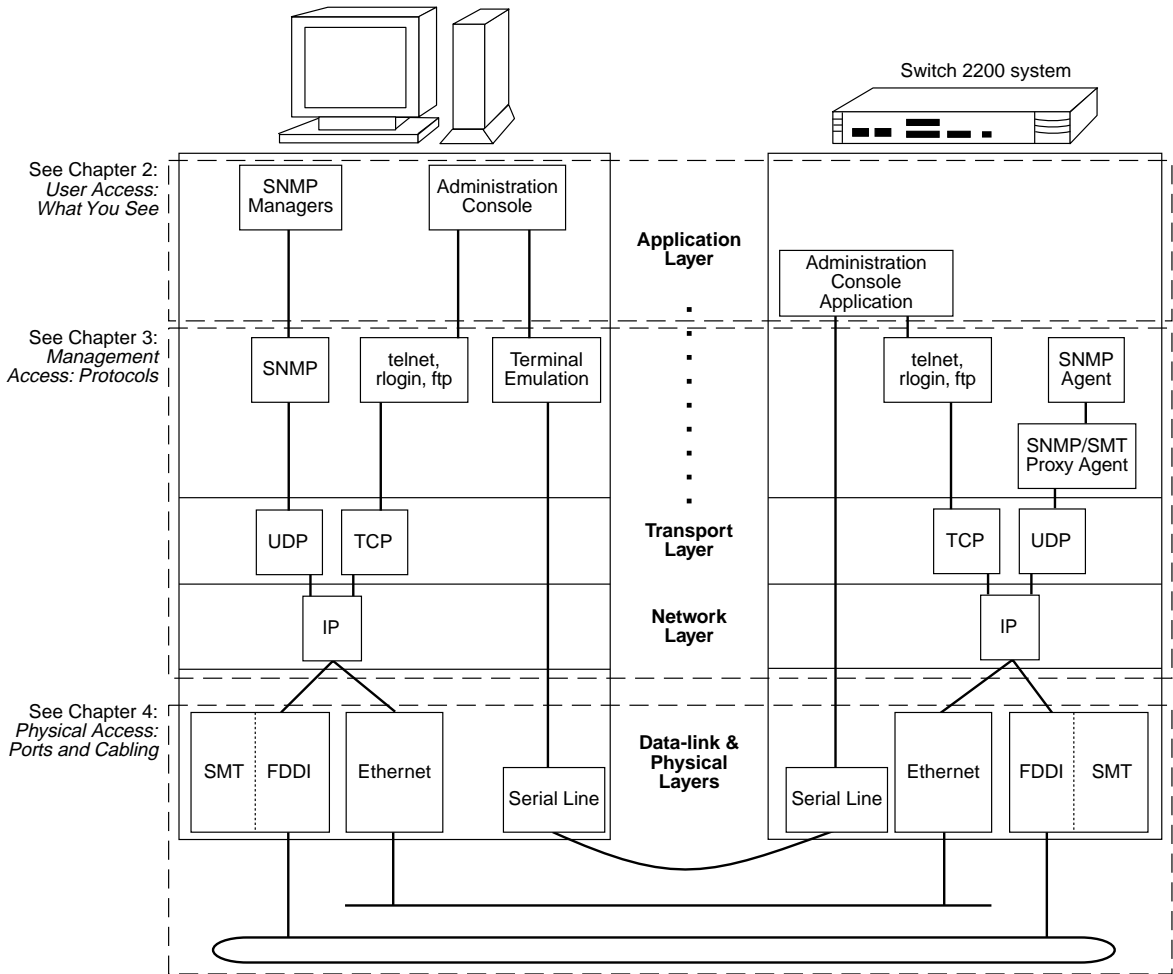


Figure 1-1 User Interfaces and Protocols for Accessing the Switch 2200 System

2

USER ACCESS: WHAT YOU SEE

This chapter describes the applications you can use to gain access to your SuperStack™ II Switch 2200 system and to perform administrative and management functions.

About the User Interfaces to the Switch 2200

The following applications provide a user interface to the Switch 2200 system:

- Administration Console
- External SNMP-based network management applications, such as Transcend® Enterprise Manager

Figure 2-1 highlights the Switch 2200 system user interfaces.

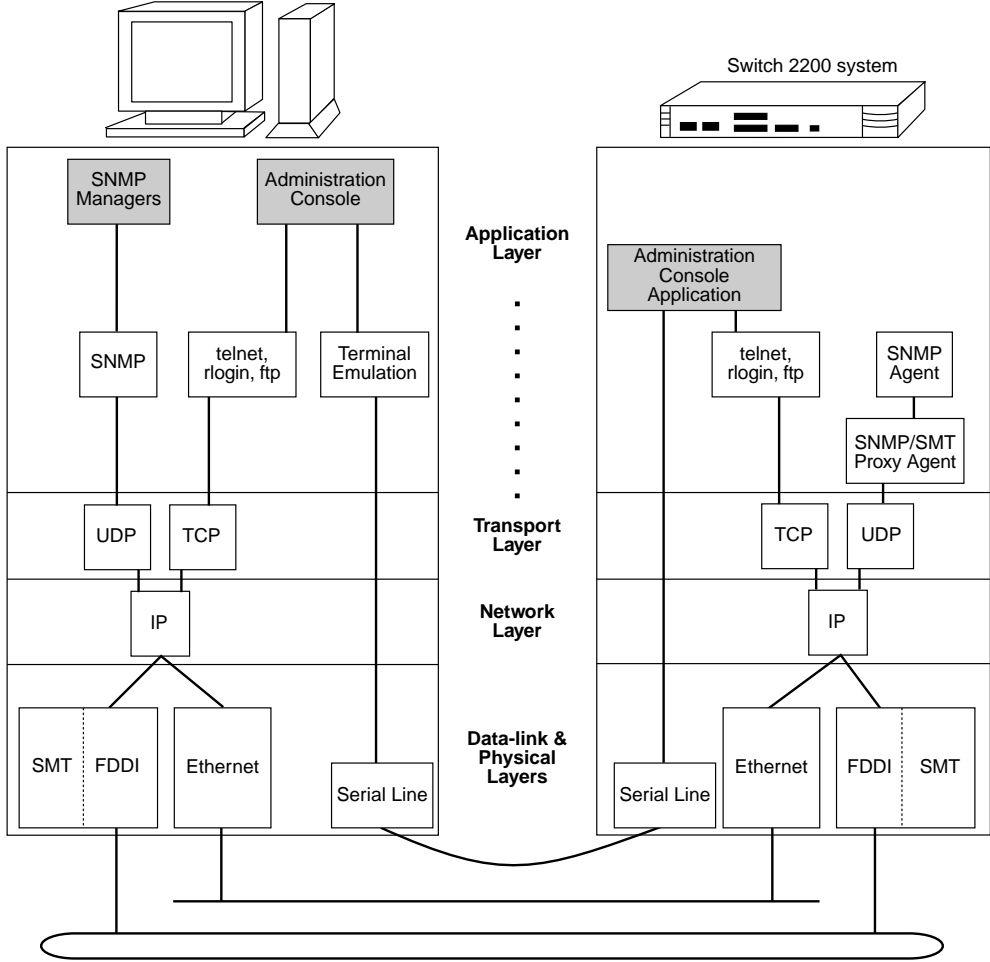


Figure 2-1 User Interfaces for the Switch 2200 System

Switch 2200 Administration Console

You can use the Administration Console to configure your Switch 2200 system to operate effectively in your networking environment. You can also use the Administration Console to display network statistics.

You can view the Administration Console from a terminal, a workstation, a Macintosh, or a PC. See Figure 2-2.

```

Menu options: -----
system                - Administer system-level functions
ethernet              - Administer Ethernet ports
fdci                  - Administer FDDI resources
bridge                - Administer bridging
ip                    - Administer IP
snmp                  - Administer the SNMP
analyzer              - Administer Roving Analysis
script                - Run a script of console commands
logout                - Logout of the Administration Console

Type ? for help.
-----
Select a menu option:bridge
Menu options: -----
display               - Display bridge information
ipFragmentation       - Enable/Disable IP Fragmentation
ipxSnmpTranslation    - Enable/Disable IP 802.3-FDDI SNAP Translation
addressThreshold      - Set the bridge address threshold
agingTime             - Set the bridge aging time
stpState              - Enable/Disable Spanning Tree on a bridge
stpPriority            - Set the Spanning Tree bridge priority
stpMaxAge             - Set the Spanning Tree bridge maximum age
stpHelloTime          - Set the Spanning Tree bridge hello time
stpForwardDelay       - Set the Spanning Tree bridge forward delay
stpGroupAddress       - Set the Spanning Tree bridge group address
port                  - Administer bridge ports
packetFilter          - Administer packet filters

Type ESC to return to the previous menu and ? for help.
-----
Select menu option (bridge):port
Menu options: -----
summary               - Display summary information
detail                - Display detailed information
multicastLimit        - Set the multicast packet rate limit
stpState              - Enable/Disable Spanning Tree on a port
stpCost               - Set the Spanning Tree path cost
stpPriority            - Set the Spanning Tree port priority
address               - Administer bridge addresses

Type ESC to return to the previous menu and ? for help.
-----
Select menu option (bridge/port):address
Menu options: -----
list                  - List addresses
add                    - Add a statically configured address
remove                - Remove an address
find                  - Find an address
flushall              - Flush all static and dynamic addresses
flushDynamic          - Flush all dynamic addresses
freeze                - Make all dynamic addresses static

Type ESC to return to the previous menu and ? for help.

```




Figure 2-2 Administration Console for the Switch 2200 System

For more information about the Administration Console, see the *SuperStack™ II Switch 2200 Administration Console User Guide*.

External Network Management Applications

3Com's Transcend® Enterprise Manager is a network management software that runs on UNIX and MS-DOS platforms. It provides network management for a wide range of 3Com products, including the Switch 2200. With Transcend Enterprise Manager software, you get a device view of the Switch 2200 so you can display the operating status, configure, and get statistics about each device.

Transcend Enterprise Manager software helps you monitor and manage the performance of your switching hub-based network. You can also display statistical graphs showing your network's status and analyze historical data. See Figure 2-3. To order Transcend Enterprise Manager, contact your sales representative.

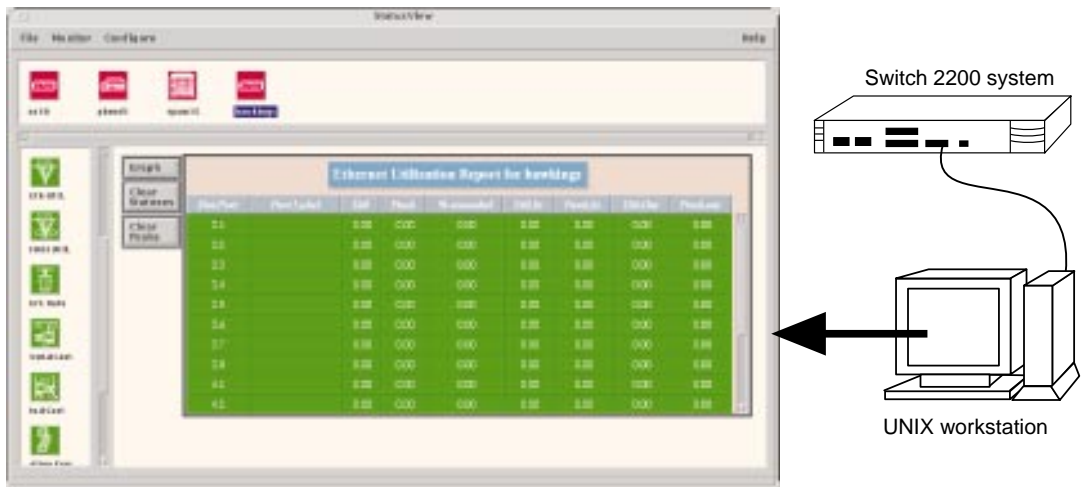


Figure 2-3 3Com's Transcend® Enterprise Manager (SNMP Manager) Main Window

Because the Switch 2200 system is based on standards, you can also use other SNMP-based network manager applications, such as Sun Microsystems' SunNet Manager™, Hewlett-Packard OpenView™, or IBM's NetView AIX®.

3

MANAGEMENT ACCESS: PROTOCOLS

This chapter describes the underlying communication and management protocols used to deliver management and administration data to and from your SuperStack™ II Switch 2200 system.

About Switch 2200 Protocols

The Switch 2200 uses the following protocols:

- “Virtual terminal” protocols, such as rlogin and telnet
- Simple Network Management Protocol (SNMP)
- FDDI Station Management (SMT) protocol

Figure 3-1 highlights these protocols and puts them into perspective in the network environment.

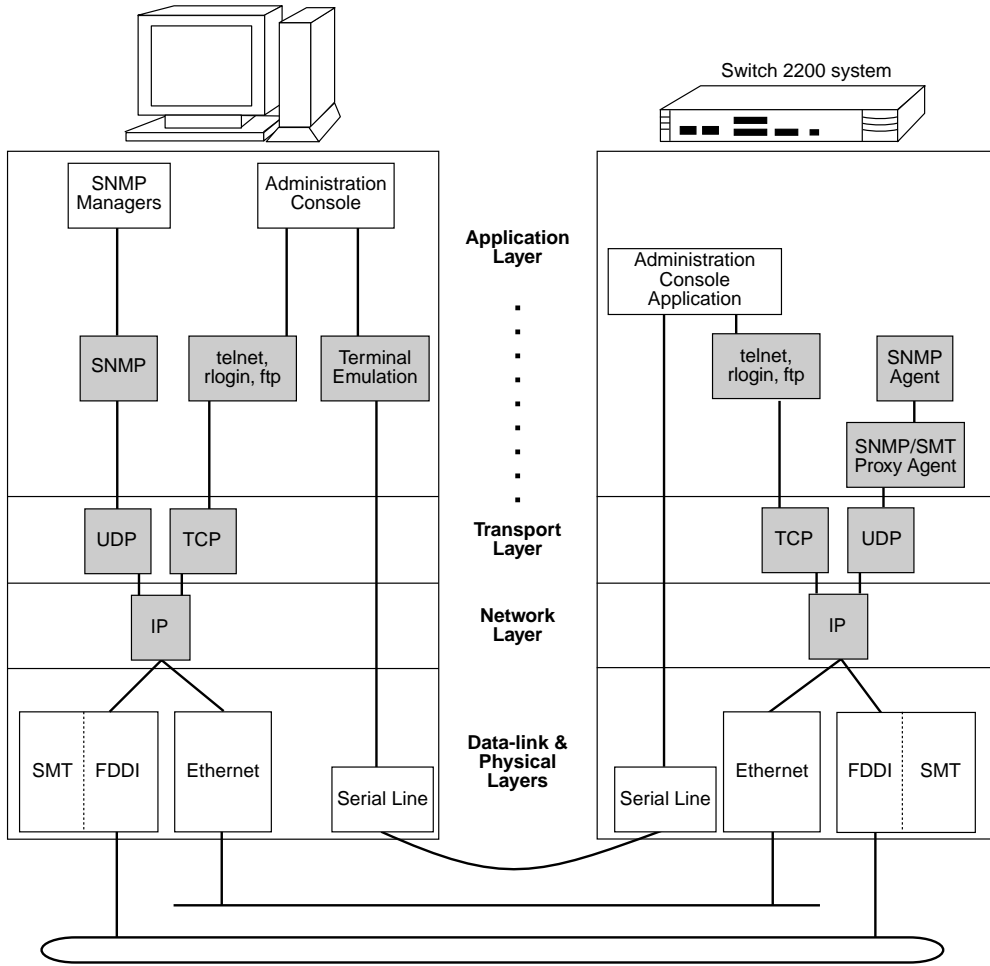


Figure 3-1 Protocol Stacks for the Switch 2200 System

Virtual Terminal Protocols

A virtual terminal protocol is a software program, such as rlogin or telnet, that allows you to establish a management session from a PC or a UNIX workstation. Because rlogin and telnet run over TCP/IP, you must have at least one IP address configured on the Switch 2200 system before you can establish access to it with a virtual terminal protocol. Within the Administration Console, you configure an IP address by defining an IP interface.

Terminal emulation differs from a virtual terminal protocol in that it connects a terminal directly to the serial line.

Figure 3-2 shows a UNIX workstation connecting to a Switch 2200 system through a virtual terminal protocol and a terminal connecting through a null modem cable.




Figure 3-2 Administration Console Access for the Switch 2200 System

SNMP

Simple Network Management Protocol (SNMP) is the standard management protocol for multivendor IP networks. SNMP supports transaction-based queries that allow the protocol to format messages and to transmit information between reporting devices and data-collection programs. It runs on top of the User Datagram Protocol (UDP), offering a connectionless-mode service.

SNMP Agent Each Switch 2200 system has an SNMP agent that provides access to management information maintained by the system. The SNMP agent responds to requests from an external manager, such as Transcend Enterprise Manager software. The agent also reports network events.

SNMP MIBs You can access the information that is defined in industry-standard and enterprise-specific (proprietary) Management Information Bases (MIBs) supported by the Switch 2200. These MIBs are collections of related managed objects (abstract representations of resources that are capable of being managed). Some examples of these resources are Ethernet and FDDI ports and bridges. The Switch 2200 supports the following SNMP MIBs:

- MIB II
- Ethernet MIB
- FDDI SMT 7.3 MIB
- Bridge MIB
- LANplex Systems MIB
- LANplex Optional FDDI MIB

For more information on which MIBs are supported, see Appendix A: *SNMP MIB Support* and the software *Release Notes*.

SNMP Traps

An SNMP trap is an asynchronous report of one of several events. To receive reports, you must configure the IP address of the management station to which the reports are sent; otherwise, the reports are discarded. Through SNMP and your network management software or through the Administration Console, you can configure which traps are sent to which IP addresses.

Table 3-1 lists the SNMP traps supported by the Switch 2200 system.

Table 3-1 SNMP Traps

Group	Trap
MIB II	coldStart
	authenticationFailure
Bridge MIB	newRoot
	topologyChange
LANplex® System MIB	lpsSystemOverTemperatureEvent
	lpsBridgeAddressThresholdEvent
LANplex® Optional FDDI MIB (SMT 7)	ISMTHoldCondition
	ISMTPeerWrapCondition
	MACDuplicateAddressCondition
	MACFrameErrorCondition
	MACNotCopiedCondition
	MACNeighborChangeEvent
	MACPathChangeEvent
	PORTLerCondition
	PORTUndesiredConnAttemptEvent
	PORTEBErrorCondition
PORTPathChangeEvent	

For descriptions of the traps, see the ASN.1 MIB definition files included with your software release.

Access Control Access to system information through SNMP is controlled by community strings. A community string is a character string included in each SNMP protocol message sent between your Switch 2200 system and external management applications like Transcend Enterprise Manager.

A community string identifies a particular group of SNMP managers with certain access rights. The SNMP agent in the Switch 2200 system allows the configuration of two community strings: one that provides access to read system information but not to change system parameters and one that provides access to read system information *and* to configure system parameters. To set up the Switch 2200 system to work with an SNMP manager, you must configure the Switch 2200 system's SNMP community strings to match those used by the SNMP manager.

For information on how to configure community strings, see the *SuperStack™ II Switch 2200 Administration Console User Guide*.

SMT

Station Management (SMT) for FDDI is a standard that specifies a set of services and signalling mechanisms dedicated to FDDI LAN management. It is responsible for managing the services of an FDDI station that are specific to the MAC, PHY, and PMD layers of the OSI Reference Model. The goal of SMT is to define shared medium management services to guarantee the interoperability of FDDI network equipment from multiple vendors.

The Switch 2200's implementation of SMT supports the full SMT MIB as defined by ANSI X3T9.5, including the many optional attributes. This MIB is accessible remotely by using SMT frames or SNMP frames.

You can set some SMT FDDI MIB parameters through the Administration Console. See the *SuperStack™ II Switch 2200 Administration Console User Guide*.

SNMP and SMT Proxy Agents

A proxy agent acts as a management gateway. It converts requests and event reports from one protocol and object format to another protocol and object format.

Your Switch 2200 system contains a proxy agent that translates between SNMP and FDDI SMT. It allows a network management station that is not necessarily connected directly to an FDDI LAN to manage FDDI end-stations on that LAN, even if the FDDI end-stations do not support SNMP. For all the stations in that FDDI LAN to be managed, only one Switch 2200 proxy agent needs to be active on each FDDI LAN within your network.

4

PHYSICAL ACCESS: PORTS AND CABLING

This chapter explains how you can manage your SuperStack™ II Switch 2200 system through its physical interfaces. Figure 4-1 highlights the system's physical access options.

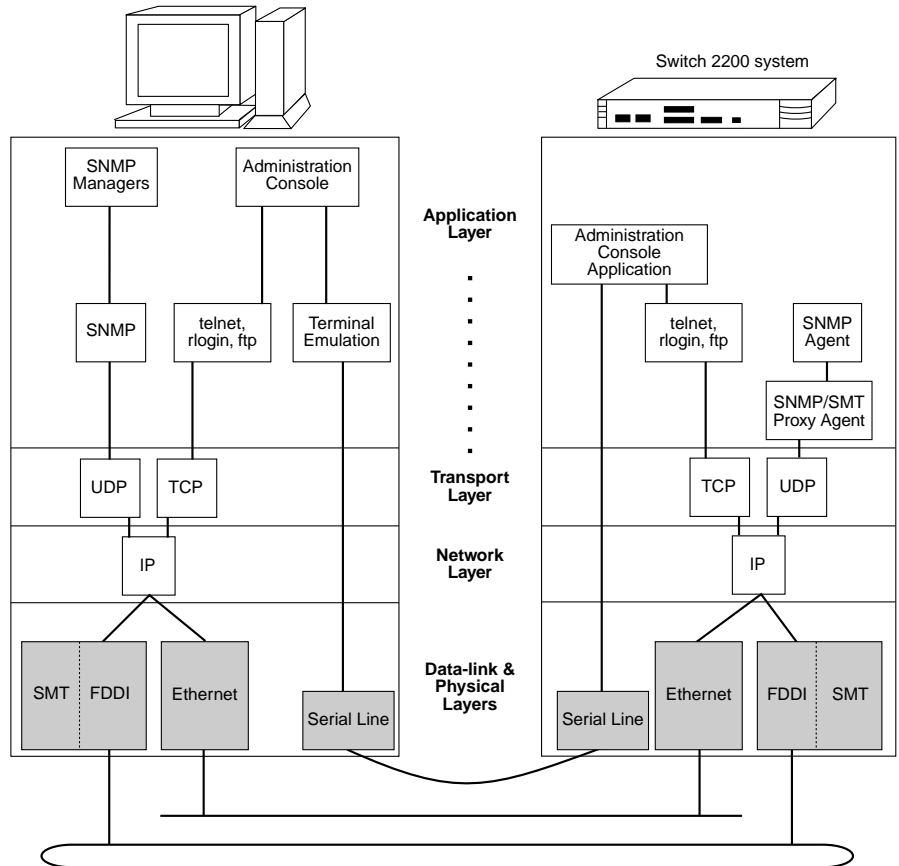


Figure 4-1 Physical Access Options for the Switch 2200 system

In-band and Out-of-band Management

If you manage your Switch 2200 system and its attached LANs over the same network that carries your regular data traffic, then you are managing your network *in-band*. This is often the most convenient and inexpensive way to access your Switch 2200 system. The disadvantage of using in-band management is that if your data network is faulty, you might not be able to diagnose the problem because the management requests are sent over that same faulty network. The Switch 2200 system supports in-band management by default.



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

If you are using a dedicated network outside your Switch 2200 system and its attached LANs for management data, then you are managing your network *out-of-band*. For more information on system management, see the *SuperStack™ II Switch 2200 Administration Console User Guide*.

Management Access

You can access the Switch 2200 through the Console port or through an Ethernet or an FDDI port. These methods are described next.

Console Serial Port

Direct access through the console serial port is often preferred because it allows you to stay attached during system boots. A Macintosh or PC attachment can use any terminal emulation program when connecting to the console serial port. A UNIX workstation can use the emulator TIP. See Figure 4-2.

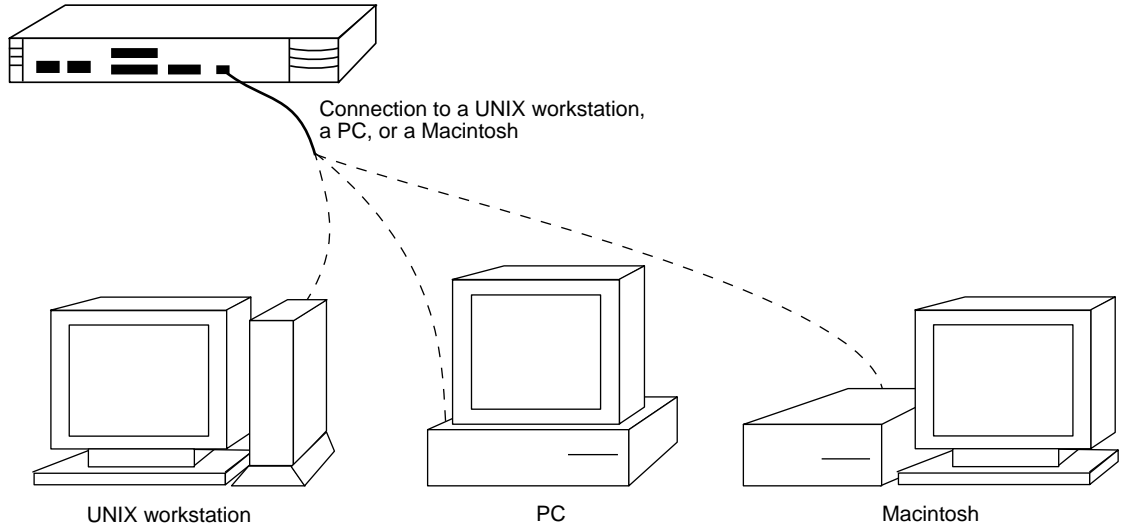


Figure 4-2 Access Through the Console Port

Ethernet and FDDI Ports

Using the rlogin or telnet interfaces, you can access the Administration Console through any Ethernet or FDDI port if an IP address is assigned to it. The SNMP agent can also be accessed through these interfaces.

Figure 4-3 shows access through an Ethernet or an FDDI port.






Figure 4-3 Access to the Switch 2200 Through an Ethernet or an FDDI Port



BRIDGING

- Chapter 5** Transparent Bridging
- Chapter 6** User-defined Packet Filtering
- Chapter 7** Bridging Extensions



5

TRANSPARENT BRIDGING

This chapter describes the operation of a transparent bridge, including how a transparent bridge:

- Learns addresses
- Ages addresses
- Forwards packets
- Prevents loops in a network

About Transparent Bridging

A transparent bridge allows two or more LANs to be interconnected and to communicate as if they were one LAN. The bridge listens promiscuously to packets on another LAN. A packet is never retransmitted onto the LAN from which it was sourced.

Transparent bridging has been adopted for standardization by the IEEE and is defined in the IEEE 802.1d specification.

What Makes a Bridge 802.1d Compliant?

The IEEE 802.1d bridging standard specifies many requirements with which a transparent bridge must comply. An 802.1d bridge must:

- Learn source addresses from packets transmitted by stations on attached LANs
- Age addresses of stations on attached LANs that have not transmitted a packet for a prolonged period of time
- Store and forward packets from one LAN to another
- Use the Spanning Tree Protocol for loop detection

The Switch 2200 system complies with all IEEE 802.1d bridging requirements.

How a Bridge Learns Addresses

Bridges learn addresses so that they can make intelligent decisions about which packets to forward from one bridge port to another. A bridge automatically learns addresses by listening on the network. For a bridge to learn the address of a station on the network, that station must transmit a packet. Each bridge maintains a dynamic table, called the address table, which contains all learned source addresses.

When a bridge receives a packet, it looks up the packet's source address in the address table, and does one of the following:

- *If the source address is known* to the bridge, then the bridge updates the source address entry in the address table and verifies the port on which the packet was received.
- *If the source address is not known* to the bridge, then the bridge stores the packet's source address in the address table, along with the port on which the packet was received. See Figure 5-1.

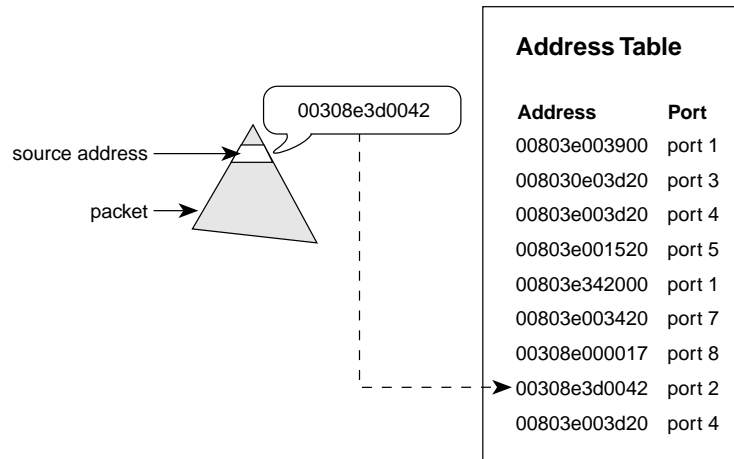


Figure 5-1 Learning Source Addresses

How a Bridge Ages Addresses

A source address remains in the address table as long as the station to which it relates regularly transmits through the bridge. If the station does not regularly transmit, the source address is "aged out" of the bridge's table. Address aging is primarily implemented to ensure that if a station moves to a different segment on the network, its address will be forgotten at the old location and packets will no longer be forwarded to that location. Address aging is also necessary because a bridge can learn only a finite number of addresses. The Switch 2200 system, when configured as an IEEE 802.1d bridge, can learn up to 8K addresses in its address table.

Address aging, although typically an efficient means of maintaining a current address table, can create problems when regularly used stations on the network do not transmit periodically. For instance, printers only transmit when they are powered on, yet printing is a function performed frequently on a network. In this case, the printer's address is aged out of the address table and the bridge no longer has the information it needs to send packets directly to that station.

To handle this situation, the Switch 2200 system allows you to statically configure the addresses of these stations. Because a statically configured address is not aged out of memory, it must be manually flushed when the station is removed from the network. Static configuration of Ethernet addresses and flushing static Ethernet addresses are described in the *SuperStack™ II Switch 2200 Administration Console User Guide*.

Packet Forwarding

A bridge either filters, floods, or forwards packets by comparing the packet's destination address to the addresses in the bridge's address table, and by comparing the destination bridge port (if known) to the port on which the packet was received. This process is described and shown in Figure 5-2.

The bridge compares the destination address to the addresses in the address table and does one of the following:

- *If the destination address is known* to the bridge, then the bridge identifies the port on which the destination address is located.
 - If the destination bridge port is *different* from the bridge port on which the packet was received, then the packet is forwarded to the destination bridge port.
 - If the destination bridge port is the *same* as the port on which the packet was received, then the packet is filtered (discarded) by the bridge.
- *If the destination address is not known* to the bridge, the packet is forwarded to all active bridge ports other than the bridge port on which the packet was received. This process is called flooding. For a port to be active, it must be enabled and in the forwarding state. See the section "Spanning Tree Port States" on page 5-16 for more information about states.

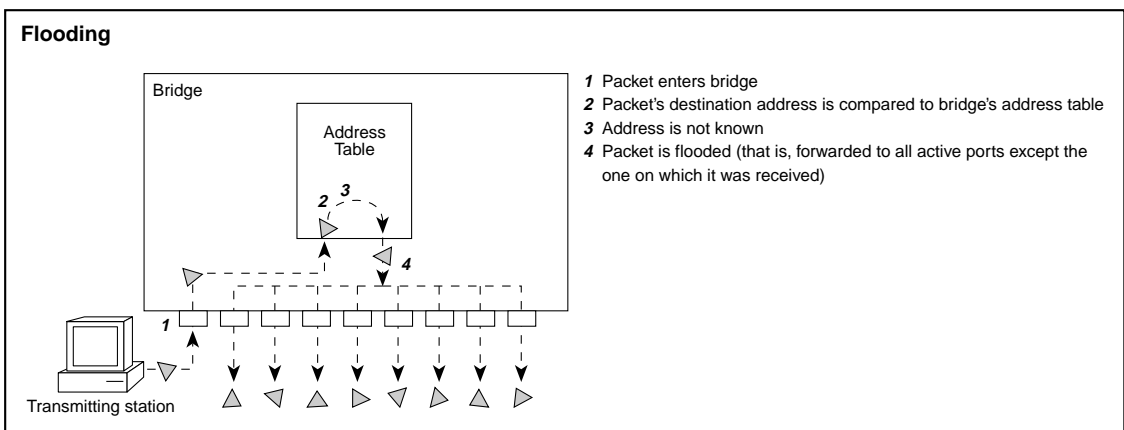
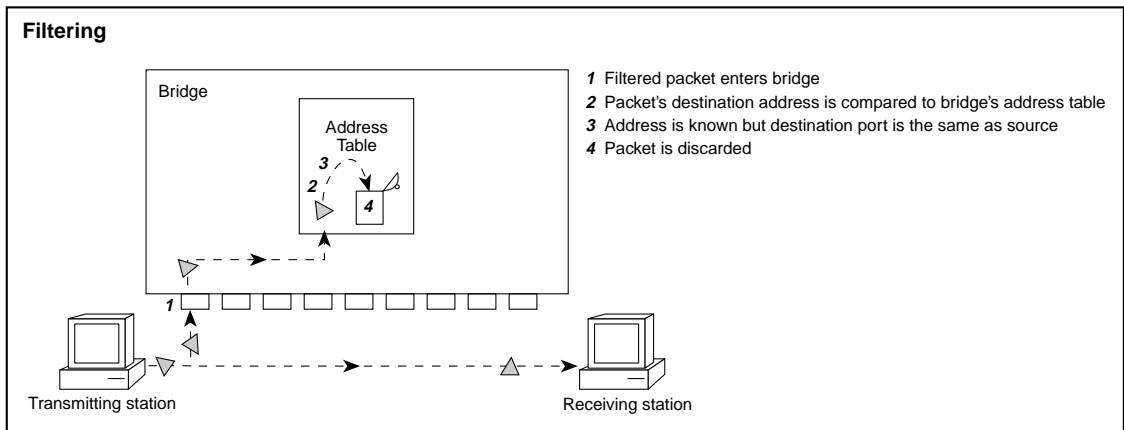



Figure 5-2 Forwarding, Filtering, and Flooding Packets

Spanning Tree and the Bridged Network

When transparent bridges are used to attach networks with redundant links, packets can loop and rapidly multiply on the attached LANs. These additional packets create traffic that might unnecessarily clog the LAN.

A loop exists if more than one path can be used to forward a packet from one station to another. To solve this problem, IEEE 802.1d bridging includes Spanning Tree Protocol, an algorithm that dynamically maps out a loopless network topology (a subset of the entire topology), ensuring that only one active path exists between every pair of LANs.

Packet Looping in a Bridged Network

Loops can occur on a bridged network for various reasons. In a network where reliability is key, network administrators often implement redundant links so that, although individual bridges might fail, the “networks” (data pathways) between stations remain active. Loops can also occur by accident. For instance, when more than one bridge is used to connect various LANs, the network manager might inadvertently configure the extended network with loops, causing packets to be circulated indefinitely.

A potential example of packet looping is shown in Figure 5-3. In this example:

- 1 Packet 1 is transmitted on LAN 1.
- 2 Bridges A, B, and C (connected to both LAN 1 and LAN 2) receive Packet 1 and forward it onto LAN 2, creating packets 1a, 1b, and 1c, respectively.
- 3 Bridge A receives Packets 1b and 1c on LAN 2 and forwards them onto LAN 1; at the same time, Bridge B receives Packets 1a and 1c on LAN 2 and forwards them onto LAN 1. Bridge C follows this same pattern.

When multiple bridges receive the same packet, they each transmit a new copy of the packet onto the attached LANs. Consequently, the packets will loop and multiply indefinitely as they traverse the bridges.

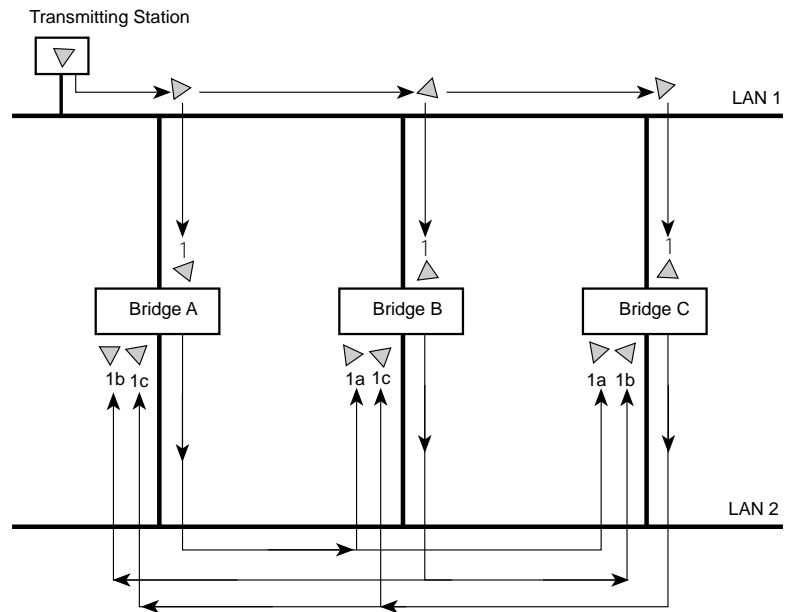


Figure 5-3 Packets Looping and Multiplying without Spanning Tree

The Spanning Tree Algorithm

The Spanning Tree algorithm detects loops and logically blocks (eliminates) redundant paths by putting some bridge ports in the blocking state so that only one path exists between any two LANs and, therefore, between any two stations. See Figure 5-4. A port in the blocking state neither forwards nor receives data packets.

After the algorithm eliminates extra paths, the network configuration stabilizes. When one or more of the bridges or communication paths in the stable topology fail, the protocol automatically recognizes the changed configuration and activates redundant links. This ensures that all stations remain connected.

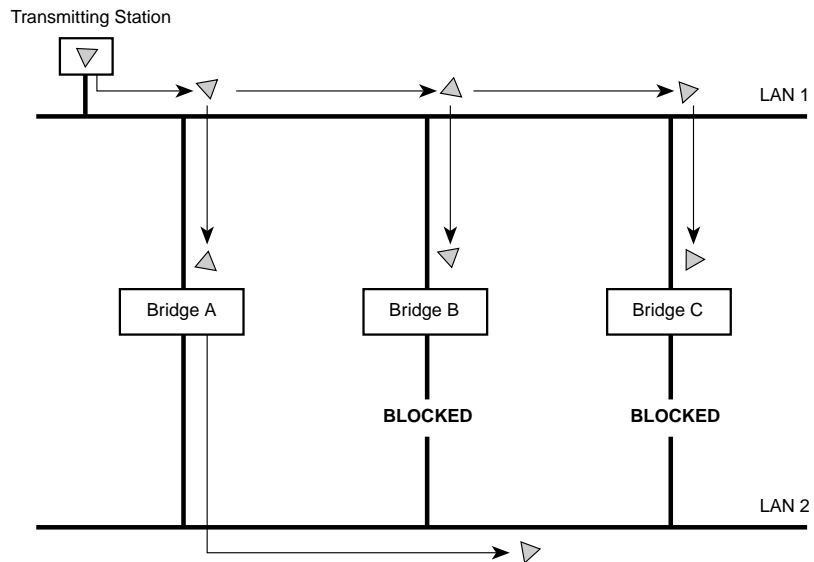


Figure 5-4 Spanning Tree Implemented to Block Redundant Links

How the Spanning Tree Algorithm Works

The Spanning Tree algorithm is based on the idea that bridges transmit messages to each other that allow them to calculate the Spanning Tree topology. These messages are special packets called *Configuration Bridge Protocol Data Units (CBPDUs)*, or configuration messages. CBPDUs are not propagated through the bridge like regular data packets. Instead, each bridge behaves as an end-station for these packets — receiving and interpreting them.

CBPDUs at work

The CBPDUs help the bridges establish a hierarchy among themselves (or a calling order) for the purposes of creating a loopless network. Based on the information in the CBPDUs, the bridges elect a *root bridge*, which is at the top level of the hierarchy. The bridges then choose the best path on which to transmit information to the root bridge.

The bridges chosen as the best path, called *designated bridges*, are the second level of the hierarchy. A designated bridge “relays” the network transmissions to the root bridge through its *root port*. Any port that transmits to the root bridge is a root port. The designated bridges also have *designated ports* — the ports attached to the LANs from which the bridge is

receiving information. Figure 5-5 shows the hierarchy of the Spanning Tree bridges and their ports.




Figure 5-5 Hierarchy of the Root Bridge and the Designated Bridge

From the information that the CBPDUs provide, the bridges:

- Elect a single bridge to be the *root bridge*. The root bridge has the lowest bridge ID among all the bridges on the extended network.
- Calculate the best path between themselves and the root bridge.
- Elect a *designated bridge* on each LAN from among the bridges residing on that LAN. This is the bridge with the least cost path to the root bridge. Its function is to forward packets between that LAN and the path to the root bridge. For this reason, the root bridge is *always* the designated bridge for its attached LANs. The port through which the designated bridge is attached to the LAN is elected the *designated port*.
- Choose a *root port* that gives the best path from themselves to the root bridge.
- Select ports to be included in the Spanning Tree topology. The ports selected include the root port plus any designated ports. Data traffic is forwarded to and from ports selected for inclusion in the Spanning Tree topology. Data traffic is never forwarded to or received on ports that are not selected for inclusion in the Spanning Tree topology.

Figure 5-6 shows a bridged network with its Spanning Tree elements.

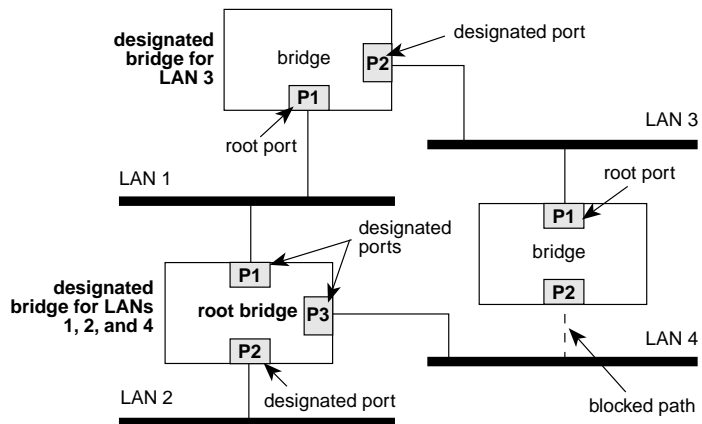


Figure 5-6 Root and Designated Bridges and Ports in a Spanning Tree Topology

CBPDU's contents

The specific information that bridges receive from the CBPDU allows them to calculate a Spanning Tree topology:

- **Root ID** — The identification of the bridge assumed to be the root
- **Cost** — The cost of the least-cost path to the root from the transmitting bridge. One of the determining factors in cost is the speed of the bridge's network interface. In this case, the faster the speed, the lower the cost.
- **Transmitting bridge ID** — The identification of the bridge transmitting this CBPDU. The bridge ID consists of the bridge address and the bridge priority
- **Port identifier** — The port priority plus the number of the port from which the transmitting bridge sent a CBPDU. It is only used in the Spanning Tree calculation if the root IDs, transmitting bridge IDs, and costs (when compared) are equal. In other words, the port identifier is a tie breaker in which the lowest port identifier takes priority. This field is primarily useful for selecting the preferred port when two ports of a bridge are attached to the same LAN or when two routes are available from the bridge to the root bridge.

Comparing CBPDUs

Here are some examples showing how the best CBPDU is determined by the bridge. The root ID is the most important determining factor. If the root ID fields are equal, then the cost is compared. The last determining factor is the transmitting bridge ID. If the CBPDUs all have the same root ID, cost, and transmitting bridge ID, then the port identifier is used as a tiebreaker.

Example 1. Message 1 has a lower root ID, so it is saved by the bridge.

Message 1			Message 2		
root ID	cost	transmitter	root ID	cost	transmitter
12	15	35	31	12	32

Example 2. Root ID is the same for both messages, but cost is lower in Message 1. Message 1 is saved.

Message 1			Message 2		
root ID	cost	transmitter	root ID	cost	transmitter
29	15	80	29	18	38

Example 3. Root ID and cost are the same for both messages, but the transmitting bridge ID is lower in Message 1. Message 1 is saved.

Message 1			Message 2		
root ID	cost	transmitter	root ID	cost	transmitter
35	80	39	35	80	40

How a bridge handles CBPDUs

The following case describes how one bridge interprets CBPDUs, thus contributing to the Spanning Tree configuration. For purposes of this case, the following convention is used to depict a CBPDU: *root ID.cost.transmitter ID*.

- 1 When Spanning Tree is first started on a network, the bridge thinks that it is the root bridge and transmits a CBPDU from each of its ports with the following information:
 - Its own bridge ID as the root ID (for example, 85)
 - Zero (0) as the cost (because it thinks it is the root bridge)
 - Its own bridge ID as the transmitting ID (for example, 85)

This CBPDU looks like: *85.0.85*.

- 2 The bridge receives CBPDUs on each of its ports from all other bridges. It saves the "best" CBPDU from each port. The best one is determined by comparing the information in each message arriving at a particular port to

the message the bridge currently has stored at that port. In general, the lower the values of the CBPDU, the “better” it is. When the bridge comes across a better CBPDU than it has stored, it replaces the old message with the new one.

- 3 From the messages received, the bridge determines which bridge is the root bridge. For example, if the bridge receives a CBPDU with the contents 52.0.52, then it would assume that the bridge with the ID 52 is the root (because its root ID is smaller).
- 4 Because the bridge now knows the root bridge, it can determine its distance to root and elect a root port. It examines CBPDUs from all ports to see which port has received a CBPDU with the smallest cost to the root. This port becomes the root port.
- 5 Now that the bridge knows what its own CBPDU contains, it can compare this updated CBPDU with the ones received on its other ports. If the bridge’s message is better than the ones received on any of its ports, then the bridge assumes that it is the designated bridge for the attached LANs.

If the bridge receives a better CBPDU on a port than the message it would transmit, it no longer transmits CBPDUs on that LAN. When the algorithm stabilizes, only the designated bridge transmits CBPDUs on that LAN.

How Spanning Tree Is Calculated for the Network

The following example illustrates how the Spanning Tree algorithm determines the Spanning Tree configuration on an entire network.

Determining the root bridge and root ports

In Figure 5-7, the network topology consists of six bridges connecting six LANs. The topology is designed with redundant links for backup purposes, which creates four loops in the extended network. When the Spanning Tree algorithm first runs, each bridge transmits a CBPDU that contains its bridge ID as both the *root ID* and the *transmitting bridge ID*, and zero as the *cost*.

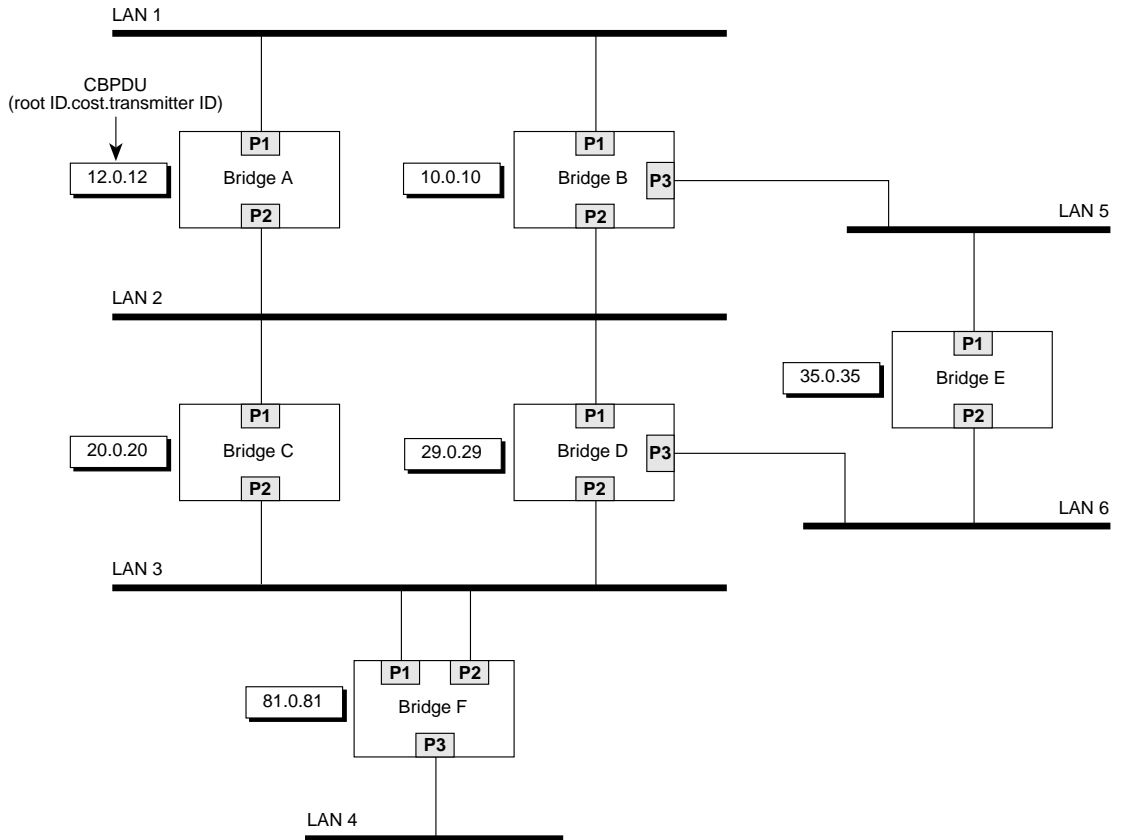


Figure 5-7 Starting the Spanning Tree Calculation

The root ID portion of the CBPDUs determines which bridge will be the root bridge. The bridges transmit their CBPDUs, receive each other's CBPDUs, and compare the CBPDUs to each other. Because Bridge B has the lowest root ID of all the bridges, it becomes the root. See Figure 5-8.

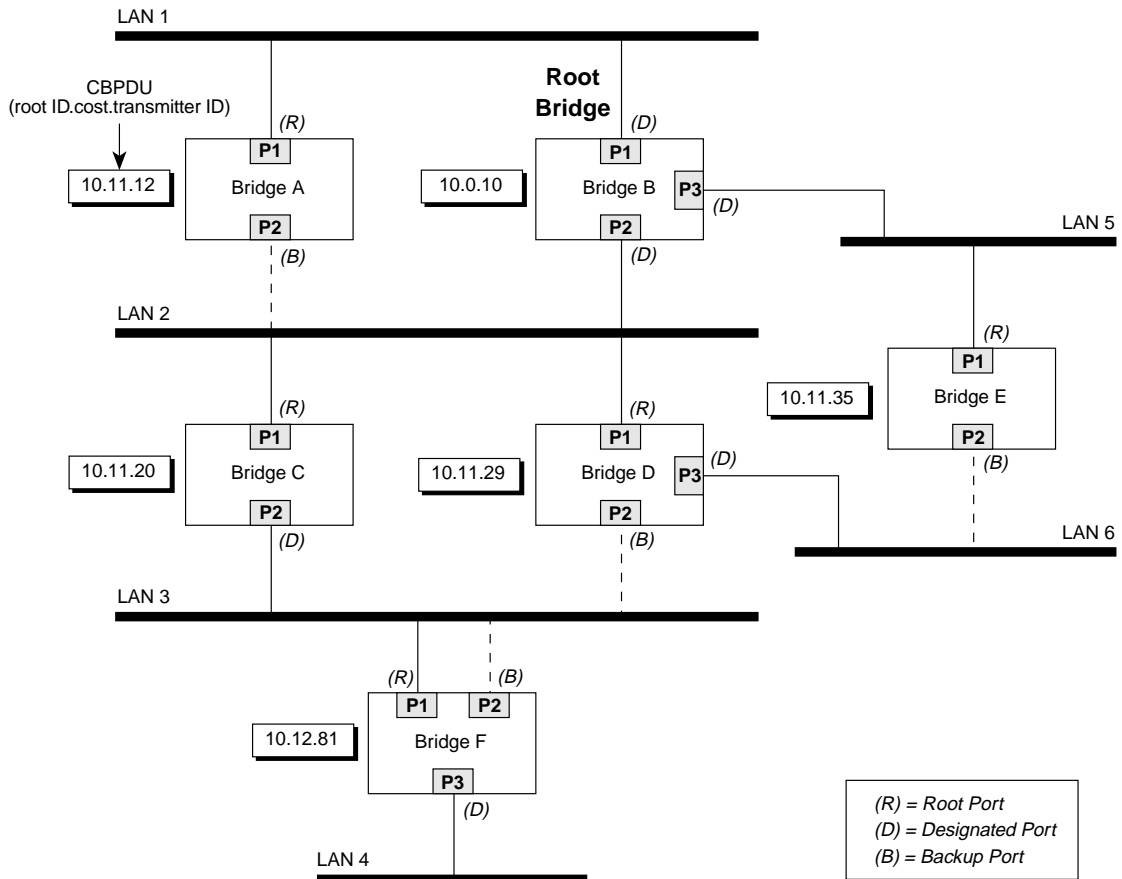


Figure 5-8 Spanning Tree Topology Calculated

Each bridge, except for the root bridge, must select a root port. To do this, each bridge determines the most cost-effective path for packets to travel from each of its ports to the root bridge. The cost depends on 1) the port's path cost, and 2) the root path cost of the designated bridge for the LAN to which this port is attached.

If the bridge has more than one port attachment, the port with the lowest cost becomes the root port, and the other ports become either designated or backup ports. If bridges have redundant links to the same LAN, then the port with the lowest port identifier becomes the root port. In Figure 5-8, Bridge F has two links to LAN 3 (through port 1 and port 2). Because the lowest port identifier for Bridge F is port 1, it becomes the root port, and port 2 becomes a backup port to LAN 3.

*Determining the
designated bridge
and designated ports*

If a LAN is attached to a single bridge, that bridge is the LAN's designated bridge. For a LAN that is attached to more than one bridge, a designated bridge must be selected from among the attached bridges. The root bridge is automatically the designated bridge for all the attached LANs.

For example, Bridge B, the root bridge in Figure 5-8, is also the designated bridge for LANs 1, 2, and 5. A designated bridge must still be determined for LANs 3, 4, and 6. Because Bridges C, D, and F are all attached to LAN 3, one of them must be the designated bridge for that LAN. The algorithm first compares the root ID of these bridges, which is the same for all. The cost is then compared. Bridge C and Bridge D both have a cost of 11. Bridge F, with a cost of 12 is eliminated as the designated bridge. Finally, the transmitting bridge ID is compared between Bridge C and Bridge D. Because Bridge C's ID (20) is smaller than Bridge D's (29), Bridge C becomes the designated bridge for LAN 3.

The designated bridge for LAN 6 is either Bridge D or Bridge E. Because Bridge D's transmitting bridge ID (29) is lower than Bridge E's (35), Bridge D becomes the designated bridge for that LAN. Finally, the designated bridge for LAN 4 is the only bridge attached to that LAN, Bridge F.

The designated port is determined by the port that attaches the designated bridge to the LAN. If there is more than one port attached to the LAN, then the port identifier determines which port is the designated port.

Spanning Tree Port States

As the Spanning Tree algorithm determines the Spanning Tree configuration, it places ports in the following states: listening, learning, forwarding, blocking, or disabled. As changes occur in the network, the port may transition in and out of these states to maintain a loopless network. These states are described in Table 5-1.

Table 5-1 Spanning Tree Port States

Port State	Description
Listening	<p>When Spanning Tree is configuring, all ports are placed in the listening state. Each port remains in this state until the root bridge is elected. While in the listening state, the bridge continues running the Spanning Tree algorithm and transmitting CBPDUs on the port; however, it discards data packets received on that port and does not transmit data packets from that port.</p> <p>The listening state should be long enough for a bridge to hear from all other bridges on the network (this time can be adjusted if necessary). After the time of the listening state, the bridge ports that are to proceed to the forwarding state go into the learning state. All other bridge ports go into the blocking state.</p>
Learning	<p>The learning state is similar to the listening state except that data packets are received on that port for the purpose of learning stations attached to that port. After spending the specified time in this state, if the bridge has still not heard any information that would make it transition the port back to the blocking state, then the bridge transitions the port to the forwarding state.</p> <p>The time the port spends in both the listening and learning states is determined by the value of the <i>forward delay</i> parameter. Forward delay is a timer that 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 the network. This delay gives all links that need to be turned off in the new topology time to do so before new links are turned on.</p>
Forwarding	<p>Once in the forwarding state, the bridge performs standard bridging functions. It receives packets and either forwards or does not forward them, depending on address comparisons between the packet's destination address and the addresses in the bridge's address table.</p>
Blocking	<p>When a port is put in a blocking state, the bridge continues to receive CBPDUs on that port (monitoring for network reconfigurations), but it does not transmit them. Additionally, the bridge does not receive data packets from the port, learn locations of station addresses from it, or forward packets onto it.</p>

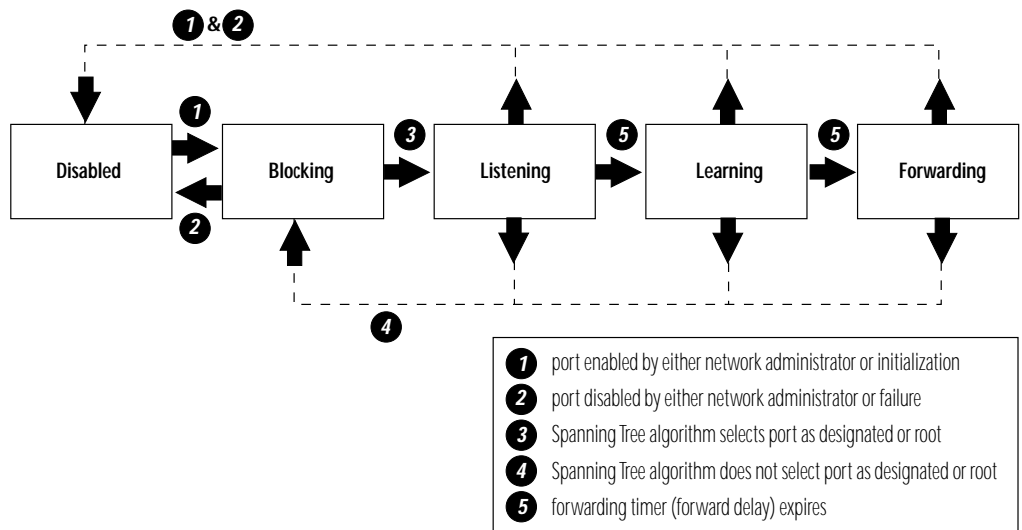
(continued)

Table 5-1 Spanning Tree Port States (continued)

Port State	Description
Disabled	A port is disabled when Spanning Tree has been turned off for that specific port or when the port has failed. In the disabled state, the port does not participate in the Spanning Tree algorithm. If Spanning Tree has been turned off for a specific port, that port will continue to forward frames only if Spanning Tree is disabled for the entire bridge.

Figure 5-9 illustrates the factors that cause a port to transition from one state to another. The arrows indicate the direction of movement between states. The numbers correspond to the factors that affect the transition.

For example, for a port in the blocking state to transition to the listening state, the Spanning Tree algorithm must select that port as a designated or root port. Once in the listening state, forward delay must expire before the port can transition to the learning state. When in listening, learning, and forwarding states, if a port is disabled by the network administrator or by a failure or initialization, then that port becomes disabled.

**Figure 5-9** Factors Involved in Spanning Tree Port State Transitions

Reconfiguring the Bridged Network Topology

The Spanning Tree algorithm reconfigures the bridged network topology when 1) bridges are added or removed, 2) the root bridge fails, or 3) the network administrator changes the bridging parameters that determine the topology.

Whenever a designated bridge detects a topology change, it sends out a Topology Change Notification Bridge Protocol Data Unit (BPDU) through its root port. This information is eventually relayed to the root bridge. The root bridge then sets the Topology Change Flag in its CBPDU so that the information is broadcast to all the bridges. It transmits this CBPDU for a fixed amount of time to ensure that all bridges are informed of the topology change.

If a port transitions from the blocking state to the forwarding state as a result of the topology change, the algorithm ensures that it sends the topology information to all of the ports before that port starts forwarding data. This delay prevents temporary data loops.

As a result of a network reconfiguration, the bridge flushes all addresses from the address table. This action ensures that each active port still forwards packets to the right network after a topology change.

Bridging References

IEEE 802.1d MAC Bridges. D9, July 14, 1989.

Perlman, Radia. *Interconnections: Bridges and Routers*. Reading, Massachusetts: Addison-Wesley Publishing Company, Inc., 1992.

6

USER-DEFINED PACKET FILTERING

This chapter contains the following information:

- A description of user-defined packet filtering
- A discussion of how to use address and port groups in packet filters
- Examples of packet filters

About User-defined Packet Filtering

The SuperStack™ II Switch 2200 system allows you to add a second layer of packet filtering on top of the standard filtering provided by a traditional transparent bridge. This *user-defined* packet filtering further restricts which packets are forwarded through the bridge. By taking advantage of this powerful feature, you can improve network performance, provide additional security, or logically segment your network to support virtual workgroups.

Designing a Packet Filter

The packet filtering mechanism supported on the Switch 2200 is very flexible. You can define complex filters comprising many simple comparisons. This flexibility allows you to use packet filters in several unique applications on your network.

You specify the packet filter using a *packet filter language*. This language consists of operands and operators that you use to compose your filters. This language is described in detail in the *SuperStack™ II Switch 2200 Administration Console User Guide*. Table 6-1 describes the two simplest operands.

Table 6-1 Packet Filter Operands

Operand	Description
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.
Constant	A literal value. As with a field, a constant can be 1, 2, 4, or 6 bytes.

The operators that you specify in the packet filter allow the filter to make a logical decision about whether to forward or discard the packet. Table 6-2 describes these operators.

Table 6-2 Packet Filter Operators

Operator	Result
equal	true if operand 1 = operand 2
not equal	true if operand 1 \neq operand 2
less than	true if operand 1 < operand 2
less than or equal	true if operand 1 \leq operand 2
greater than	true if operand 1 > operand 2
greater than or equal	true if operand 1 \geq operand 2
not	true if operand 1 = false
and	operand 1 bit-wise AND operand 2
or	operand 1 bit-wise OR operand 2
exclusive or	operand 1 bit-wise XOR operand 2
shift left	operand 1 SHIFT LEFT operand 2
shift right	operand 1 SHIFT RIGHT operand 2



*The operators **and**, **or**, and **exclusive or** are bit-wise operators, which means that the corresponding bits of each of the operands are logically compared to produce the resulting bit.*

Assigning Packet Filters to Paths

For a packet filter to be used by the bridge, you must install it on a specific bridge port. You can put the filter on the bridge port's *receive* or *transmit* paths depending on when you want the filter to be applied. Placing the 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 is discarded if it does not meet the forwarding criteria defined in the filter. See Figure 6-1.

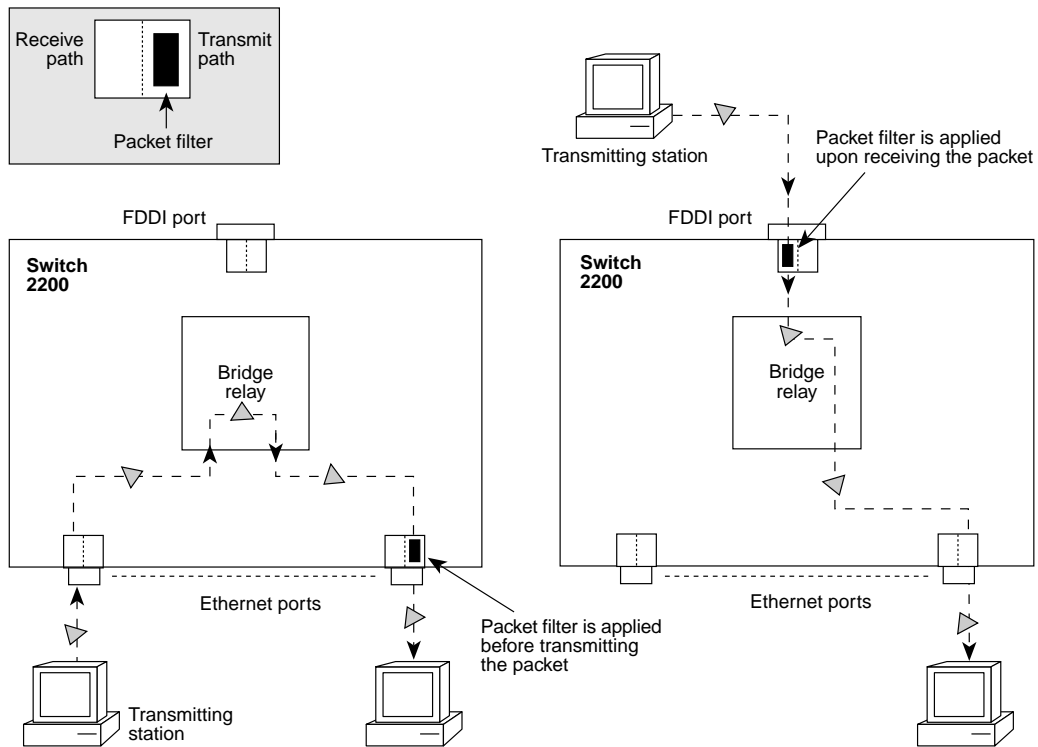


Figure 6-1 Assigning Filters to Paths

Packet Filter Examples

The following examples show you how to define and apply packet filters to ports and paths on your network. Example 1 isolates protocols in a network. Example 2 shows a more complex version of the filter in the first example, showing how expressions are evaluated in a packet filter.

Example 1: Isolating IP Segments

The network shown in Figure 6-2 is composed of two types of protocols:

- The internet protocol (IP), over which Sun workstations and a compute server communicate
- AppleTalk Phase I protocol, over which Apple Macintosh workstations and servers communicate

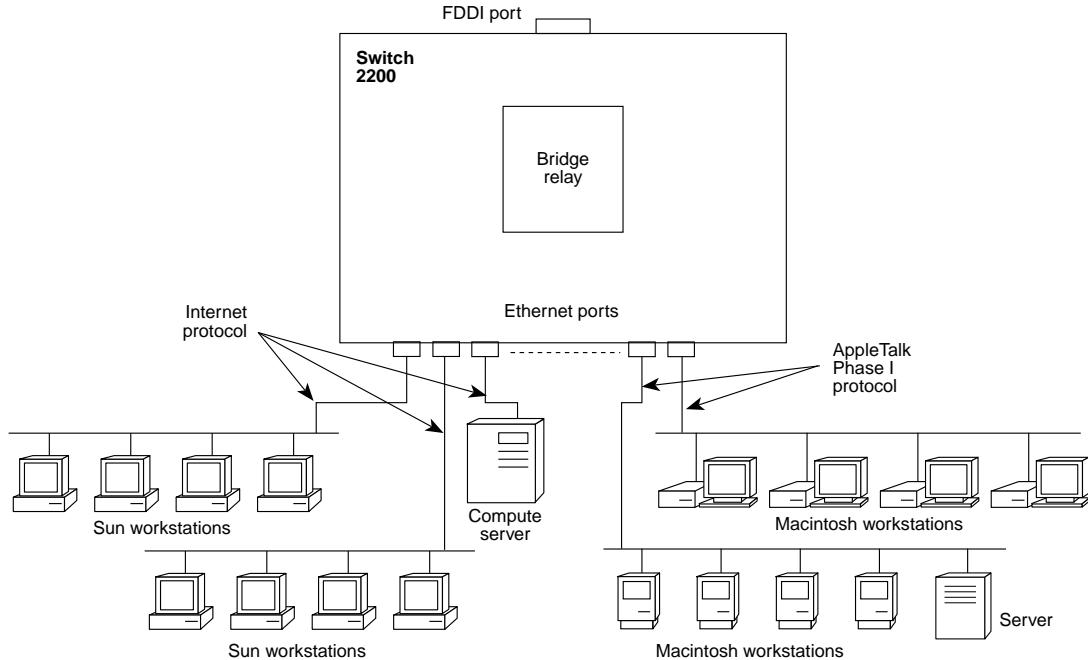


Figure 6-2 A Network with Two Protocols

The Macintosh computers can use Internet Protocol (IP) to communicate with the Sun workstations, but the Sun workstations cannot use AppleTalk protocol to communicate with the Macintoshes.

Without any packet filtering, AppleTalk broadcasts would be bridged to the IP segments. Even some unicast traffic would occasionally be flooded to those segments if the destination station had not been learned. Because the Sun workstations cannot interpret AppleTalk packets, you want to isolate these packets from the IP segments.

To isolate the IP segments, define a packet filter that discards all AppleTalk packets received on the transmit path of ports that have only IP stations connected to them. (*All* ports would need a packet filter if the filter were installed on the receive path.)

The filter definition is:

If **type field = AppleTalk** then discard packet

In the example:

- Operand 1 is the type field, which is a 2-byte value at offset 12 in the Ethernet packet.
- Operand 2 is the type field constant value for AppleTalk.
- Operator is equal.

The filter is installed as shown in Figure 6-3.

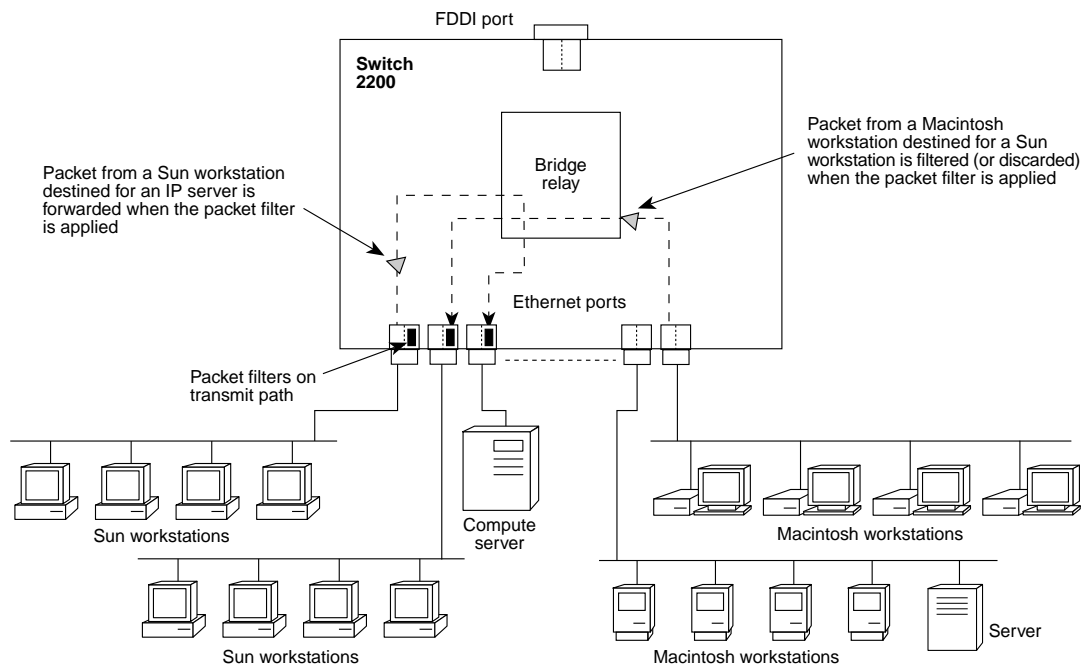


Figure 6-3 Example of AppleTalk Filter

Example 2: Filtering AppleTalk Phase II Packets

If the Macintosh computers use the AppleTalk Phase II protocol instead of the AppleTalk Phase I protocol (as shown in Example 1), then the filter needs to be slightly more complicated.

AppleTalk Phase II uses 802.3 protocol instead of Ethernet as the physical layer protocol. Ethernet and 802.3 packets are distinguished using the 2-byte field at offset 12 in the packet. If that field is greater than 1500, then the packet is an Ethernet packet and the value is interpreted as the type

field. If that field is less than or equal to 1500, then the packet is an 802.3 packet and the value is interpreted as the data length. The filter must first ensure that the packet is an 802.3 packet.

In an AppleTalk Phase II packet, a Subnetwork Access Protocol (SNAP) “header” follows the 802.3 header. See Figure 6-4. The filter needs to verify that the contents of this SNAP field match the AppleTalk packet’s SNAP field.




Figure 6-4 AppleTalk Phase II Packet Fields

The filter definition for filtering AppleTalk Phase II packets is:

```
if (type field <= 1500) AND (SNAP = 0x03080007809b) then discard
packet
```

In this example, several simple expressions are combined to form the complete complex logical expression. The expressions can be separated as follows:

Expression 1: type field <= 1500

- Operand 1 is the type field, which is a 2-byte value at offset 12 in the AppleTalk Phase II packet.
- Operand 2 is the literal constant 1500.
- Operator is greater than or equal.

Expression 2: SNAP = 0x03080007809b

- Operand 1 is the SNAP field, which is a 6-byte value at offset 16 of the AppleTalk Phase II packet.
- Operand 2 is the constant value for the AppleTalk Phase II SNAP field: 0x03080007809b.
- Operator is equal.

- **Expression 3:** Expression 1 result AND Expression 2 result
- Operand 1 is the result of Expression 1.
- Operand 2 is the result of Expression 2.
- Operator is bit-wise AND.

Figure 6-5 illustrates how the simple expressions form a complete packet filter definition.




Figure 6-5 Packet Filter Expressions Evaluated

Using Address Groups and Port Groups in a Packet Filter

The section "About User-defined Packet Filtering" described how you can use packet filters to restrict the flow of packets based solely on the contents of the packet. The Switch 2200 also allows you to set up groups of addresses or ports and then combine these group definitions with the packet filter definitions to control which stations can communicate with each other. For instance, you can define:

- A group of stations that can communicate only with other stations in that group
- A group of stations that have access only to a specific network resource
- A group of stations that have access only to a group of network segments
- A group of network segments whose attached stations can communicate only with each other

What Is an Address Group?

An address group is simply a list of MAC addresses. You can configure up to 32 address groups per Switch 2200. The same address group can be associated with multiple systems.

When an address is added to a group, the address is inserted into the address table on each system that is associated with that group. Each address table entry has a 32-bit *group mask* associated with it. Each bit in the mask specifies *one* of the 32 groups. For example, bit 1 could specify group 1, and bit 2 could specify group 2. When an address is added to a group, the corresponding bit in the mask is set. See Figure 6-6.

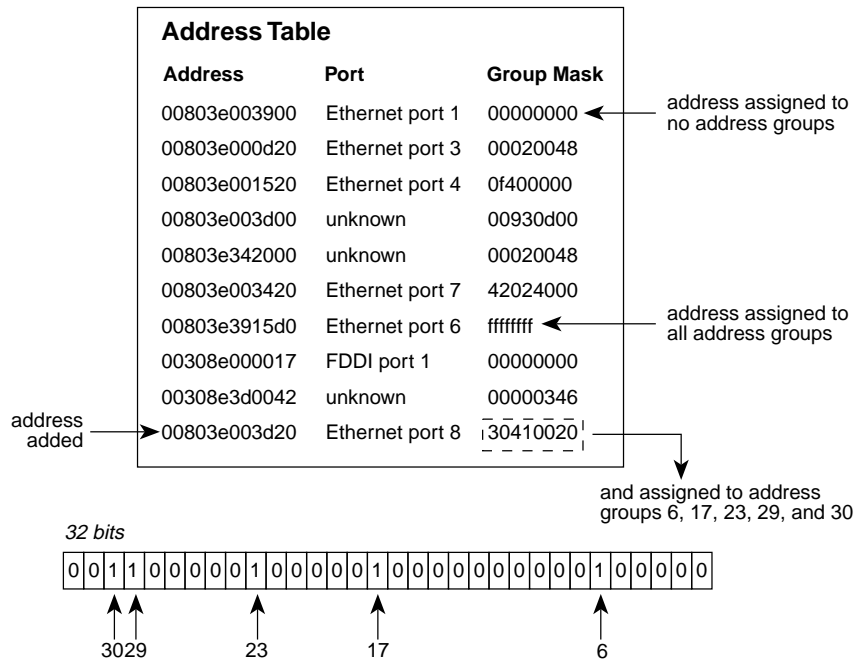


Figure 6-6 Adding an Address to an Address Group



If the address has never been learned or has been aged, the port ID associated with the address is set to “unknown.”



Broadcast and other multicast addresses are assumed to be in all groups.

Referencing Address Groups and Port Groups from a Packet Filter

After you configure address and port groups, you can reference them in a packet filter. The packet filter language defines several operands that relate to address and port groups. These operands are described in Table 6-3.

Table 6-3 Packet Filter Operands for Address Groups and Port Groups

Operand	Description
Source address group mask	The group mask associated with the source address in the packet
Destination address group mask	The group mask associated with the destination address in the packet
Source port group mask	The group mask associated with the bridge port on which this packet was received
Destination port group mask	The group mask associated with the bridge port for which this packet is destined

Example: Using Address Groups in a Packet Filter

The following example shows how a packet filter references the address groups. The process is similar for port groups.

This example shows how you can use packet filtering to restrict which end-stations have access to a specified server. The network is the one shown in Figure 6-8. This network has the following groups and servers:

- Accounting group spread over three segments
- Engineering A group spread over four segments (three of which are direct attach)
- Engineering B group spread over four segments (two of which are direct attach)
- Accounting data server that contains accounting information, such as payroll, revenue data, purchase orders
- Compute servers that are used to compile and link software programs
- Mail Server that is used to store and distribute electronic mail

For purposes of this example, the MAC address for each station in Figure 6-8 is in the form: **00-01-02-03-04-xx**, where "xx" is the station number. For example, Compute server A has the MAC address of **00-01-02-03-04-01**.

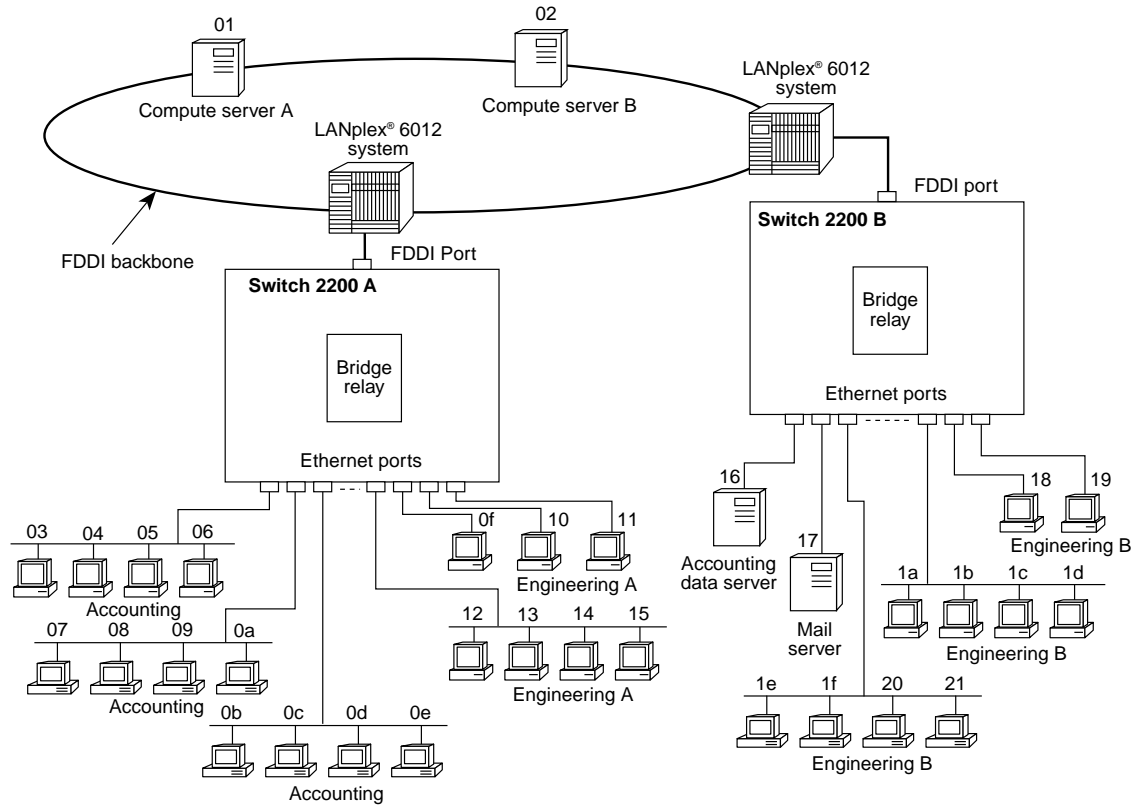


Figure 6-8 Network Needing Filtering to Restrict Server Access

The packet filter is designed to limit network traffic in these ways:

- Users in the Accounting group can communicate with each other and the Accounting Data Server.
- Users in Engineering A group can communicate with each other and Compute Server A.
- Users in Engineering B group can communicate with each other and Compute Server B.
- Users in all three groups can communicate with the Mail server.

To implement packet filtering for this network, take these steps:

- 1 Set up address groups as follows.

Address group 1 — Accounting

00-01-02-03-04-03	00-01-02-03-04-0a
00-01-02-03-04-04	00-01-02-03-04-0b
00-01-02-03-04-05	00-01-02-03-04-0c
00-01-02-03-04-06	00-01-02-03-04-0d
00-01-02-03-04-07	00-01-02-03-04-0e
00-01-02-03-04-08	00-01-02-03-04-16
00-01-02-03-04-09	00-01-02-03-04-17

Address group 2 — Engineering A

00-01-02-03-04-0f	00-01-02-03-04-14
00-01-02-03-04-10	00-01-02-03-04-15
00-01-02-03-04-11	00-01-02-03-04-01
00-01-02-03-04-12	00-01-02-03-04-17
00-01-02-03-04-13	

Address group 3 — Engineering B

00-01-02-03-04-18	00-01-02-03-04-1e
00-01-02-03-04-19	00-01-02-03-04-1f
00-01-02-03-04-1a	00-01-02-03-04-20
00-01-02-03-04-1b	00-01-02-03-04-21
00-01-02-03-04-1c	00-01-02-03-04-02
00-01-02-03-04-1d	00-01-02-03-04-17

Note that the Mail server address (00-01-02-03-04-17) is included in each group. These address groups yield an address table as shown in Figure 6-9 on page 6-14.

The address table on each system contains the same address listing and group masks, but the port numbers are different. The address table in Figure 6-9 on page 6-14 is for Switch 2200 A.

2 After setting up the address groups, you generate the following filter:

if (**source address group mask AND destination address group mask**) = 0 then discard packet

The expressions used in this example filter can be separated as follows:

Expression 1: source address group mask AND destination address group mask

- Operand 1 is the source address group mask.
- Operand 2 is the destination address group mask.
- Operator is bit-wise AND.

Expression 2: Expression 1 result = 0

- Operand 1 is the result of Expression 1.
- Operand 2 is the value 0.
- Operator is equal.

The filter would be installed on the receive paths of the user group ports as shown in Figure 6-10. The packet is examined as soon as it is received by the port and discarded if the destination address of the packet is not in the same address group as the source address of the packet.

Address Table		
Address	Port	Group Mask
00010203040e	Ethernet Port 3	00000001
000102030403	Ethernet Port 1	00000001
000102030421	FDDI Port 1	00000004
00010203040c	Ethernet Port 3	00000001
000102030411	Ethernet Port 8	00000002
000102030419	FDDI Port 1	00000004
00010203041a	unknown	00000004
000102030412	Ethernet Port 5	00000002
00010203041d	FDDI Port 1	00000004
000102030418	FDDI Port 1	00000004
000102030401	FDDI Port 1	00000002
00010203040f	Ethernet Port 6	00000002
000102030409	Ethernet Port 2	00000001
000102030407	Ethernet Port 2	00000001
000102030404	Ethernet Port 1	00000001
00010203040d	Ethernet Port 3	00000001
000102030414	Ethernet Port 5	00000002
000102030417	FDDI Port 1	00000007
00010203041c	FDDI Port 1	00000004
000102030410	Ethernet Port 7	00000002
000102030415	Ethernet Port 5	00000002
00010203041e	unknown	00000004
000102030402	FDDI Port 1	00000004
00010203040a	Ethernet Port 2	00000001
00010203041f	FDDI Port 1	00000004
000102030408	Ethernet Port 2	00000001
00010203041b	FDDI Port 1	00000004
000102030405	Ethernet Port 1	00000001
000102030416	FDDI Port 1	00000001
000102030406	Ethernet Port 1	00000001
000102030420	FDDI Port 1	00000004
00010203040b	unknown	00000001
000102030413	Ethernet Port 5	00000002

This station, on Ethernet port 8, is included in the Engineering A address group.

The Mail server is included in all of the address groups.

Compute server B is included in the Engineering B address group.

The Accounting data server is included in the Accounting address group.

Figure 6-9 Address Table for Restricting Server Access

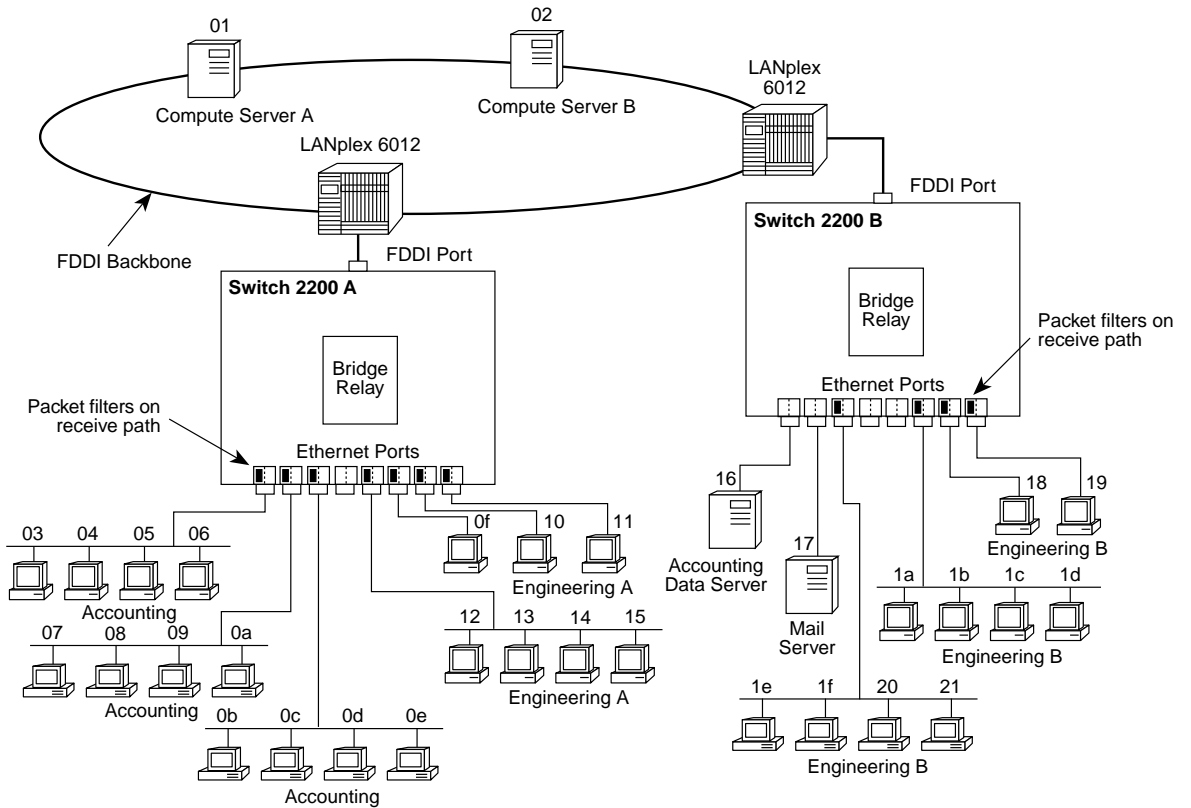


Figure 6-10 Address Group Filtering Example

Globally Administering Packet Filters

You can create packet filters and group definitions locally using the Administration Console. Alternatively, you can define them on an external computer by creating files that contain the necessary information in the specified format and loading them onto the Switch 2200. When you create the definitions externally in this way, multiple Switch 2200s can share the same definition. This capability is especially important for address group distribution because the related stations are often distributed across many different network segments. See the *SuperStack™ II Switch 2200 Administration Console User Guide*.

7

BRIDGING EXTENSIONS

This chapter describes bridging extensions — additional functionality that enhances a SuperStack™ II Switch 2200's bridging performance. These extensions include:

- Multicast packet firewalls
- IP fragmentation
- Reduced packet flooding
- Network security enhancements

Multicast Packet Firewalls

A network error condition that can significantly disrupt attached stations is a *multicast storm*. This term refers to the repeated transmission of a high rate of broadcast (or other multicast) packets onto the network. Several scenarios can result in a multicast storm. These include:

- Faulty protocol implementations
- Undetected network loops
- Faulty network equipment

As a result of these storms, the network and its attached stations are stressed, often causing end-stations to hang or fail. The Switch 2200 system supports a mechanism, called the *multicast packet firewall*, that limits the rate at which multicast packets are forwarded. You can adjust the threshold rate to control the effects of multicast storms on your network.

See Figure 7-1 for an illustration of the threshold mechanism. For information on setting this threshold, see the *SuperStack™ II Switch 2200 Administration Console User Guide*.

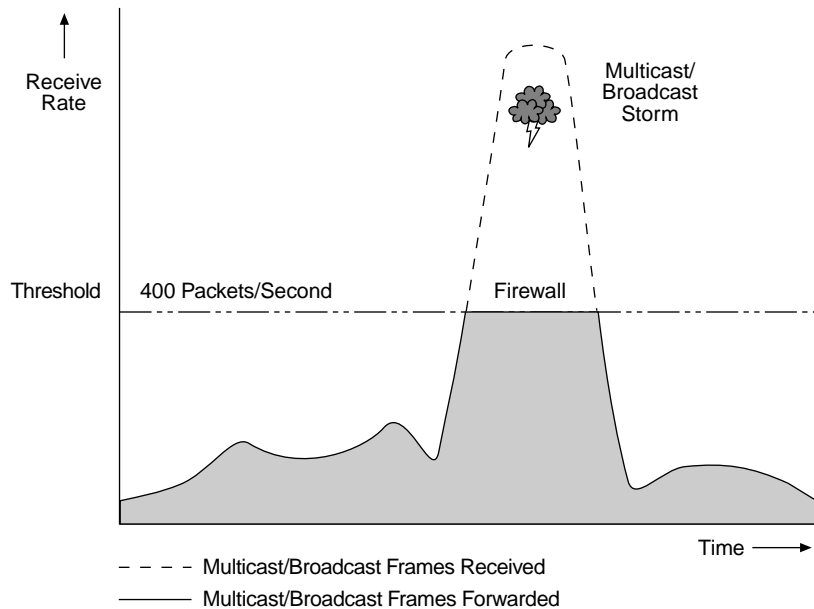


Figure 7-1 Multicast Packet Firewall Threshold Mechanism

IP Fragmentation

The maximum length of the information field in an FDDI packet is 4478 bytes, but the maximum length of the information field in an Ethernet packet is only 1500 bytes. Therefore, any packet sourced from an FDDI station and destined for an Ethernet station must have an information field that does not exceed 1500 bytes to bridge the packet in a conventional fashion.

To overcome this limitation, the Internet Protocol (IP) specifies a procedure, called *IP fragmentation*. This allows a large FDDI packet to be “fragmented” into smaller packets. With IP fragmentation, FDDI and Ethernet stations connected to a Switch 2200 can communicate using IP even if the FDDI stations are transmitting packets that would typically be too large to bridge.

IP fragmentation is specified in RFC 791 (Internet Protocol) and RFC 1122. For information on enabling IP fragmentation, see the *SuperStack™ II Switch 2200 Administration Console User Guide*.

Reduced Packet Flooding

The Switch 2200 has functionality that enhances IEEE 802.1d's traditional timer-based address aging mechanism to reduce packet flooding significantly.

When a station is moved from one bridged segment to another, its address must be learned on the new bridge port and forgotten on the old one. If the address is learned on both ports at the same time, a packet sent to that address may be directed to the old port instead of to the new one. Therefore, it is important that the address be moved in a timely fashion.

Traditional bridges use the address aging process to forget the address at its former location. As a result, you would typically set the bridge aging timer to a relatively short interval to reduce the likelihood that an address is learned on more than one bridge port at a time. Unfortunately, the side effect of shortening the aging interval is that all station addresses are forgotten and re-learned more frequently. This results in increased packet flooding.

The Switch 2200 supports the traditional aging mechanism, but it also contains logic that monitors the source address in every packet, ensuring that the port associated with that address has not changed. If the Switch 2200 detects that a station has moved, it immediately re-assigns the associated port. This enhancement allows you to lengthen the Switch 2200 address aging interval, knowing that the Switch 2200 rapidly detects most station moves independent of the aging process. This results in reduced packet flooding and improved network performance.


For information on setting the address aging timer, see the *SuperStack™ II Switch 2200 Administration Console User Guide*.

Enhanced Network Security

In addition to using packet filters to improve network security (as described in Chapter 6), the Switch 2200 allows you to use statically configured addresses as a form of network security.

From the Administration Console or an SNMP manager, you can manually assign an address to an Ethernet port. This address is then permanently tied to the specified port unless you manually remove it. This implies that the address is never aged and can never be learned on a different Ethernet port. If a packet with a statically configured source address is received on a port that differs from the address's assigned port, the packet is discarded and a management event is generated. The event can be used to expose the imposter station.


You can also convert dynamic addresses to static addresses. It is often much more convenient to let the bridge first learn all of the addresses on your network and then convert these learned, or dynamic, addresses to static addresses to improve your network security.




FDDI TECHNOLOGY

Chapter 8 FDDI Overview and Implementation

Chapter 9 FDDI Networks





FDDI OVERVIEW AND IMPLEMENTATION

This chapter provides you with general FDDI concepts, terms, and background information, and describes how FDDI is implemented in the Switch 2200.

About FDDI

Fiber Distributed Data Interface (FDDI) is a standards-based solution that provides fast and reliable data transfer on a local area network (LAN). FDDI's sophisticated technology, which supports data transfer of 100 million bits per second (100 Mbps), was developed by the American National Standards Institute (ANSI). FDDI meets the demands of today's powerful and data-intensive computing environments by providing increased throughput, greater network size, improved reliability, and superior fault tolerance. Here are some facts about FDDI:

- FDDI uses optical fiber as its transmission medium, providing security, low signal loss, and high bandwidth data communication.
- FDDI can support simultaneous connection of over 500 nodes on a ring, with up to two kilometers between adjacent nodes, and up to 200 kilometers of total fiber length.
- FDDI uses a token-passing protocol for access to the network.
- FDDI uses a dual-ring approach — a combination of two independent counter-rotating rings, each running at a data rate of 100 Mbps.
- FDDI is the first LAN technology to provide an embedded network management capability.

FDDI is divided into four major standards:

- **Physical Medium Dependent (PMD)** — PMD specifies the characteristics of the fiber optic medium, the connectors that attach stations to the fiber optic medium, the transmission wavelength, the power requirements for transmitters, and the methods for optically bypassing inactive stations.
- **Physical (PHY)** — PHY specifies data encoding and decoding, clock speed and clocking scheme, data framing, and the control symbols used in the network.
- **Media Access Control (MAC)** — MAC specifies access to the medium, token passing, addressing, data checking, frame generation and reception, error detection and recovery, and the bandwidth allocation among the stations.
- **Station Management (SMT)** — SMT specifies the FDDI station and ring configurations, initialization and maintenance of station-to-station connections, and the control required for the proper operation of stations in an FDDI ring.

These four standards are described in relation to the Open Systems Interconnect (OSI) Reference Model. This model was established by the International Standards Organization (ISO) to standardize digital data communications. Each FDDI station is made up of logical entities that conform to the four standards. These entities represent the active services or management elements within OSI.

Figure 8-1 illustrates the relationship of FDDI entities to the OSI Reference Model. Network attachments communicate with each other using predetermined protocols. The model segments these communication protocols into seven layers, defined so that each layer only requires services from the layer below it.

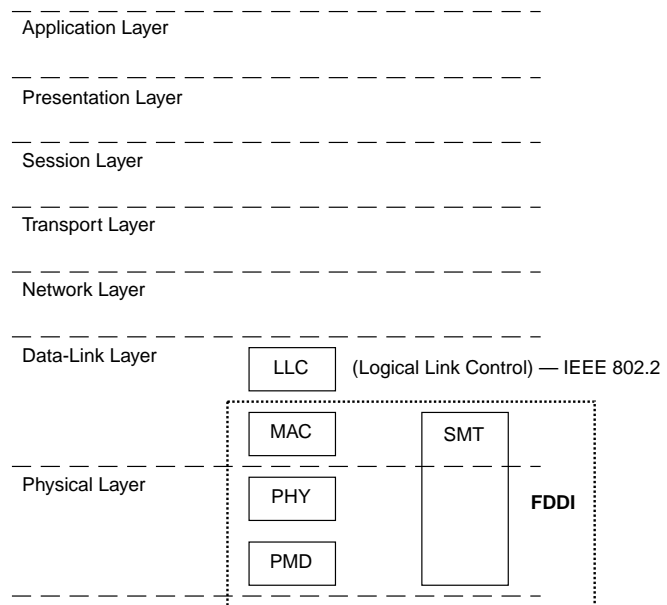


Figure 8-1 FDDI Relationship to OSI Reference Model

Ports

As parts of the Physical Layer, the PHY and PMD entities work together to support each link between FDDI stations. These entities provide the protocols that support the transmission and reception of signals between stations, as well as the optical fiber hardware components that link FDDI stations together. Within an FDDI station, the PHY and PMD entities make up a *port*. Together, they create a PHY/PMD pair that connects to the fiber media and that provides one end of a physical connection with another station.

Ports are located at both ends of a physical connection and determine the characteristics of that physical connection. The protocols that are executed at each port determine whether the connection is accepted or rejected. A connection is accepted if at least one station's policy is to allow such a connection. A connection is rejected if both stations have a policy that disallows the connection.

Each port belongs to one of four types: A, B, M, and S.

- **A port** — This port connects to the primary ring on the incoming fiber and the secondary ring on the outgoing fiber. A properly formed FDDI dual ring is composed of a set of stations with the A port of one station connected to the B port of the neighboring station.
- **B port** — This port connects to the incoming fiber of the secondary ring and the outgoing fiber of the primary ring.
- **M port** — This port, also referred to as Master port, is used by a concentrator station to provide connections within a concentrator tree.
- **S port** — This port, also referred to as Slave port, is used by a single attachment station to provide attachment to an M port within a concentrator tree.

MACs

The Media Access Control (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 to their destination by scheduling and performing all data transfers.

MAC Services

Some of the services that the MAC performs include frame repetition and reception, frame removal, frame validity criteria checking, token capture, token rotation, ring initialization, and the beacon process. MAC services are provided by all conforming stations attached to the FDDI network.

MAC Operation

The MAC controls access to the physical medium by passing a token around the ring. When the token is received by a station, the station may transmit a frame or a sequence of frames. When a station wants to transmit, it removes the token from the ring and transmits the queued frames. After transmission, the station issues a new token, which is used by the downstream station.

Stations that are not transmitting only repeat the incoming symbol stream. When repeating, the station determines whether the information was or was not destined for it by comparing the destination address to its own address. If a match occurs, the MAC processes subsequent received symbols or sends them to the Logical Link Control (LLC) in the data-link layer for translation.

Paths

FDDI's dual, counter-rotating ring is made up of a primary and 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 or station failure occurs, the ring "wraps" around the location of the failure, creating a single logical ring.

Paths represent the segments of a logical ring that pass through a station. An FDDI station can contain three paths: the primary, the secondary, and the local. They are defined as follows:

- **Primary path** — The segment or segments of the primary ring that pass through a station. Conditions may exist in parts of the network that may cause the path to be in a different ring. The primary path must be present in all nodes on the network.
- **Secondary path** — The segment or segments of the secondary ring that pass through a station. Conditions may exist in parts of the network that may cause the path to be in a different ring.
- **Local path** — Represents the segment or segments of the rings other than the primary ring and secondary ring that pass through the station.

Nodes and Attachments

An FDDI network is made up of stations and concentrators that contain active services or management elements that conform to the ANSI FDDI standards. These stations and concentrators are connected to optical fiber medium and are attached in the prescribed manner set forth in the FDDI standards to allow reliable data transmission. These connections are made through FDDI ports and are managed by FDDI MACs.

Nodes An FDDI network is made up of logically connected *nodes*. This generic term is used to refer to any active *station* or *concentrator* in an FDDI network.

- A *station* is any addressable node on an FDDI network that can transmit, repeat, and receive information. A station contains only one SMT, and *at least one* MAC, one PHY, and one PMD.
- A *concentrator* is an FDDI station with additional PHY/PMD entities, beyond those required for its own connection to an FDDI network. These additional PHY/PMD entities (M ports) are for connecting other FDDI stations, including other concentrators, in a tree topology.

Attachments Attachments refer to how a node, station, or concentrator is connected to an FDDI network. They are classified as *single attachment* and *dual attachment*. Concentrators can be classified as *null attachment* when the A and B ports are either not present or not used.

- **Single attachment** — A station or concentrator that has only one physical connection to an FDDI network. It cannot accommodate a dual (counter-rotating) ring. A single attachment station or concentrator has an S port that attaches to an M port within a concentrator tree.
- **Dual attachment** — Any station or concentrator that has two physical connections to an FDDI network. This type of attachment can accommodate a dual (counter-rotating) ring. A dual attachment station has one A-B port pair; a dual attachment concentrator has an A-B port pair and at least one M port.
- **Null attachment** — Concentrators that have one or more M ports, but do not contain A, B, or S ports.

Node Types Six station and concentrator types are used to describe station configurations and topologies. Table 8-1 lists these node types and their abbreviations. An example of how these six node types may connect to an FDDI dual ring are shown in Figure 8-2.

Table 8-1 Node Types and Abbreviations

Node Type	Abbreviation
Single MAC-Dual Attachment Station	SM-DAS
Dual MAC-Dual Attachment Station	DM-DAS
Single Attachment Station	SAS
Dual Attachment Concentrator	DAC
Single Attachment Concentrator	SAC
Null Attachment Concentrator	NAC



The Switch 2200 only supports SM-DAS and SAS topologies.

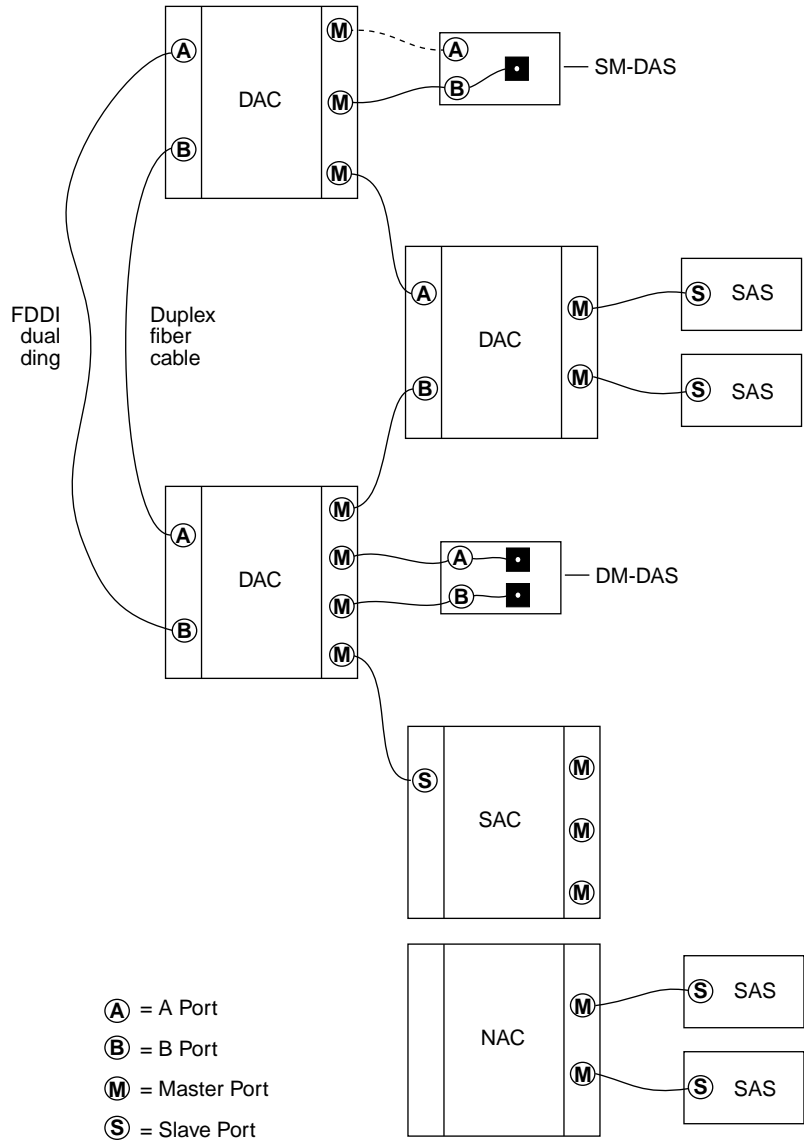


Figure 8-2 Examples of Possible FDDI Configurations

Station Management

Each FDDI station has one Station Management (SMT) entity to provide connection management, ring management, and operational management to the FDDI network. SMT specifies a set of services and signaling mechanisms dedicated to FDDI network management. It manages those services of each station on the FDDI network that are specific to the Physical Layer and the MAC portion of the Data Link Layer.

The goal of SMT is to completely define shared medium-management services to guarantee the interoperability of FDDI network equipment from multiple vendors.

SMT Operation

The operation of SMT can be divided into three broad categories: Physical Connection Management, Configuration Management, and Ring Management.

- **Physical Connection Management (PCM)** — PCM establishes and maintains point-to-point physical links between neighboring ports. It provides all the signaling necessary to initialize connections, withhold marginal connections, and support maintenance.
- **Configuration Management (CFM)** — CFM interconnects PHYs and MACs on paths to achieve proper station configuration and network topology.
- **Ring Management (RMT)** — RMT manages a MAC's operation in an FDDI ring. RMT detects stations that are "stuck" in the beacon process and initiates the trace function. RMT locates duplicate addresses that may prevent the ring from operating.

The FDDI MIB

The FDDI Management Information Base (MIB) defines the collection of information available to network management about an FDDI station. The MIB uses an object-oriented approach similar to that used in OSI management standards. FDDI-managed objects include SMT (that is, the SMT of the station), MACs, paths, and ports. Each of these objects has a collection of attributes such as statistics, error counters, configuration information, event notifications, and actions.

You can access a station's MIB locally through a local management interface or remotely through a management protocol such as Parameter Management Frame (PMF) or Simple Network Management Protocol (SNMP). The SMT standard specifies the meaning and encoding of each MIB attribute.

Frame-based Protocols

SMT provides a number of frame-based services that are used by higher-level management functions to manage stations on the network and to gather information about them. Frame-based protocols perform these functions. The purposes of these protocols are to:

- Gather network statistics
- Detect, isolate, and resolve faults in the network
- Tune FDDI configuration and operational parameters to meet application and connectivity requirements

The key frame-based protocols for SMT are:

- **Neighbor Notification** — This protocol allows SMT to learn the addresses of the logical neighbors of each MAC in a station. This information is useful in detecting and isolating network faults.
- **Parameter Management** — This protocol performs the remote management of station attributes. It operates on all SMT MIB attributes, attribute groups, and actions.
- **Status Reporting** — This protocol allows a station to notify network managers of the occurrence of events, such as station configuration changes and network errors.
- **Status Polling** — This protocol provides a mechanism to obtain station status remotely through a request/response protocol.
- **Echo** — This protocol performs loopback testing on the FDDI dual ring.
- **Synchronous Bandwidth Allocation** — This protocol allocates synchronous bandwidth and monitors both synchronous and total bandwidth.

FDDI and the Switch 2200 System


Your Switch 2200 system brings you the power of FDDI by combining Ethernet switching and Ethernet-to-FDDI transparent bridging into one system. This combination dramatically enhances LAN performance and increases the capacity of your existing Ethernet network. With all this power, you can accommodate both the significant demands of client/server computing, and the addition of high-performance workstations, applications, and super servers. For example, when you place your super servers on FDDI and your clients on switched Ethernet ports, you immediately get the speed and capacity of FDDI without having to upgrade all clients to FDDI. You can install your system into many possible FDDI configurations. For FDDI configuration examples, see the *SuperStack™ II Switch 2200 Getting Started* guide.



3Com strongly recommends that equipment that can be powered on and off, such as workstations, be connected only through concentrators. Intermediate systems that are seldom powered off, such as bridges and routers, should be connected to the FDDI dual ring only if they are equipped with an optical bypass switch. These precautions protect the integrity of the dual ring.



For additional information on FDDI, see Chapter 9: FDDI Networks, which discusses physical and logical topologies, FDDI connection rules, and dual homing.



FDDI NETWORKS

This chapter provides general information about FDDI networks and describes the difference between physical and logical topologies. This chapter also covers FDDI connection rules and dual homing.

About FDDI Networks

FDDI networks have many important differences from other types of LANs. FDDI networks can provide a network backbone between buildings on a campus or within a multilevel high-rise building. A major advantage of an FDDI network is that it can meet the networking needs of today's high-performance workstations that produce large quantities of data. FDDI networks also offer the speed, distance, and capacity required for the powerful workstations, applications, and super servers of the '90s. And FDDI networks are capable of handling the significant demands of client/server computing.

FDDI Network Topologies

The term *network topology* refers to the ways that stations are interconnected within a network. An FDDI network topology may be viewed at two distinct levels: physical and logical.

- **Physical topology** — A network's physical topology is defined by the arrangement and interconnection of its nodes. The FDDI physical topology is a *ring of trees*. See Figure 9-1.




Figure 9-1 Physical Topology Example

- **Logical topology** — A network's logical topology is defined by the paths through which tokens and data flow in the network. The FDDI logical topology is a *dual ring*. See Figure 9-2.




Figure 9-2 Logical Topology Example

Physical Topology: The Ring of Trees

The FDDI trunk ring consists of DASs and DACs. The DACs on the ring allow trees to be attached. The trees consist of branches of SASs and DASs that are star-wired off the concentrators. There are several advantages to creating this kind of network. In addition to being highly reliable, the ring of trees provides a single, fault-tolerant ring, offers fault isolation, and allows centralized management. See Figure 9-3.

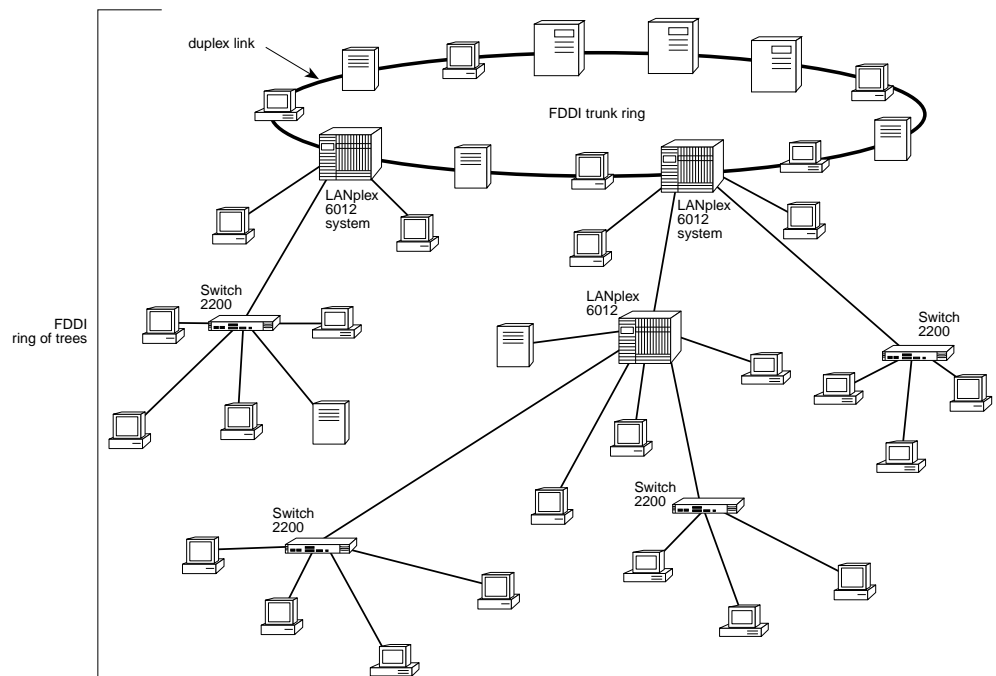


Figure 9-3 Ring of Trees

All physical connections in an FDDI topology are *duplex links* (a pair of insulated fiber optic conductors). Both the FDDI trunk ring and the ring of trees created through concentrators are made up of duplex links. The nodes in an FDDI network must be interconnected to form at *most* one trunk ring.

If a topology is legal, when physical connections and nodes fail or are removed from the network, one or more legal FDDI topologies are formed. This means that subsets of legal topologies are also legal. Some examples of legal FDDI topologies include the dual ring with trees, the dual ring without trees, and the single tree. For information on legal topologies, see the section "FDDI Connection Rules" later in this chapter.

Logical Topology: The Dual Ring

A legal FDDI topology consists of at most two separate logical rings: the primary ring and the secondary ring. These logical rings are formed from the physical links that make up the Physical Layer connections. For example, a set of DASs connected into a closed loop form an FDDI dual ring. Each ring is a logical ring, that is, a separate data path with its own token.

Functionally, one of the major characteristics of the FDDI network is its dual ring, which provides a high degree of reliability to a LAN. When an FDDI network is in normal operation, only the primary ring is used to transmit and receive data. The secondary ring may also be used to carry data, but it is typically used as a backup in case there is a connectivity problem in the primary ring or in one of the nodes on the ring.

When a single fault takes place on an FDDI dual ring, recovery can be made by joining the two rings between the two nodes adjacent to the fault. This creates a single logical ring resulting in a wrapped configuration. A wrapped ring is a legal FDDI topology. In the same way, when many faults take place, several disjointed logical rings are created, producing multiple FDDI topologies.

FDDI Connection Rules

SMT follows specific connection rules to ensure that only desired physical connection types are included in the network topology. A connection's type is determined by the types of the ports at either end of the connection. There are three categories of connection types:

- **valid** — Always accepted
- **illegal** — Always rejected
- **undesired** — Either accepted or rejected as determined by connection policies established by the network manager

SMT notifies network management when undesired connection types are attempted, regardless of whether the connection is accepted or rejected. The FDDI SMT standard cites detailed connection rules for a specific port ("this Port") to other ports, which are shown in Table 9-1.

Table 9-1 Port Connection Rules

Port Connection	Connection Rules
A to A	Undesirable peer connection that creates twisted primary and secondary rings; notify SMT
A to B	Normal dual ring peer connection
A to S	Undesirable peer connection that creates a wrapped ring; notify SMT
A to M	Tree connection with possible redundancy. Node will not go to THRU state in Configuration Management (CFM). My B port (the port you are connected to) will have precedence (with defaults) for connection to an M port in a single MAC node.
B to A	Normal dual ring peer connection
B to B	Undesirable peer connection that creates twisted primary and secondary rings; notify SMT
B to S	Undesirable peer connection that creates a wrapped ring; notify SMT
B to M	Tree connection with possible redundancy. Node will not go to THRU state in CFM. My B port will have precedence (with defaults) for connection to an M port in a single MAC node.
S to A	Undesirable peer connection that creates a wrapped ring; notify SMT
S to B	Undesirable peer connection that creates a wrapped ring; notify SMT
S to S	Connection that creates a single ring of two slave stations
S to M	Normal tree connection
M to A	Tree connection that provides possible redundancy
M to B	Tree connection that provides possible redundancy
M to S	Normal tree connection
M to M	Illegal connection that creates a tree of rings topology

Table 9-2 provides a connection rule matrix summarizing the validity of most types of connections.

Table 9-2 Connection Rule Matrix

		Other Port			
		A	B	S	M
This Port	A	V, U	V	V, U	V
	B	V	V, U	V, U	V
	S	V, U	V, U	V	V
	M	V	V	V	I, U

* V — A valid connection
 * I — An illegal connection
 * U — An undesirable connection (with notification to SMT required)

Dual Homing

When the operation of a dual attachment node is critical to your network, a configuration called dual homing can provide added reliability. When using dual homing, a network administrator can determine a station's operation by setting the appropriate configuration policy. The dual-homed station can be configured in one of two ways: 1) with both links active or 2) with one link active and one connection withheld as a backup, becoming active only if the primary link fails. See Figure 9-4.





Figure 9-4 Dual Homing


For additional information about dual homing, see the *SuperStack™ II Switch 2200 Getting Started* guide.




APPENDIXES

Appendix A SNMP MIB Support

Appendix B Technical Support





SNMP MIB SUPPORT

This appendix lists the SNMP MIBs supported by the SuperStack™ II Switch 2200 system software and describes the supported SNMP compilers.

SNMP MIBs

SNMP MIB files are shipped with the system software as ASN.1 files. The currently supported version of each MIB is listed in this section. All applicable MIB attributes are supported unless otherwise specified.



MIB version changes and attribute additions and deletions may occur from release to release. They are documented in the release notes.

Copies of ASN.1 files are provided for each of the supported compilers described in this appendix.

- **bridge.mib** — Bridge MIB, RFC 1493

The following Bridge MIB attributes are not supported:

dot1dBase Group

- dot1dBasePortDelayExceedDiscards

dot1dSr Group

dot1dTp Group

- dot1dFdbTable

dot1dStatic Group

- **ethernet.mib** — Ethernet MIB, RFC 1398

The following Ethernet MIB attributes are not supported:

dot3StatsTable

- dot3StatsMultipleCollisionFrames
- dot3StatsSQETestErrors
- dot3StatsDeferredTransmissions

dot3CollTable

- **fddiSmt7.mib** — FDDI SMT 7.3 MIB, RFC 1512
- **Ip.mib** — LANplex Systems MIB, version 1.2.1

The following LANplex Systems MIB trap is not supported:

- IpsSystemFanFailure

- **IpOpFddi.mib** — LANplex Optional FDDI MIB, version 1.2.1, based on SMT 7.3

The following LANplex Optional FDDI MIB attributes are not supported:

IpOptMAC Group

- IpOptMACPriTable

IpOptPATH Group

- IpOptPATHSbaTable
-

- **mib2.mib** — MIB-II, RFC 1213

The following MIB-II attributes are not supported:

interfaces Group

- ifLastChange

egp Group

- **srbridge.mib** — Source Routing MIB RFC1525

The following generic SNMP traps are not supported:

-
- warmStart
 - linkDown
 - linkUp
 - egpNeighborLoss
-

SNMP MIB Compilers

ASN.1 MIB files are provided for each of the MIB compilers listed in this section. Any warnings or exceptions related to a compiler are listed with it.

- SMIC (version 1.0.9)
- MOSY (version 7.1)


For the MIB file *IpOPFddi.mib*, the MOSY compiler reports warnings for counter names that do not end in "s". This report has no effect on the output produced by the MOSY compiler.

- HP Openview (version 3.1)
- mib2schema (with SunNet Manager™ version 2.0)

The MIB file *fddiSmt7.mib* produces the following warning messages when compiled using mib2schema:

```
Translating....
Warning: The following INDEX entries in
fddimibMACCountersTable not resolved:
    fddimibMACSMTIndex
    fddimibMACIndex
Translation Complete.
Schema file in "fddiSmt7.mib.schema"
Oid file in "fddiSmt7.mib.oid"
```

However, these warning messages have no effect on the ability of SNM to use the schema file generated with versions SNM 2.0 or later.



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. Dial 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.

3ComFacts 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. International access numbers are:

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

OPERATION GLOSSARY

- A port** Each DAS contains two ports, one designated A and one designated B. The A port is connected to the primary ring on the incoming fiber and the secondary ring on the outgoing fiber. A properly formed trunk ring is composed of a set of stations with the A port of one station connected to the B port of the neighboring station. See also *B port*.
- ANSI** American National Standards Institute. ANSI, the primary group that defines standards in the United States, developed the Fiber Distributed Data Interface (FDDI) standard.
- AppleTalk®** Apple Computer Corporation's networking specifications for the physical layer (LocalTalk, EtherTalk, and TokenTalk), network and transport functions (Datagram Delivery Protocol and AppleTalk Session Protocol), addressing (Name Binding Protocol), file sharing (AppleShare), and remote access (AppleTalk Remote Access).
- application layer** The uppermost layer of the OSI model and the only layer that users can directly communicate with. Users interact with the applications layer through such services as e-mail and file transfer. See also *OSI*.
- ATM** Asynchronous Transfer Mode. A transfer method used by Broadband ISDN. ATM carries voice, video, and data at speeds up to 2.2 Gbps and can integrate geographically distant disparate networks. Also called cell relay.
- B port** Each DAS contains two ports, one designated A and one designated B. The B port is intended to be connected to the incoming fiber of the secondary ring and the outgoing fiber of the primary ring. A properly formed trunk ring is composed of a set of stations with the A port of one station connected to the B port of the neighboring station. See also *A port*.
- bridge** Equipment that connects different LANs, allowing communication between devices on separate LANs. Bridges are protocol independent, but hardware specific, with communication limited to the data link layer and physical layer

of the ISO reference model. Bridges connect LANs with different hardware and different protocols. An example is a device that connects an Ethernet network to an FDDI network. This bridge allows the two networks to send signals to each other. The Switch 2200 can operate as a translation/transparent 802.1d bridge. See also *Spanning Tree Protocol*.

broadcast packet A single packet that is sent to all stations in a network. See also *multicast packet*.

broadcast/multicast storm The network congestion that results when many stations, responding to a transmission by one station, transmit a large number of frames. This condition can overstress a network and cause end-stations to hang or fail.

cell relay See *ATM*.

client A single-user computer that requests application or network services from a server.

client-server A distributed system model of computing that brings computing power to the desktop, where users ("clients") access resources from servers.

community string A character string included in each SNMP protocol message sent between external management applications, such as between Transcend® Enterprise Manager and the Switch 2200 system.

DAS Dual attachment station. A station directly attached to FDDI's dual token rings. A DAS has four fiber attachments, one receive and one transmit fiber for each ring. Rather than an individual user workstation, a DAS is most likely to be the device controlling LAN operation, such as an FDDI concentrator, bridge, router, server, minicomputer, or mainframe. A DAS can be either single-MAC or dual-MAC and contains one A-B port pair.

data-link layer The second layer of the OSI model, which contains the Media Access Control (MAC) sublayer and the Logical Link Control (LLC) sublayer. The data-link layer defines how data is divided into packets and transmitted within a network. See also *ISO, OSI*.

dual homing A method of cabling concentrators and stations that allows an alternate path to the FDDI network. Dual homing creates a more stable ring of concentrators.

- Ethernet** A CSMA/CD, 10Mbps, local area data network, developed by Digital Equipment Corp., Intel, and Xerox Corporation. It is one of the most popular baseband LANs in use.
- Express switching** A positive filtering algorithm that automatically learns the addresses of stations attached to each Ethernet port and forwards only packets specifically destined for learned stations. This operational mode of the Switch 2200 eliminates superfluous traffic created by the flooding that results from IEEE 802.1d address learning and aging.
- FDDI** Fiber Distributed Data Interface. A high-performance, fiber optic token ring LAN that operates at 100Mbps over distances of up to 200 kilometers with up to 1000 connected stations.
- FDDI dual ring** The pair of counter-rotating, logical rings (primary and secondary) common to the FDDI network. This architecture provides a high degree of reliability. In normal operation, only the primary ring carries data. The second or backup ring is used for automatic recovery in case of failure. If a network fault occurs, only the stations on either side of the fault are affected. They detect the fault and automatically bypass it to maintain continuous transmission of data.
- FDDI paths** The segments of an FDDI ring that pass through a station. Every FDDI station must contain a primary path. The primary path represents, to the best of the station's knowledge, the segments of the primary ring that pass through the station. In addition, a station may optionally contain a secondary path representing the segments of the secondary ring that pass through the station. A station may contain additional paths representing segments of rings other than the primary and secondary. Such paths are called local paths.
- FDDI standard** A standard by the X3T9.5 Committee of the American National Standards Institute (ANSI) that addresses the need for more speed and reliability than is currently available in other standard LANs. Its recent completion is a major factor contributing to the expected acceptance and widespread use of optical fiber as a LAN transmission medium. The standard has four parts. See also *PHY standard*, *PMD standard*, and *SMT*.
- Flash EPROM** Erasable Programmable Read-Only Memory.

- frame buffer memory** In data communications, a storage medium used for holding one or more blocks of data during transfer of that data from one device to another.
- gateway** A hardware and software device, operating at the fourth through seventh levels of the OSI model, that connects two dissimilar systems.
- hostname** A meaningful, easy-to-remember name or title assigned to a machine on the Internet that is associated with the IP address. See also *IP address*.
- IEEE 802.1d** A bridging standard specifying that a transparent bridge must learn source addresses, age addresses, store and forward packets, and participate in the Spanning Tree Protocol.
- in-band management** Network management performed using the same network normally used for data transmission. See also *out-of-band management*.
- interoperability** The ability of computer equipment from one vendor to communicate and exchange information with dissimilar equipment from other vendors.
- IP address** Internet Protocol address. A unique identifier for a machine attached to a network that's made up of two or more interconnected local area or wide area networks.
- IP fragmentation** The process of breaking up larger IP frames on one network to a size compatible with the network to which they will be forwarded.
- ISO** International Standards Organization. The ISO is a multinational organization that sets computer, communication, and other standards. The ISO defined the OSI seven-layer reference model for computer communications.
- LAN** Local Area Network. A data communications network spanning a limited geographical area, such as a single building or campus. It provides communication between computers and peripherals. LANs are distinguished by their small geographical size, high data rate, and low error rate.
- LLC** Logical Link Control. The upper sublayer of the data-link layer of the OSI seven-layer reference model. The LLC handles error control, flow control, and frames transmission between stations. The IEEE 802.2 standard is the most widely implemented LLC protocol.

- local management** Management of a station by software running on the station.
- MAC** Media Access Control. A station resource that specifies the lower sublayer of the data-link layer for FDDI. It presents the specifications and services provided for conforming FDDI attachment devices. MAC specifies access to the medium, including addressing, data checking, and data framing.
- MIB** Management Information Base. Stores a device's managed characteristics and parameters. MIBs are used by Simple Network Management Protocol (SNMP) and Common Management Information Protocol (CMIP) to contain attributes of their managed systems. The Switch 2200 contains its own internal MIB.
- multicast packet** A single packet that is copied to a subset of addresses in a network. See *broadcast packet*.
- multicast packet firewall** A mechanism in the Switch 2200 that limits the rate at which multicast packets are forwarded through the system. This threshold is configurable.
- nonvolatile memory** Computer memory that is preserved when power is lost. Also called NVRAM.
- operating system** A program that manages and provides access to system resources.
- OSI** Open Systems Interconnect. A reference model, developed by the ISO, that divides computer communications into seven layers: physical, data-link, network, transport, session, presentation, and application. See also *ISO*.
- out-of-band management** Network management accomplished through a network or connection other than the one normally used for data transmission. See also *in-band management*.
- packet filtering, user-defined** A second layer of filtering on top of the standard filtering provided by a traditional transparent bridge. This filtering can improve network performance, provide additional security, and logically segment a network to support virtual workgroups.
- PHY standard** Physical Layer standard. An American National Standard (ANSI X3) that specifies the data-encoding mechanism and the clock recovery and data framing parameters.

- PMD standard** Physical Layer Medium Dependent standard. An American National Standard (ANSI X3) that specifies the lower sublayer of the physical layer for FDDI, including the power levels and characteristics of the optical transmitter and receiver; interface optical signal requirements including jitter; the connector receptacle footprint; the requirements of conforming FDDI optical fiber cable plants; and the permissible bit error rates.
- primary ring** One of two counter-rotating, fiber optic rings that serve as the root of an FDDI network. The primary ring normally enters each station on the trunk ring through the A port and exits through the B port. See also *secondary ring*.
- protocol** A set of rules for communicating between computers. The rules dictate format, timing, sequencing, and error control.
- proxy agent** Acts as a management gateway, converting requests and event reports from one protocol and object format to another protocol and object format.
- RS-232 serial port** The port on the system accepting a DB-9 or modem connector. It changes the parallel arrangement of data within computers to the serial (one after the other) form used on data transmissions links. This port can be used for dedicated local management access.
- SAS** Single attachment station. A station that offers one attachment to the FDDI network. A SAS has an S port to be attached to an M port within a concentrator tree.
- secondary ring** One of two counter-rotating, fiber optic rings that serve as the root of an FDDI network. The secondary ring normally enters each station on the trunk ring through the B port and exits through the A port. See also *primary ring*.
- server** A computer that provides clients with application and network services. Servers are shared by multiple users.
- SMT** Station Management. A component of the FDDI standard that specifies the control required for proper operation of a station in an FDDI ring.
- SNMP** Simple Network Management Protocol. A protocol originally designed to be used in managing TCP/IP internets. SNMP is presently implemented on a wide variety of computers and networking equipment and may be used to manage many aspects of network and end-station operation. See also *protocol*.

- Spanning Tree Protocol** An algorithm that detects loops in a network and logically blocks the redundant paths, ensuring that only one route exists between any two LANs. This protocol is used in an IEEE 802.1d bridged network.
- topology** The physical or logical placement of stations on a network in relation to one another, such as ring, mesh, star, or bus.
- transparent bridge** A bridge that allows two or more LANs to be interconnected and to communicate as if they were one LAN. The bridge listens promiscuously to packets on the attached LANs and forwards packets from one LAN to another. See also *bridge*.
- UNIX** A computer operating system, developed by AT&T, that is capable of multitasking.

INDEX

Numerics

3Com Bulletin Board Service (3ComBBS) B-1
3Com sales offices B-4
3ComFacts B-3
3ComForum B-2

A

address
 aging 5-3, 7-3
 destination 5-4
 learning 5-2
 network security 7-4
 source 5-3
 table 6-8
address groups
 defined 6-8
 example 6-10
 mask 6-8
aging addresses 5-3, 7-3
AppleTalk
 Phase I 6-3
 Phase II 6-5
attachments 8-6
authenticationFailure 3-5

B

blocking port state 5-16
bridge
 aging timer 7-3
 designated 5-8
 least cost path 5-9
 root 5-8
Bridge Protocol Data Unit (BPDU) 5-18
bridging
 configuration messages 5-8
 extensions 7-1 to 7-4
 IEEE 802.1d compliant 5-1
 SNMP traps, and 3-5
 Spanning Tree 5-6

standards 5-1
 transparent 5-1 to 5-18
bulletin board service B-1

C

cabling 4-1 to 4-3
CBPDU
 best 5-11
 information in 5-10
coldStart 3-5
CompuServe B-2
Configuration Bridge Protocol Data Unit. *See* CBPDU
conventions
 notice icons, About This Guide 2
 text, About This Guide 2

D

designated bridge 5-8
designated port 5-8, 5-9
disabled port state 5-17
documentation
 for the Switch 2200 system 3
dual homing 9-6
dual ring, FDDI 9-4

E

external network management 2-4

F

fax service B-3
FDDI 3-6
 capabilities 8-1
 connection rules 9-4
 dual homing 9-6
 dual ring 9-4
 events 3-5
 MIB 8-9
 networks 9-1 to 9-7
 nodes 8-6

- optical bypass switch 8-11
- OSI model, and 8-2
- overview 8-1
- paths 8-5
- ports 8-3
- proxy agents 3-7
- standards 8-2
- topologies 9-1
- trunk ring 9-3
- filtering. *See* packet filter
- flooding, packet 5-5, 7-3
- forwarding
 - packets 5-4
 - port state 5-16
- fragmentation. *See* IP fragmentation
- frame-based protocols 8-10

G

- groups, address. *See* address groups
- groups, port. *See* port groups

I

- IEEE 802.1d
 - bridging standards 5-1
- in-band management 4-2
- IP fragmentation
 - defined 7-2
 - RFC specified in 7-2

L

- LAN, FDDI 3-6
- LAN-layer 3-6
- layers, OSI Reference Model 8-2
- learning port state 5-16
- listening port state 5-16
- looping, packet 5-6

M

- M port 8-4
- MAC (Media Access Control) 8-2, 8-4
- management access 4-1 to 4-3
- managing the Switch 2200 system
 - from outside the network 2-4
 - protocols 3-1 to 3-4
- Master port 8-4

- MIB 3-5
 - events 3-5
 - FDDI 3-5, 8-9
 - SNMP A-1 to A-2
 - Switch 2200 3-5
- MIB II
 - traps 3-5
- multicast packet firewall
 - defined 7-1
 - multicast storm 7-2
 - receiveMulticastThreshold attribute 7-1
 - threshold mechanism 7-2
- multicast storm 7-1
- multicast/broadcast traffic 7-1

N

- network
 - FDDI 9-1 to 9-7
 - reduced packet flooding 7-3
 - security 7-4
- network supplier support B-3
- network topologies, FDDI 9-2
- newRoot 3-5
- nodes
 - FDDI, on 8-6
 - types 8-7

O

- on-line technical services B-1
- operand 6-1
- operator 6-1
- optical bypass switch
 - on intermediate systems 8-11
- OSI Reference Model
 - and FDDI 8-2
- out-of-band management 4-2

P

- packet
 - fields for operands 6-1
 - flooding 5-5, 7-3
 - forwarding 5-4
 - fragmentation 7-2
 - looping 5-6
 - reduced flooding 7-3
- packet filter
 - address groups 6-8, 6-10
 - AppleTalk network, and a 6-3
 - examples 6-3 to 6-7

- implementation on paths 6-1
- language 6-1
- purpose of 6-1
- user-defined 6-1 to 6-15
- paths
 - FDDI, on 8-5
 - receive 6-1
 - transmit 6-1
- PHY standard
 - defined 8-2
- physical access 4-1 to 4-3
- PMD standard
 - defined 8-2
- port
 - bridging states 5-16, 5-17
 - designated 5-8
 - FDDI 8-3
 - identifier 5-10
 - root 5-8, 5-9
 - table 6-9
- port groups
 - defined 6-9
 - mask 6-9
- protocols 3-1 to 3-4
 - frame-based 8-10
 - SMT 3-6
 - SNMP 3-4 to 3-6
 - virtual terminal 3-3
- proxy agents 3-7

R

- receive path 6-1
- returning products for repair B-4
- ring of trees 9-3
- rlogin 3-3
- RMA number B-4
- root bridge 5-8
- root port 5-8

S

- S port 8-4
- security
 - manual address assignment 7-4
- Slave port 8-4
- SMT 3-6, 8-2, 8-9

- SNMP 3-4 to 3-6
 - management applications 2-4
 - MIBs A-1 to A-2
 - traps 3-5
- Spanning Tree
 - algorithm 5-7
 - blocking paths 5-7
 - BPDU 5-18
 - CBPDU 5-8
 - designated bridge 5-8
 - designated port 5-8
 - packet looping, and 5-6
 - port identifier 5-10
 - reconfiguring the topology 5-18
 - root bridge 5-8
 - root port 5-8
- station
 - FDDI 8-6
- SunNet Manager™ 2-4
- Switch 2200 system
 - documentation About This Guide 3
 - external network management 2-4
 - managing from outside the network 2-4
 - overview of operations 1-1

T

- table
 - address 6-8
 - port 6-9
- TCP/IP 3-3
- technical support B-1
- telnet 3-3
- terminal
 - emulation 3-3
 - virtual 3-3
- token 8-4
- Topology Change Flag 5-18
- topology, FDDI 9-2 to 9-4
- Transcend® Enterprise Manager 3-4
 - SNMP community strings, and 3-6
- transmit path 6-1
- transparent bridging
 - aging addresses, and 5-3
 - defined 5-1
 - flooding packets, and 5-5
 - forwarding packets, and 5-4
 - IEEE 802.1d compliant 5-1
 - learning addresses, and 5-2
 - Spanning Tree 5-6
- traps, SNMP 3-5



U

UDP/IP networks 3-4
User Datagram Protocol 3-4
user interfaces on the Switch 2200 system 2-2
user-defined packet filtering 6-1 to 6-15

V

virtual terminal protocols 3-3
virtual workgroup 6-1