LANPLEX® 2500
OPERATION GUIDE

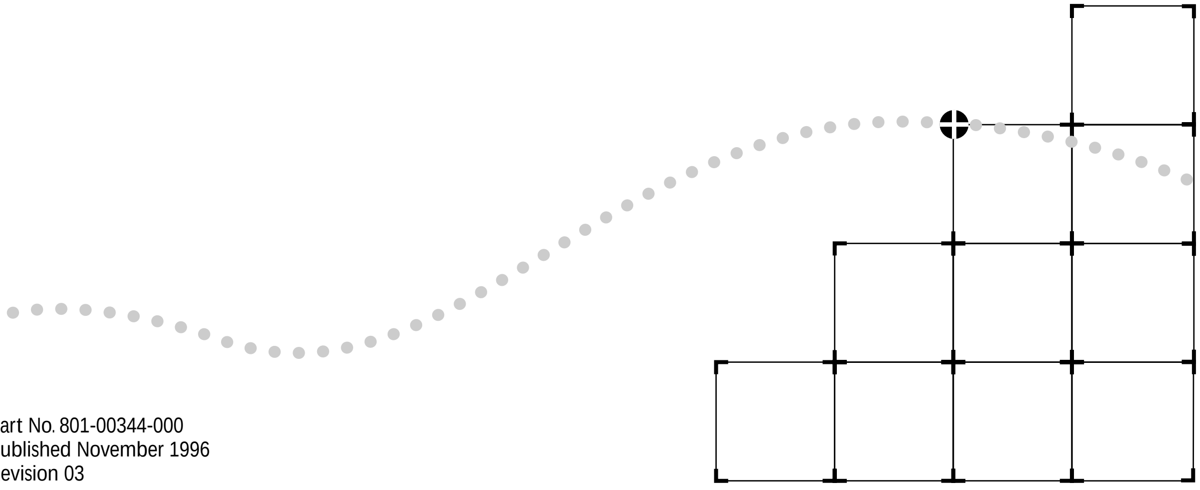

Part No. 801-00344-000
Published November 1996
Revision 03

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

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 Unpublished rights reserved under the copyright laws of the United States.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, the 3Com Logo, EtherDisk, EtherLink, EtherLink II, LANplex, LinkBuilder, NETBuilder, NETBuilder II, NetFacts, SmartAgent, TokenDisk, TokenLink, and Transcend are registered trademarks of 3Com Corporation. 3TECH, FDDILink, 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
LANplex 2500 Documentation 4
Documentation Feedback 5

PART I MANAGEMENT AND ADMINISTRATION

1 LANPLEX® MANAGEMENT AND ADMINISTRATION OVERVIEW

About the LANplex® 2500 System 1-1
User Interfaces and the LANplex® 2500 System 1-2

2 USER ACCESS: WHAT YOU SEE

About the User Interfaces to the LANplex® 2500 2-1
The Built-in LANplex® Administration Console 2-3
External Network Management Applications 2-4

3 MANAGEMENT ACCESS: PROTOCOLS

About LANplex® 2500 Protocols 3-1
Virtual Terminal Protocols 3-3
SNMP 3-4
    SNMP Agent 3-4
    SNMP MIBs 3-4
    SNMP Traps 3-5
    Access Control 3-6
SMT 3-6
SNMP and SMT Proxy Agents 3-7

4 Physical Access: Ports and Cabling

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

Part II Bridging Technology

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 5-12
Spanning Tree Port States 5-16
Reconfiguring the Bridged Network Topology 5-18
Bridging References 5-18

6 Express Switching

About Express Switching 6-1
Address Learning 6-3
Packet Filtering and Forwarding 6-3
Using Express Switching 6-3

7 USER-DEFINED PACKET FILTERING

About User-defined Packet Filtering 7-1
Designing a Packet Filter 7-1
Assigning Packet Filters to Paths 7-2
Packet Filter Examples 7-3
Example 1: Isolating IP Segments 7-4
Example 2: Filtering AppleTalk Phase II Packets 7-6
Using Address Groups and Port Groups in a Packet Filter 7-8
What Is an Address Group? 7-8
What Is a Port Group? 7-9
Referencing Address Groups and Port Groups From a Packet Filter 7-11
Example: Using Address Groups in a Packet Filter 7-11
Globally Administering Packet Filters 7-16

8 LANPLEX® BRIDGING EXTENSIONS

Multicast Packet Firewalls 8-1
IP Fragmentation 8-2
Reducing Packet Flooding 8-3
Enhanced Network Security 8-4

PART III FDDI TECHNOLOGY

9 FDDI OVERVIEW AND IMPLEMENTATION

About FDDI 9-1
Ports 9-3
MACs 9-4
MAC Services 9-4
MAC Operation 9-4
Paths 9-5
Nodes and Attachments 9-5
Nodes 9-5
Attachments 9-6
Node Types 9-6
Station Management 9-8
SMT Operation 9-8
FDDI MIB 9-8
Frame-based Protocols 9-9
FDDI and the LANplex 2500 System 9-10

10 FDDI Networks

About FDDI Networks 10-1
FDDI Network Topologies 10-2
Physical Topology: A Ring of Trees 10-3
Logical Topology: The Dual Ring 10-4
FDDI Connection Rules 10-4
Dual Homing 10-6

Part IV ATM Technology

11 ATM Networks

About ATM 11-1
ATM Basics 11-2
ATM Network Interfaces 11-2
ATM Cell Structure 11-3
Virtual Connections in ATM 11-5
Virtual Circuits 11-5
Switched Virtual Circuits 11-5
Permanent Virtual Circuits 11-6
Virtual Paths and Virtual Channels 11-6
ATM Protocol Architecture 11-7
The Physical Layer 11-8
SONET Physical Layer Interface 11-8
The ATM Layer 11-8
Quality of Service 11-8
The ATM Adaptation Layer (AAL) 11-9
AAL Sublayers 11-9
Interim Local Management Interface (ILMI) Communication Protocol 11-10
ILMI Management Information Base (MIB) 11-10
ATM Address Registration 11-10
Integrating ATM in Legacy LANs 11-11
LAN Emulation (LANE) 11-12
LANE Components 11-14
Virtual Channel Connections (VCC) 11-15
Emulated LANs (ELANs) 11-15
ATM and the LANplex System 11-16
Configuring ATM 11-16
LAN Emulation 11-16
Creating an Emulated LAN 11-17
Bridge Loops in ELAN Configurations 11-17

PART V APPENDIXES

A SNMP MIB SUPPORT

SNMP MIBs A-1

SNMP MIB Compilers A-4

B TECHNICAL SUPPORT

Online Technical Services B-1

- 3Com Bulletin Board Service B-1
  - Access by Analog Modem B-1
  - Access by Digital Modem B-2
- World Wide Web Site B-2
- 3ComForum on CompuServe® Online Service B-2
- 3ComFactsSM Automated Fax Service B-3

Support from Your Network Supplier B-3

Support from 3Com B-4

Returning Products for Repair B-4

OPERATION GLOSSARY

INDEX

About This Guide

Introduction

This *LANplex® 2500 Operation Guide* provides all the information you need to understand how your LANplex® switching hub works in FDDI, Ethernet, and ATM networking environments.

Audience

This guide is intended for the system or network administrator who is responsible for configuring, using, and managing the LANplex® 2500 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 software installation and release notes shipped with your product differs from the information in this guide, follow the release notes.

How to Use This Guide

Table 1 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 LANplex system operation	Chapter 1
Access interfaces for the user (network administrator)	Chapter 2
Management protocols	Chapter 3
Cabling the LANplex system for management access	Chapter 4
Transparent bridging issues	Chapter 5
Spanning Tree information	Chapter 5
Express Switching	Chapter 6
User-defined packet filters	Chapter 7
Bridging extensions	Chapter 7
LANplex bridging extensions	Chapter 8
FDDI overview and implementation	Chapter 9
FDDI networks	Chapter 10

(continued)

Table 1 Locating Information in This Guide (continued)

If you are looking for information on...	Turn to...
Asynchronous Transfer Mode (ATM) technology	Chapter 11
SNMP MIB support	Appendix A
3Com Technical Support	Appendix B
Definitions for operating the system	Glossary

Conventions

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

Table 2 Notice Icons

Icon	Type	Description
ℹ	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"	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: The following syntax specifies the time and date: mm/dd/yy hh:mm:ss The word "command" indicates that all variables in the command have been supplied and you must enter the command as shown. Example: The following command enables Spanning Tree: **`bridge stpState enabled`**
`screen display`	`This typeface` represents displays that appear on your terminal screen. Example: `Login:`
`commands`	**`This typeface`** represents commands that you enter. Example: **`bridge stpState disabled`**
Keys	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]. If you must press two or more keys simultaneously, the keys are linked with a plus sign (+). Example: Press [Ctrl]+[Alt]+[Del].
Italics	*Italics* are used to denote *emphasis* or new terms where they are defined.
Bold	**Bold** is used to denote key features, menus, and menu options.

LANplex 2500 Documentation

The following documents comprise the LANplex 2500 documentation set. If you want to order additional documents or one that you do not have, contact your sales representative for assistance.

- *LANplex® 2500/2016 Unpacking Instructions*

  Describes how to unpack your LANplex system. It also provides you with an inventory list of all the items shipped with your system. (Shipped with your system)

- *LANplex® 2500 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 LANplex system's documentation. (Shipped with your system)

- *LANplex® 2500 Getting Started*

  Describes all the procedures necessary for planning your configuration and for installing, cabling, powering up, configuring management access, and troubleshooting your LANplex system. (Shipped with your system/Part No. 801-00335-000)

- *LANplex® 2500 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 LANplex system. (Shipped with your system/Part No. 801-00344-000)

- *LANplex® 2500 Administration Console User Guide*

  Provides information about using the Administration Console embedded system software to configure and manage your LANplex system. (Shipped with your system/Part No. 801-00322-000)

- *LANplex® 2500 Extended Switching User Guide*

  Describes how the routing protocols are implemented in the LANplex system and provides information about using the Administration Console to configure and manage your routing protocols. (Shipped with the option package/Part No. 801-00343-000)

- *LANplex® 2500 Intelligent Switching Administration Console Command Quick Reference*

  Contains Administration Console intelligent switching commands for the LANplex system. (Folding card; shipped with your system/Part No. 801-000318-000)

- *LANplex® 2500 Extended Switching Administration Console Command Quick Reference*

  Contains all of the Administration Console commands for Extended Switching Options in the LANplex system. (Folding card, shipped with the system/Part No. 801-000319-000)

- *Module Installation Guides*

  Provide an overview, installation instructions, LED status information, and pin-out information for the particular option module. (Shipped with individual modules)

Documentation Feedback

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:

LANplex® 2500 Administration Console User Guide
Part No. 801-00322-000
Page 2-5 (chapter 2, page 5)

I

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

LANplex® Management and Administration Overview

This chapter introduces you to how your LANplex® system is managed and administered.

About the LANplex® 2500 System

The LANplex 2500 system combines high port density, Ethernet switching, Ethernet-to-FDDI bridging, Fast Ethernet switching, FDDI switching, and ATM switching in an integrated system. You can configure much of this functionality to meet your specific networking needs.

LANplex system management and administration occur through the layers of the OSI reference model. See Figure 1-1.

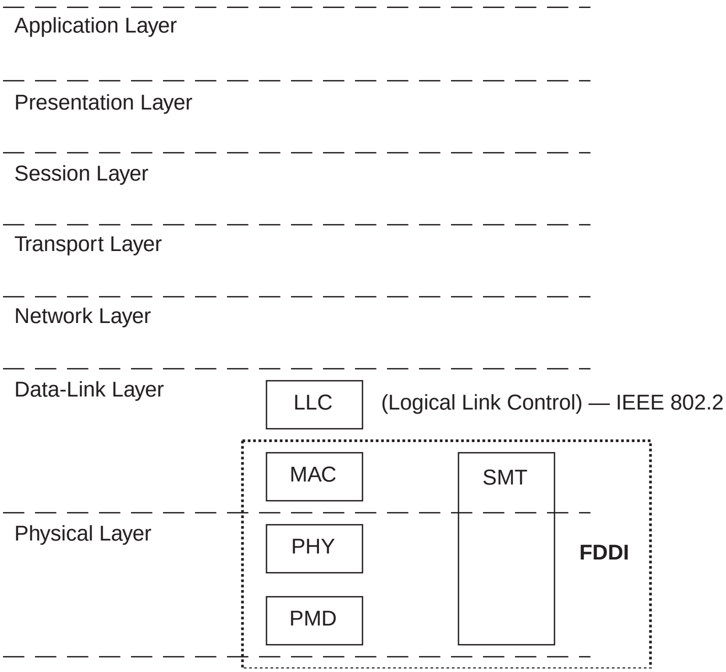

Figure 1-1 OSI Reference Model

Using the LANplex system's own management application or an external SNMP- based management application, the system can be managed and administered through various protocols and physical interfaces.

User Interfaces and the LANplex® 2500 System

Figure 1-2 shows how the levels of the OSI reference model are integrated in the management and administration device as well as in the LANplex system. The illustration also indicates where to find related information in this guide.

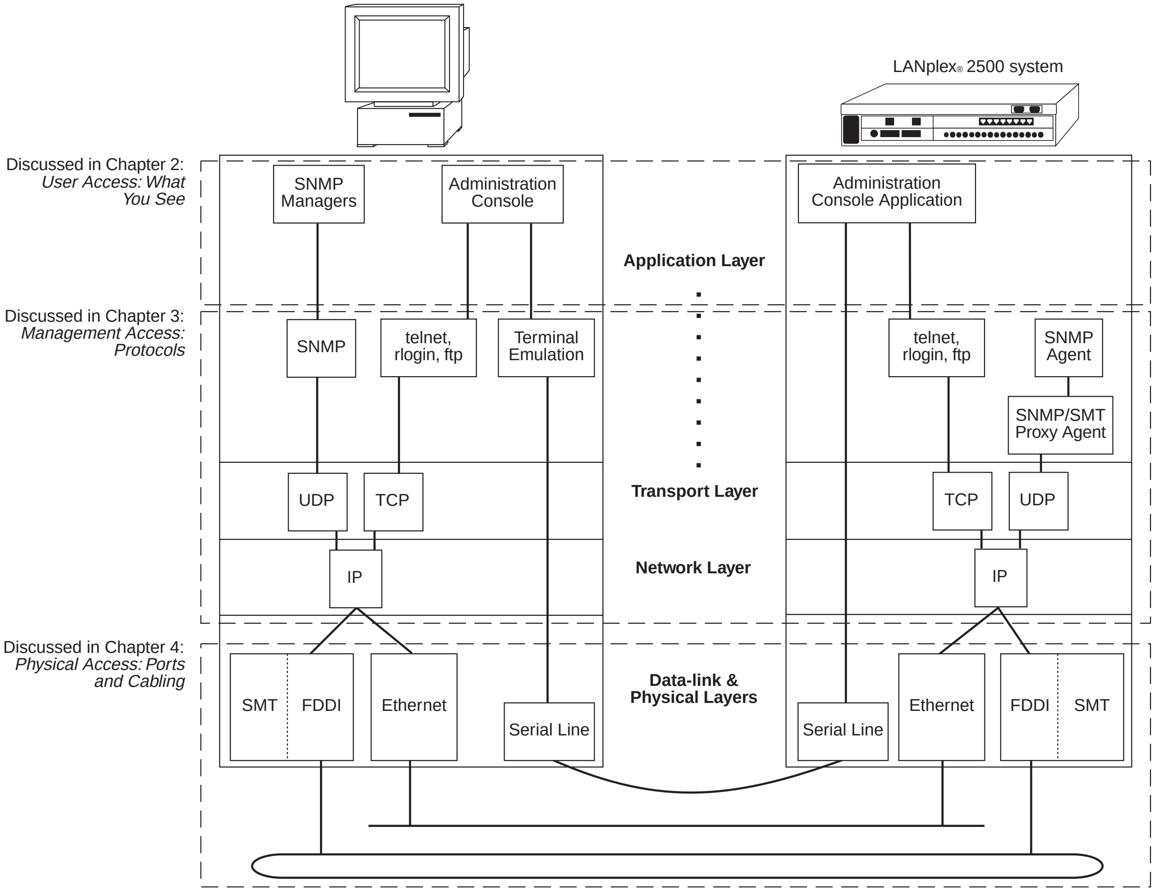

Figure 1-2 User Interfaces and Protocols Used to Access the LANplex System

2

USER ACCESS: WHAT YOU SEE

This chapter describes the applications you can use to gain access to your LANplex system and to perform administrative and management functions.

About the User Interfaces to the LANplex® 2500

You can use the following applications as the user interface to your LANplex system:

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

Figure 2-1 shows the LANplex system user interfaces in the OSI reference model.

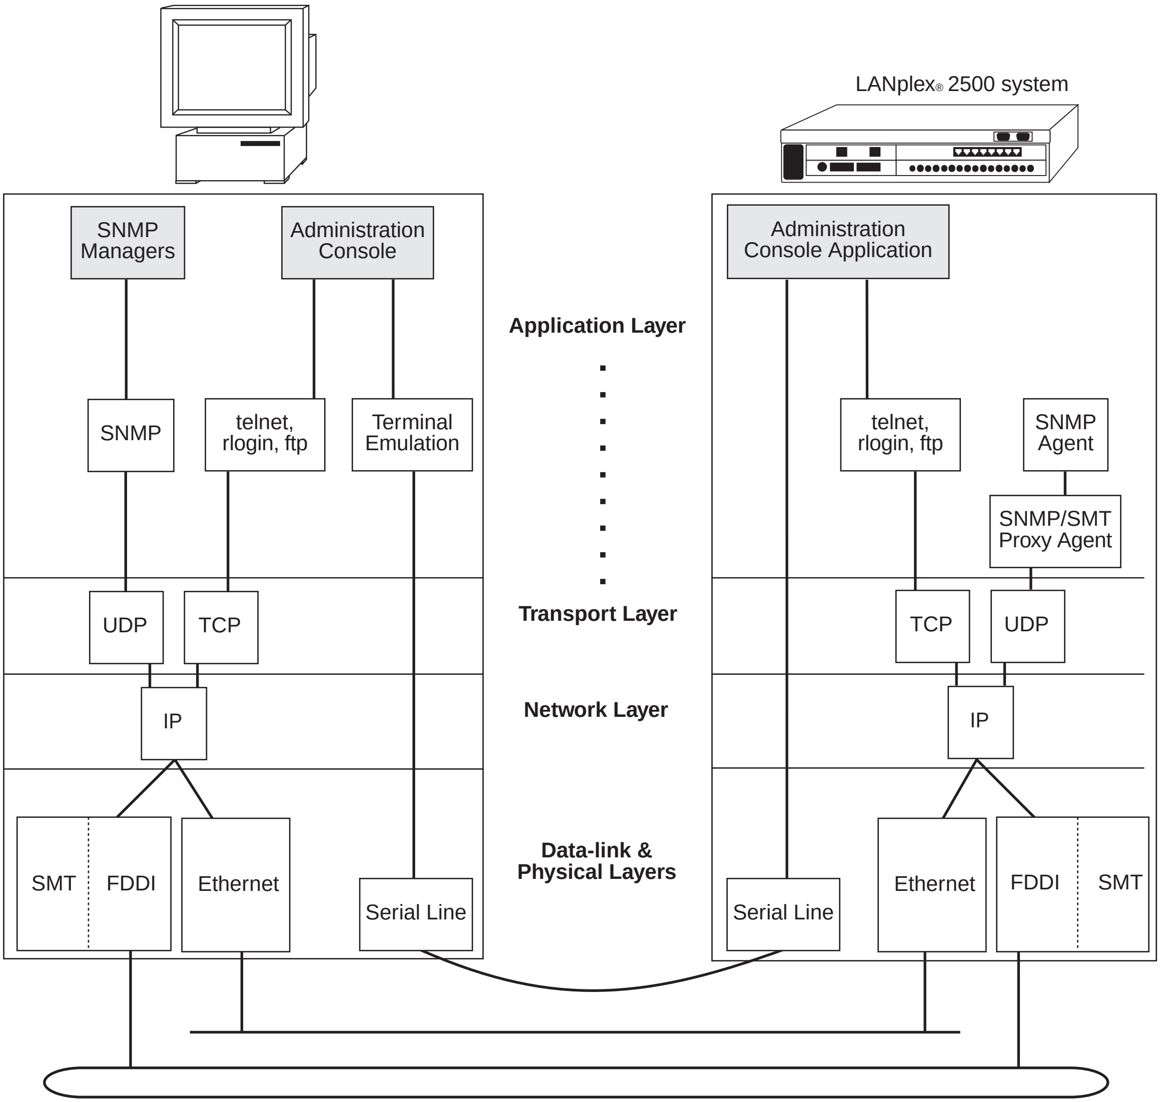

Figure 2-1 User Interfaces for the LANplex System

The Built-in LANplex® Administration Console

You can use the Administration Console to configure your LANplex 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.

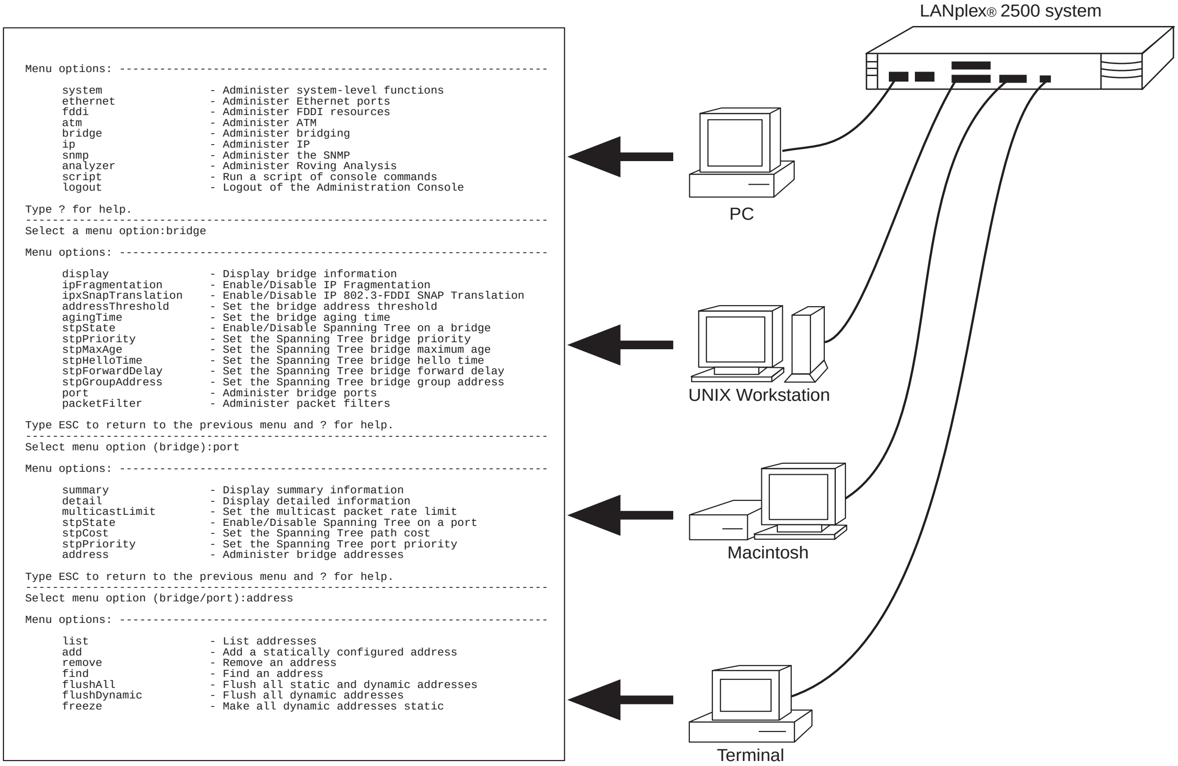

```
Menu options: ----------------------------------------------------------------

      system                - Administer system-level functions
      ethernet              - Administer Ethernet ports
      fddi                  - Administer FDDI resources
      atm                   - Administer ATM
      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
      ipxSnapTranslation    - 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
```

Figure 2-2 Administration Console for the LANplex System

For more information about the Administration Console, see the *LANplex* 2500 *Administration Console User Guide*.

External Network Management Applications

3Com's Transcend® Enterprise Manager is a network management software family suite that runs on UNIX and MS-DOS platforms. It provides network management for a wide range of 3Com products, including the LANplex 2500 system. With Transcend Enterprise Manager software, you get a device view of the LANplex 2500 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.

To order Transcend Enterprise Manager, contact your sales representative.

Figure 2-3 shows an example of a Device Manager Screen from the Transcend Enterprise Manager for UNIX software. Figure 2-4 shows an example of a Status View screen from the Transcend Enterprise Manager for UNIX software.

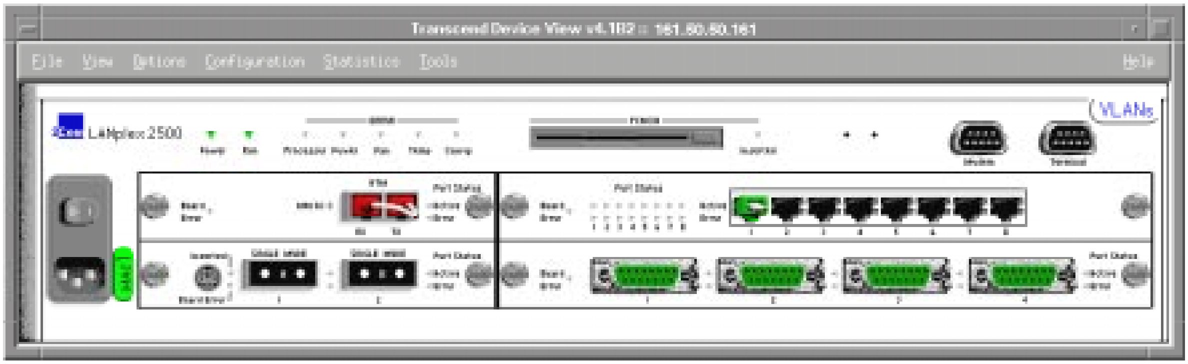

Figure 2-3 Sample Screen from Transcend® Enterprise Manager Device Manager for Hubs Software

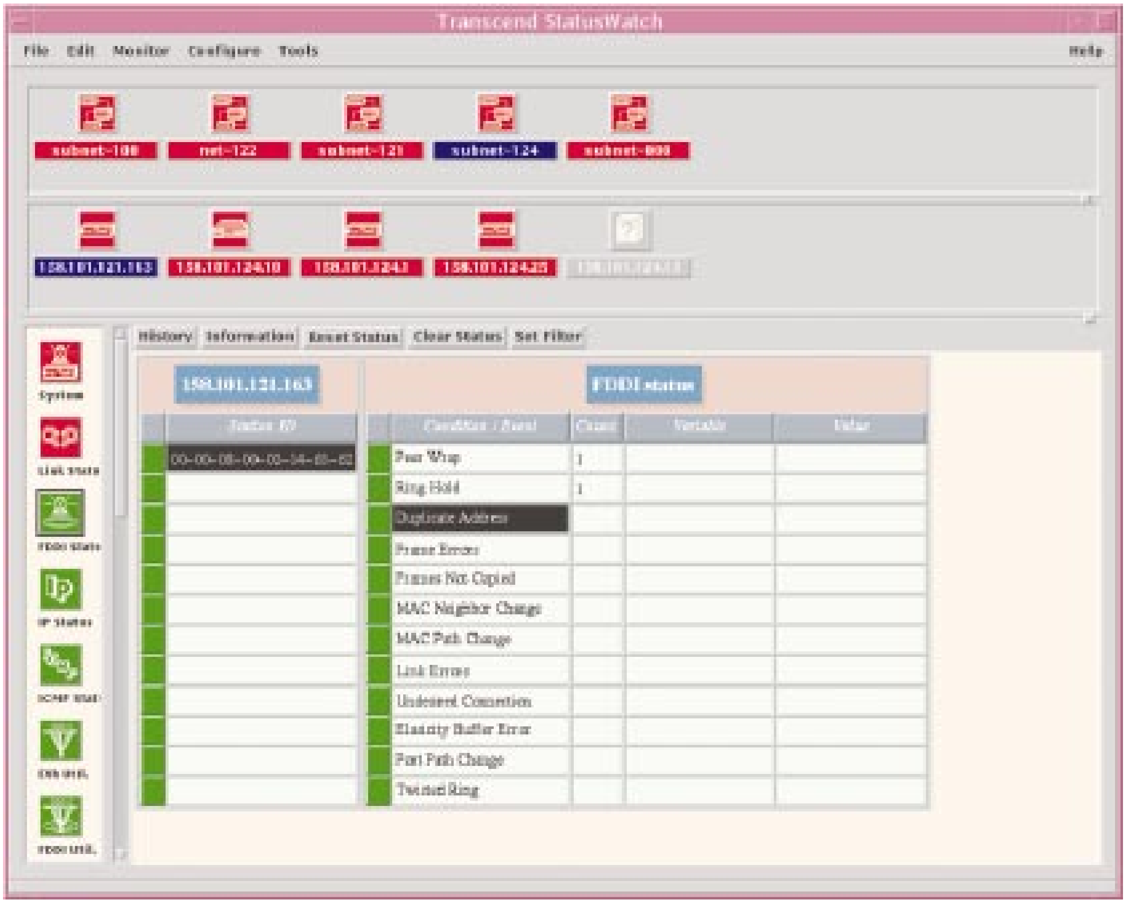

Groups and individual devices

Devices in the selected group

	Station ID
🟩	00-00-08-09-02-14-03-02
🟩	
🟩	
🟩	
🟩	
🟩	
🟩	
🟩	
🟩	
🟩	
🟩	
🟩	

	Condition / Event	Count	Variable	Value
🟩	Peer Wrap	1		
🟩	Ring Hold	1		
🟩	Duplicate Address			
🟩	Frame Errors			
🟩	Frames Not Copied			
🟩	MAC Neighbor Change			
🟩	MAC Path Change			
🟩	Link Errors			
🟩	Undesired Connection			
🟩	Elasticity Buffer Error			
🟩	Port Path Change			
🟩	Twisted Ring			

Figure 2-4 Sample Screen from Transcend® Enterprise Manager Status View Software

Because the LANplex system is based on SNMP standards, you can manage your system using third-party SNMP-based network manager applications, such as Sun Microsystems SunNet Manager™, Hewlett-Packard OpenView®, or IBM NetView for 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 LANplex system.

About LANplex® 2500 Protocols

The LANplex 2500 system 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 OSI reference model for the network environment.

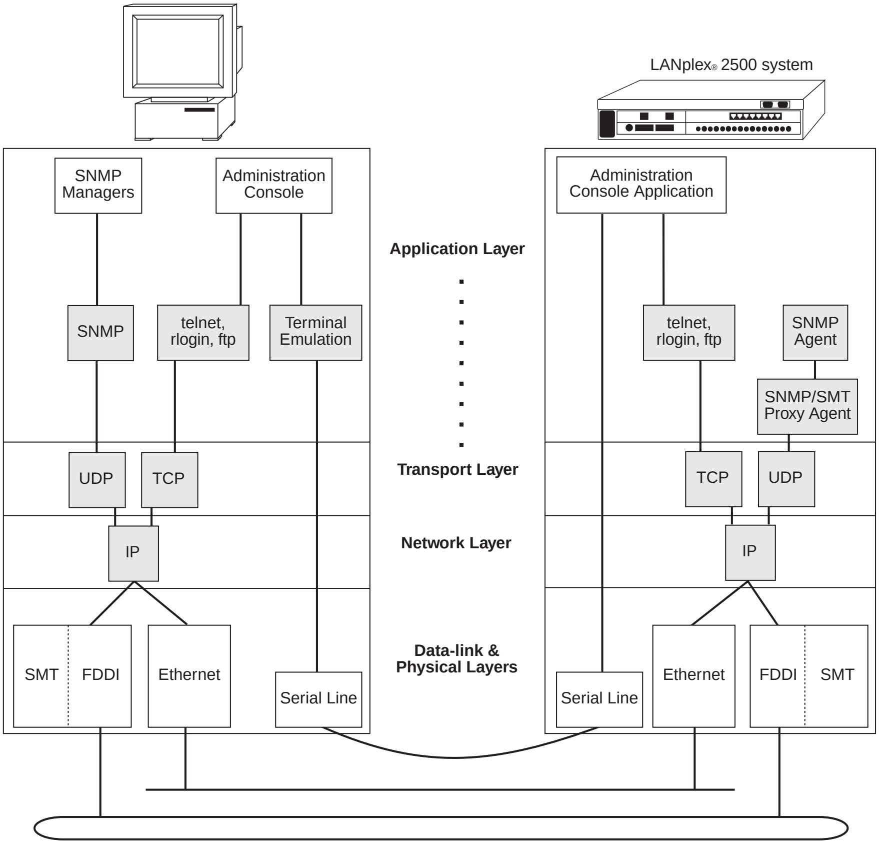

Figure 3-1 Protocol Stacks for the LANplex® System

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 LANplex 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 you must connect a terminal directly to the serial line.

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

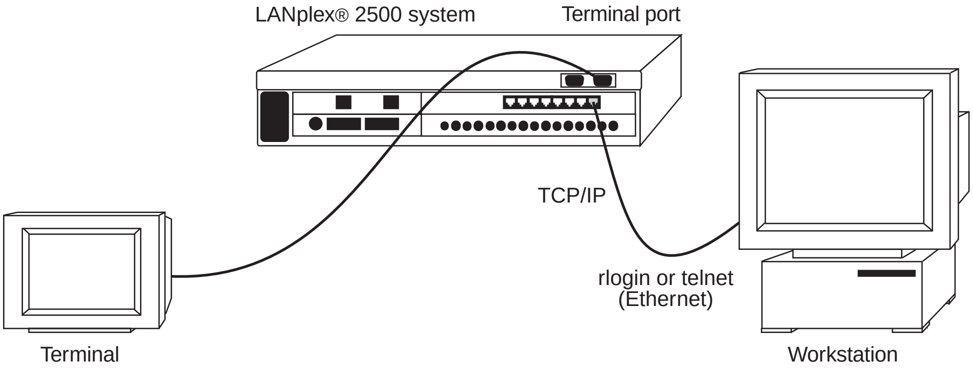

Figure 3-2 Administration Console Access for the LANplex System

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 LANplex system contains 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 LANplex system. 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 LANplex system supports the following SNMP MIBs:

- ATM MIB
- Bridge MIB
- Ethernet MIB
- IF MIB
- FDDI SMT 7.3 MIB
- LANplex Systems MIB
- LANplex Optional FDDI MIB
- LEC MIB (af-lane-0044.000)
- LES MIB (af-lane-1129.001)
- MIB II

For more information on which MIBs are supported, see Appendix A: *SNMP MIB Support* and the software *Installation and 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 network management station (NMS) 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 LANplex system.

Table 3-1 SNMP Traps in the LANplex 2500 system

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

For descriptions of these 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 LANplex 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 LANplex 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 LANplex system to work with an SNMP manager, you must configure the LANplex system's SNMP community strings to match those used by the SNMP manager.

For information on how to configure community strings, see the *LANplex® 2500 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 LANplex system'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 *LANplex® 2500 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 LANplex 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 LANplex 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 LANplex system through its physical interfaces. Figure 4-1 highlights the system's physical access options in the OSI reference Model environment.

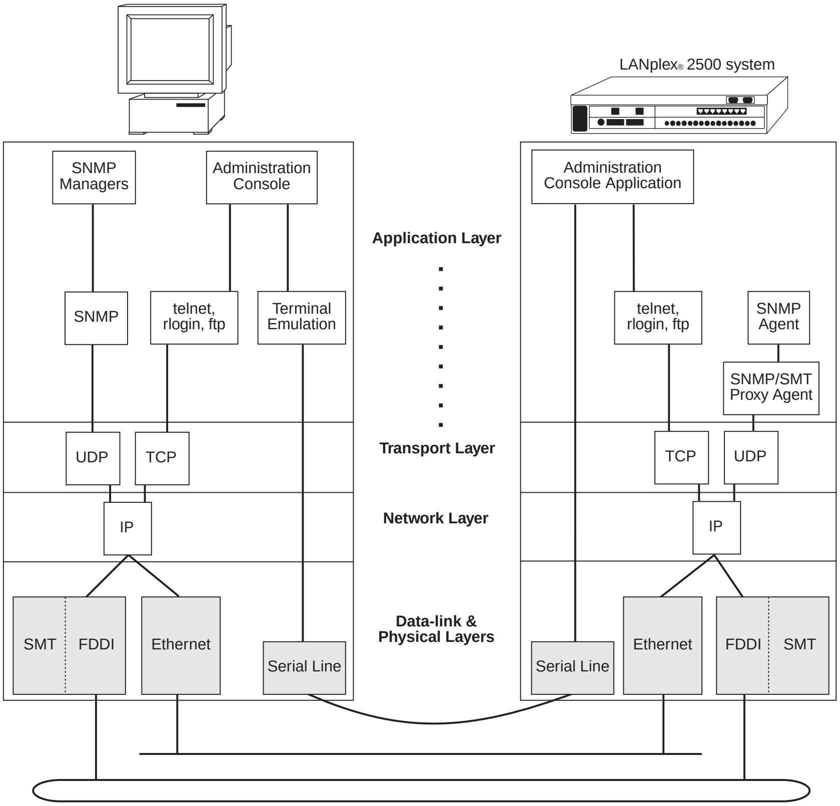

Figure 4-1 Physical Access Options for the LANplex® system

In-band and Out-of-band Management

If you manage your LANplex 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 LANplex 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 LANplex 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 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 *LANplex 2500 Administration Console User Guide*.

Management Access

You can access the LANplex system through 1) a serial port or 2) any Ethernet or FDDI port. These methods are described next.

Serial Ports

Access through each of the serial ports is described below.

- **Terminal serial port** — Direct access through the terminal 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 terminal port. A workstation attachment under UNIX can use the emulator TIP.
- **Modem serial port** — You can access the Administration Console remotely through an external modem attached to the modem port.

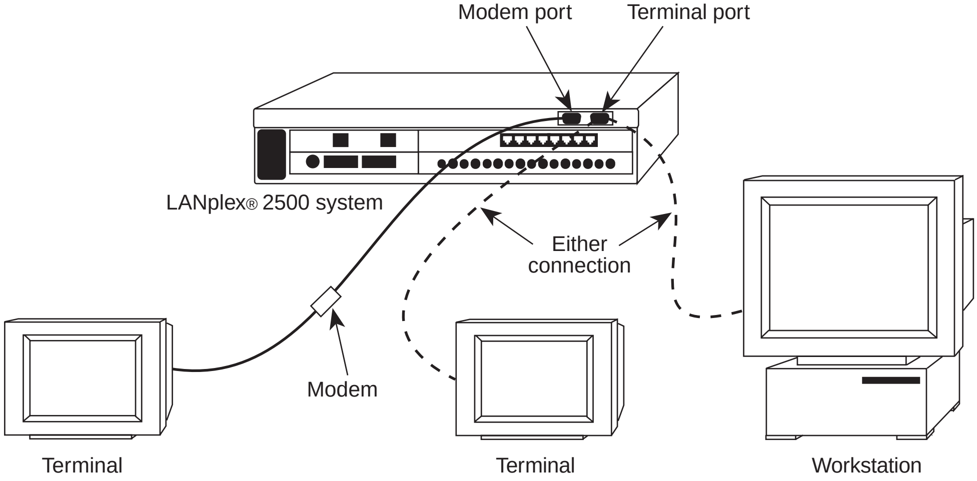

Figure 4-2 Access Through the Console Port

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 the ports. The LANplex SNMP agent can also be accessed through these interfaces.

The following Ethernet port types are possible:

- 10BASE-T (RJ-21)
- 10BASE-T (RJ-45)
- 10BASE-5 (AUI)
- 10BASE-FL
- 10BASE-2 (BNC)
- 100BASE-FX
- 100BASE-TX

The following FDDI connections are available:

- Dual Attached Station (DAS) MIC
- DAS TP

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

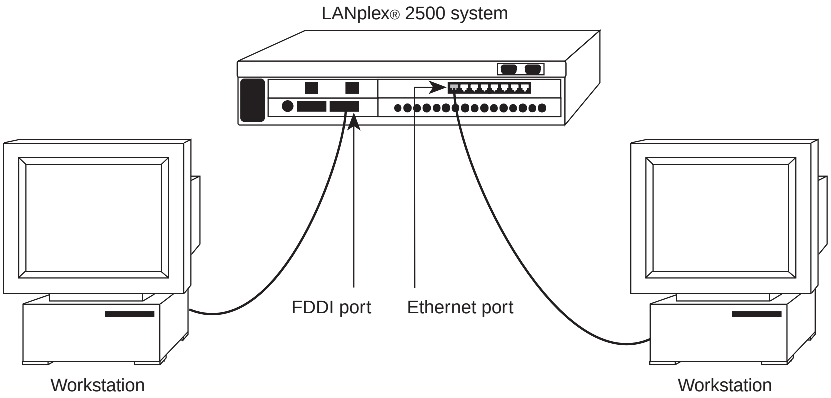

Figure 4-3 Access to the LANplex System Through any Ethernet or FDDI Port

II

BRIDGING TECHNOLOGY

Chapter 5 Transparent Bridging

Chapter 6 Express Switching

Chapter 7 User-defined Packet Filtering

Chapter 8 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.

What Makes a Bridge 802.1d Compliant?

Transparent bridging has been adopted for standardization by the IEEE and is defined in the IEEE 802.1d specification. The IEEE 802.1d bridging standard specifies many requirements with which a transparent bridge must comply. An 802.1d compliant 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 LANplex 2500 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.

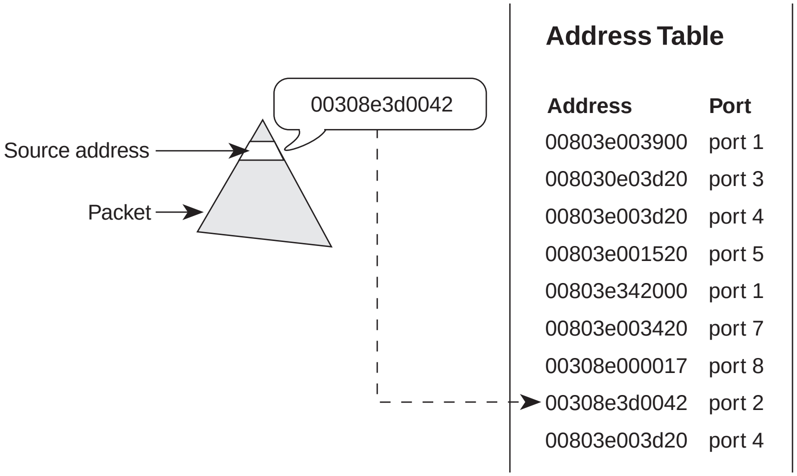

Address Table

Address	Port
00803e003900	port 1
008030e03d20	port 3
00803e003d20	port 4
00803e001520	port 5
00803e342000	port 1
00803e003420	port 7
00308e000017	port 8
00308e3d0042	port 2
00803e003d20	port 4

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 LANplex 2500 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 LANplex 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 *LANplex® 2500 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.

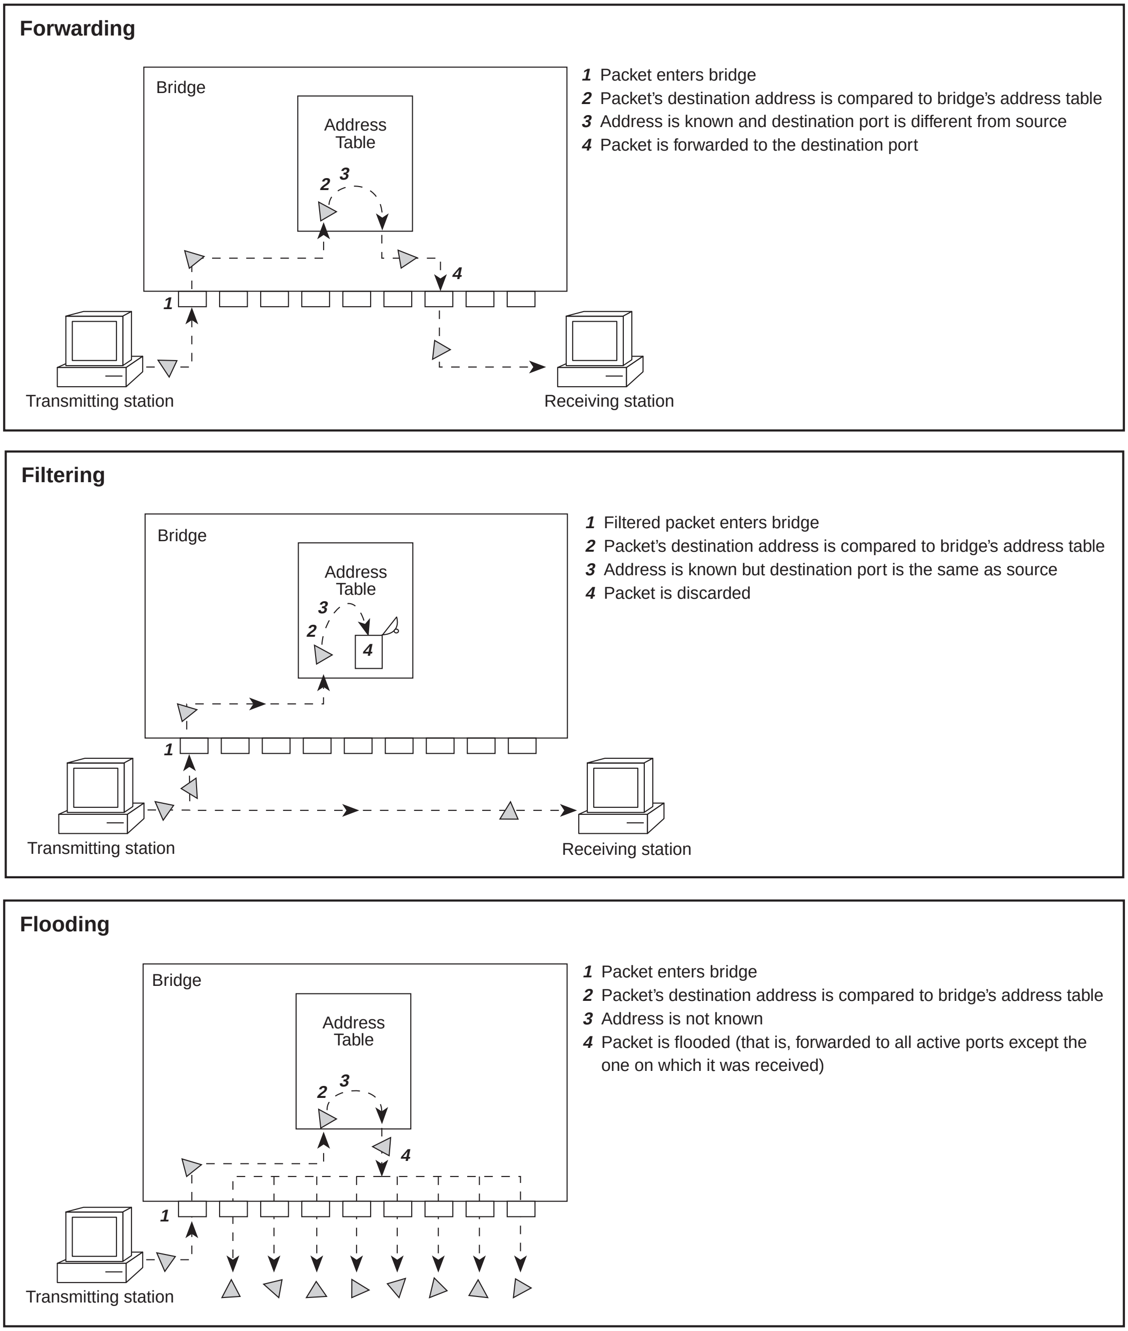

Forwarding

1 Packet enters bridge
2 Packet's destination address is compared to bridge's address table
3 Address is known and destination port is different from source
4 Packet is forwarded to the destination port

Filtering

1 Filtered packet enters bridge
2 Packet's destination address is compared to bridge's address table
3 Address is known but destination port is the same as source
4 Packet is discarded

Flooding

1 Packet enters bridge
2 Packet's destination address is compared to bridge's address table
3 Address is not known
4 Packet is flooded (that is, forwarded to all active ports except the one on which it was received)

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 standards require 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 in which 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.

In the example of packet looping shown in Figure 5-3:

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 on to 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.

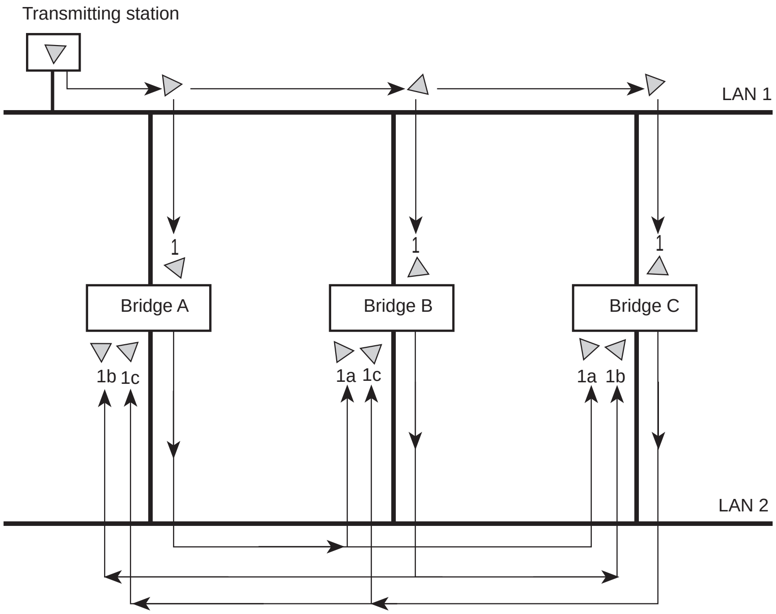

Figure 5-3 Packets Looping and Multiplying Without Spanning Tree Protocol

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 activation ensures that all stations remain connected.

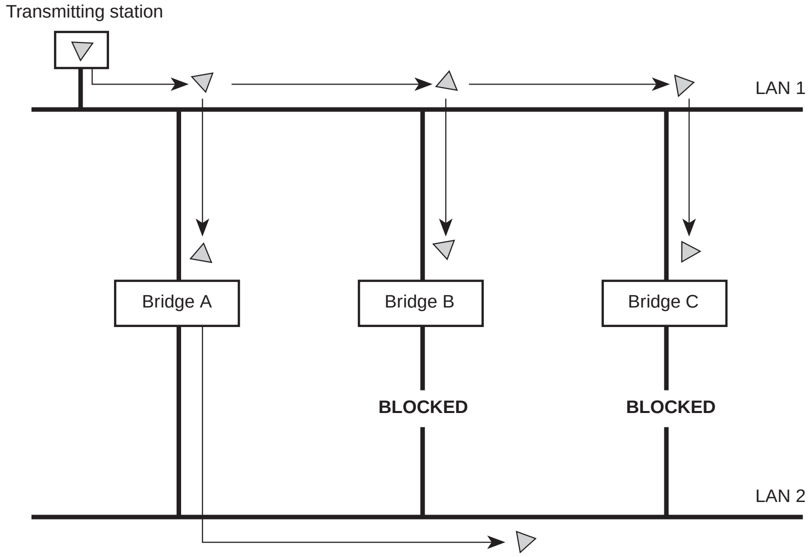

Figure 5-4 Spanning Tree Protocol Implemented to Block Redundant Links

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 as regular data packets are. 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.

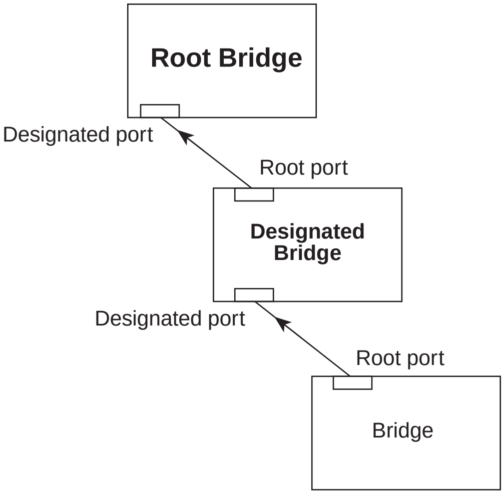

Figure 5-5 Hierarchy of the Root Bridge and the Designated Bridge

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.

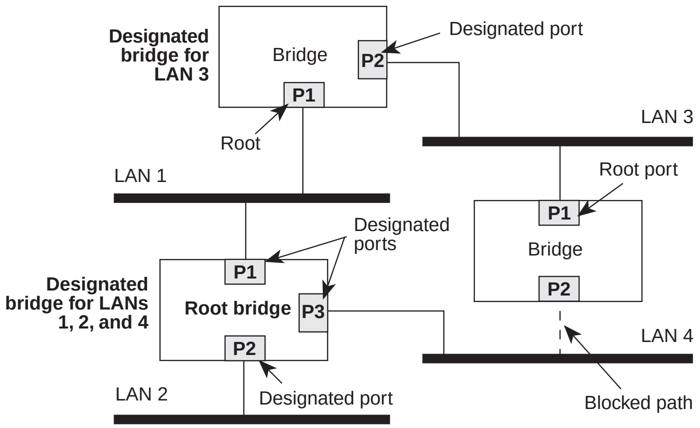

Figure 5-6 Root and Designated Bridges and Ports in a Spanning Tree Topology

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 tiebreaker in which the lowest port identifier takes priority. This field is useful primarily 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 the bridge saves the message.

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 Message 1 and Message 2, but cost is lower in Message 1. The bridge saves Message 1.

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 Message 1 and Message 2, but the transmitting bridge ID is lower in Message 1. The bridge saves Message 1.

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 this: *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 CPBDU with the contents 52.0.52, then it assumes 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 the contents of its own CBPDU, 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*.

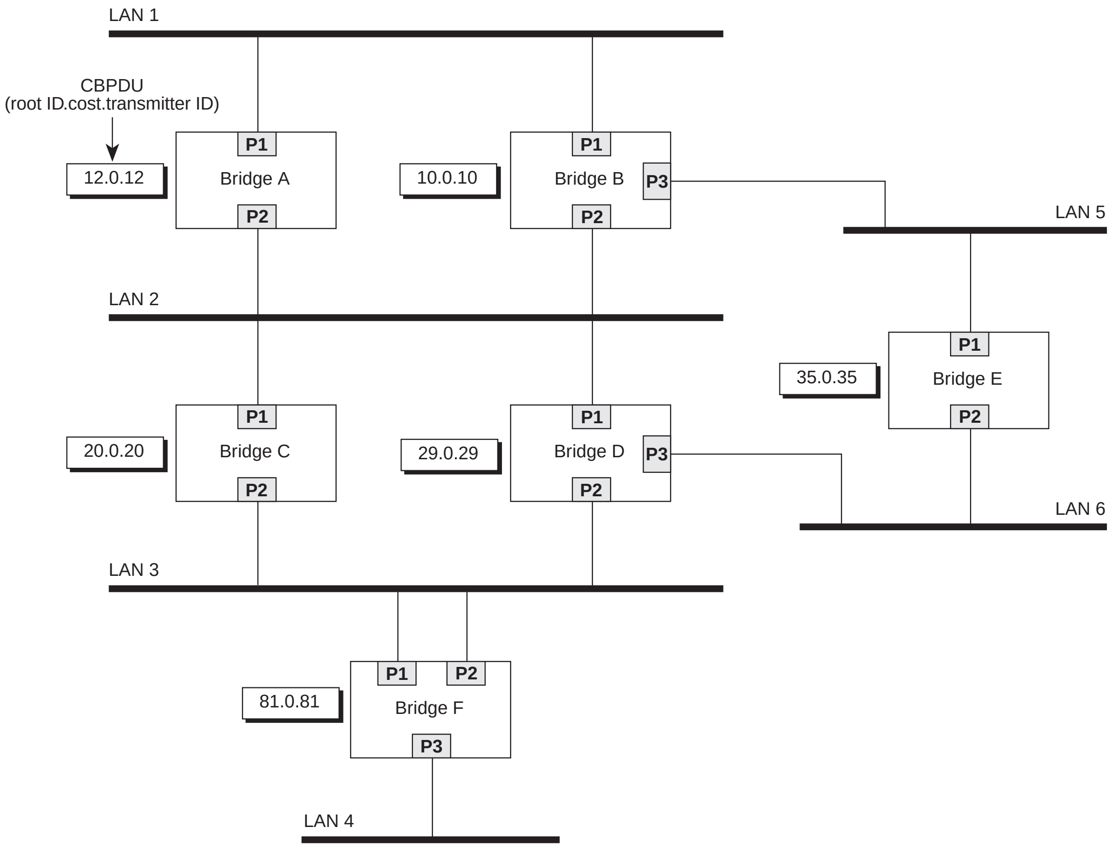

Figure 5-7 Starting the Spanning Tree Calculation

The root ID portion of the CBPDU 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.

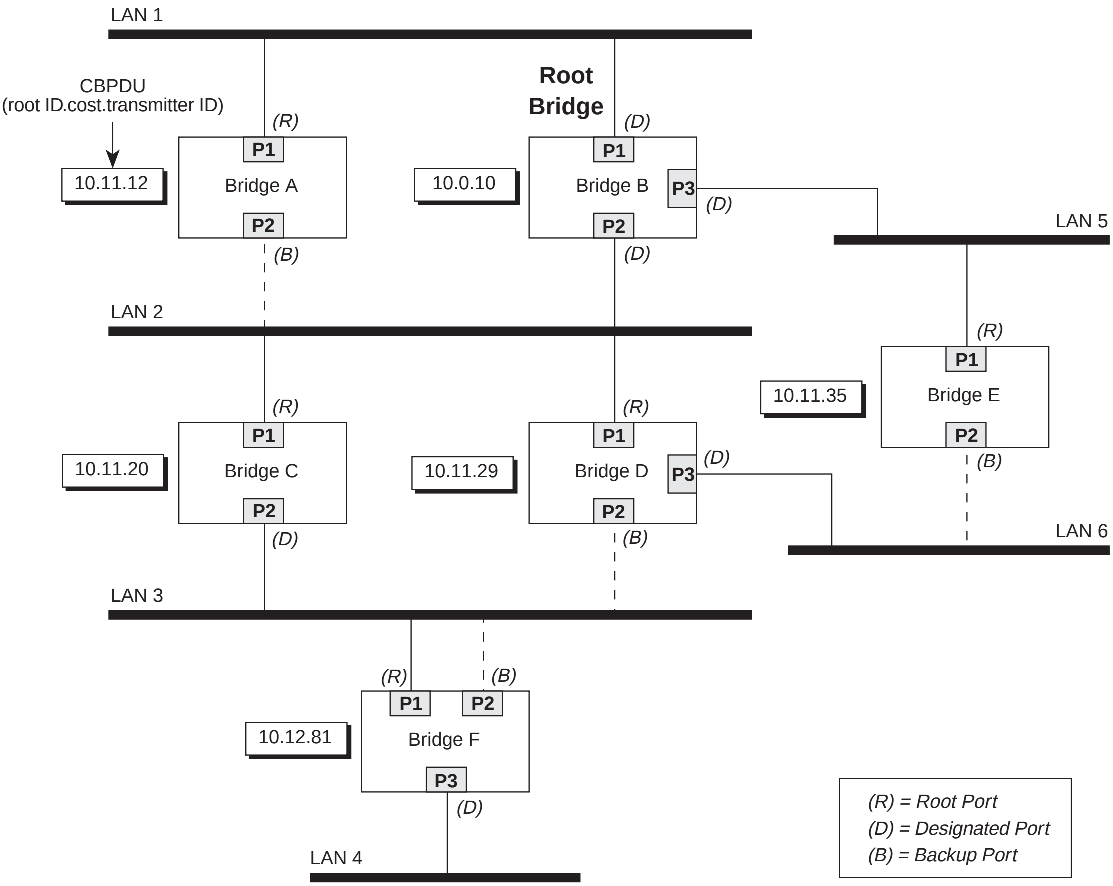

Figure 5-8 Spanning Tree Topology Calculated

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	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. 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 period of being in 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.
Learning	The learning state is similar to the listening state except that data packets are received on that port for the purpose of learning the stations that are 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 change the port back to the blocking state, then the bridge changes the port to the forwarding state. The time the port spends in both the listening and learning states is determined by the value of the *forward delay* 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.
Forwarding	After the port enters 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.
Blocking	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. In addition, the bridge does not receive data packets from the port, learn locations of station addresses from it, or forward packets onto it.

(continued)

Table 5-1 Spanning Tree Port States (continued)

Port State	Description
Disabled	A port is disabled when Spanning Tree Protocol 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 continues to forward frames only if Spanning Tree Protocol is disabled for the entire bridge.

Figure 5-9 illustrates the factors that cause a port to change 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. After the port enters the listening state, forward delay must expire before the port can transition to the learning state. If a port in listening, learning, or forwarding state is disabled by the network administrator or by a failure or initialization, then that port becomes disabled.

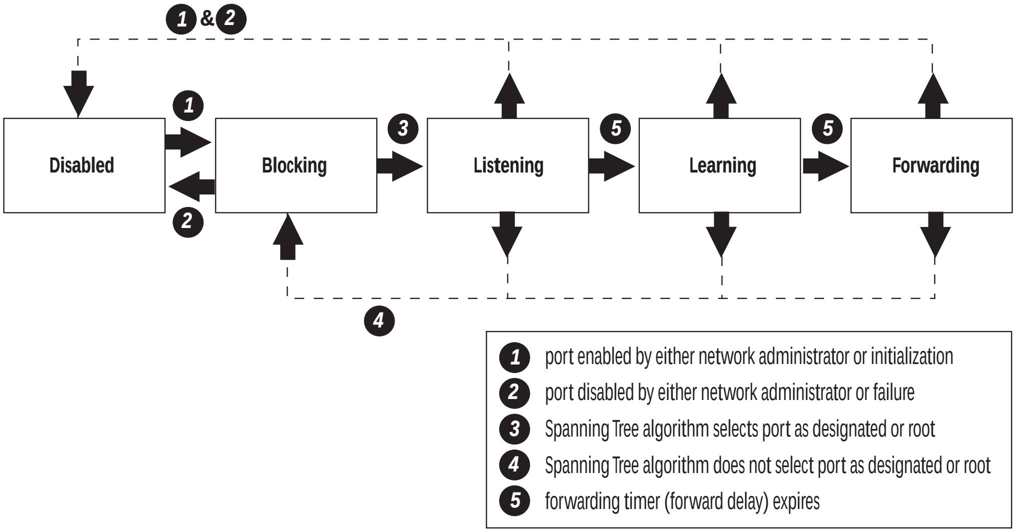
1 & 2
1
2
3
5
5
4
Disabled
Blocking
Listening
Learning
Forwarding
1 port enabled by either network administrator or initialization
2 port disabled by either network administrator or failure
3 Spanning Tree algorithm selects port as designated or root
4 Spanning Tree algorithm does not select port as designated or root
5 forwarding timer (forward delay) expires

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 changes 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

EXPRESS SWITCHING

This chapter describes these aspects of Express Switching mode:

- Technology and topology
- Address learning and aging
- Packet filtering and forwarding
- Advantages and constraints

About Express Switching

The LANplex 2500 system supports a mode of bridging called *Express Switching*. This mode is slightly different from the transparent bridging described in Chapter 5: *Transparent Bridging*.

In Express Switching mode, one of the bridge ports is designated as the *backbone port*. A unicast frame is forwarded *from* the backbone port only if the bridge has learned the destination address on another port. Conversely, a unicast frame is forwarded *to* the backbone port only if the destination address has *not* been learned on another port. As a result, the system floods no unicast frames to unknown addresses. (Multicast and broadcast frames are forwarded to all ports.)

Spanning Tree is configurable while in Express Switching mode.

Optimally, Express Switching is used in flat segmentation topologies containing a smaller number of high-performance Ethernet end-stations that are bridged directly to an FDDI backbone. See Figure 6-1.

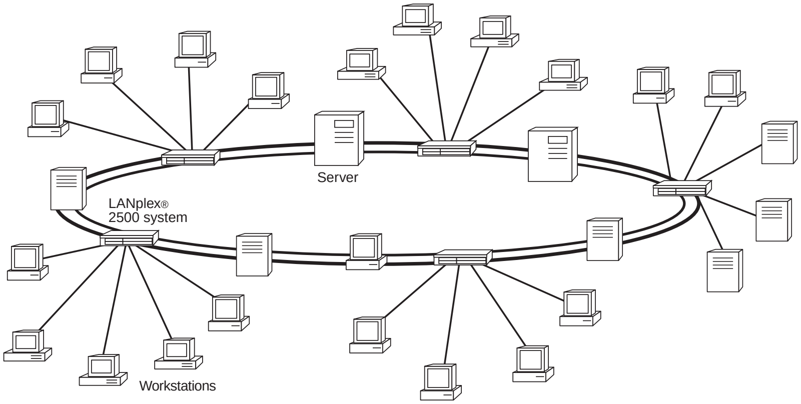

Figure 6-1 Express Switching — Flat Topology

Traditional transparent bridging is preferable in environments that contain tree or mesh topologies that mix extended Ethernet and FDDI networks. See Figure 6-2.

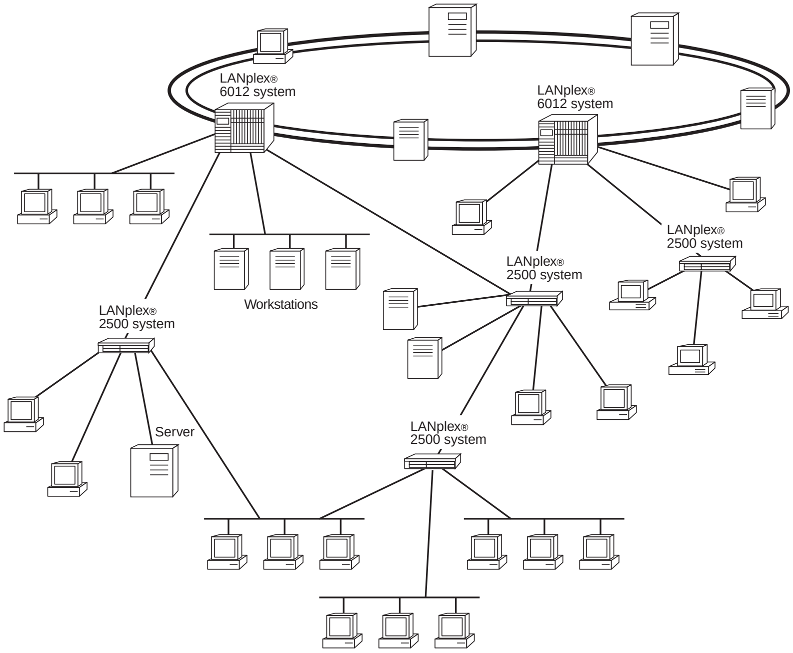

Figure 6-2 Traditional Transparent Bridging — Tree Topology

Address Learning

In Express Switching mode, the LANplex 2500 system learns the source addresses only of stations not connected to the bridge's backbone port. The bridge listens to the frames transmitted by stations and stores their addresses in a table. The address table is referenced by the filtering/ forwarding function of the bridge.

Packet Filtering and Forwarding

As with traditional bridging, described in Chapter 5, the destination system looks the address of each packet arriving at a bridge port (Ethernet or FDDI) in the address table.

- If an address match is found, then the system forwards packet directly to the destination port.
- If no match is found, the bridge behaves differently from a traditional bridge. For example:
  - A traditional bridge would flood the packet to all ports except the one on which it was received.
  - In Express Switching mode, the bridge never floods unicast or directed packets. A packet received on the backbone port is discarded. A packet received on a nonbackbone port is forwarded *only* to the backbone port.

Using Express Switching

Using Express Switching mode in your network has these advantages:

- Packets with unknown destination addresses are never flooded. Therefore, each nonbackbone segment can be assured of being private to the stations that are directly connected to it and that its bandwidth cannot be intruded upon by extraneous packets.
- There is no limitation on the number of addresses that can be seen on the backbone network. Therefore, the system can support very large backbone networks.

Using Express Switching mode in your network has these contraints:

- Directed (unicast) packets are not forwarded to an end-station that is not connected to the backbone until the bridge has learned that station's address. Therefore, in certain circumstances, a station must "speak before it can be spoken to."
- Source address group information relating to user-defined packet filtering is not available for packets received on the backbone port. For more information on packet filtering, see Chapter 7: *User-defined Packet Filtering*.

7 USER-DEFINED PACKET FILTERING

The LANplex system allows you to add a second layer of packet filtering on top of the standard filtering provided by a traditional transparent bridge.

This chapter contains the following information on user-defined packet filters:

- A general description
- How to use address groups and port groups in packet filters
- Some packet filters
- How to administer user-designed packet filters

About User-defined Packet Filtering

User-defined packet filtering further restricts the packets to be 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 LANplex system is flexible. You can define complex filters by combining many simple comparisons. This flexibility allows you to use packet filters in several helpful ways on your network.

You specify the packet filter using a *packet filter language*. This language consists of operands and operators with which to compose your filters. The language is described in detail in the *LANplex® 2500 Administration Console User Guide*. Table 7-1 describes the two simplest operands in the packet filter language.

Table 7-1 Some Simple Packet Filter Operands

Operand	Description
Constant	A literal value. A constant can be 1, 2, 4, or 6 bytes.
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. Only specify a 6 byte field when you want the filter to examine a 48-bit address.

The operators that you specify in the packet filter allow the filter to make a logical decisions about whether to forward or discard each packet. Table 7-2 describes your choices of operators.

Table 7-2 Packet Filter Operators

Operator	Result
equal	true if operand 1 = operand 2
not equal	true if operand 1 ≠ operand 2
less than	true if operand 1 < operand 2
less than or equal	true if operand 1 ≤ operand 2
greater than	true if operand 1 > operand 2
greater than or equal	true if operand 1 ≥ 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 the packet filter on a specific bridge port. You can install 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 network 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 network 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 7-1Figure 7-1.

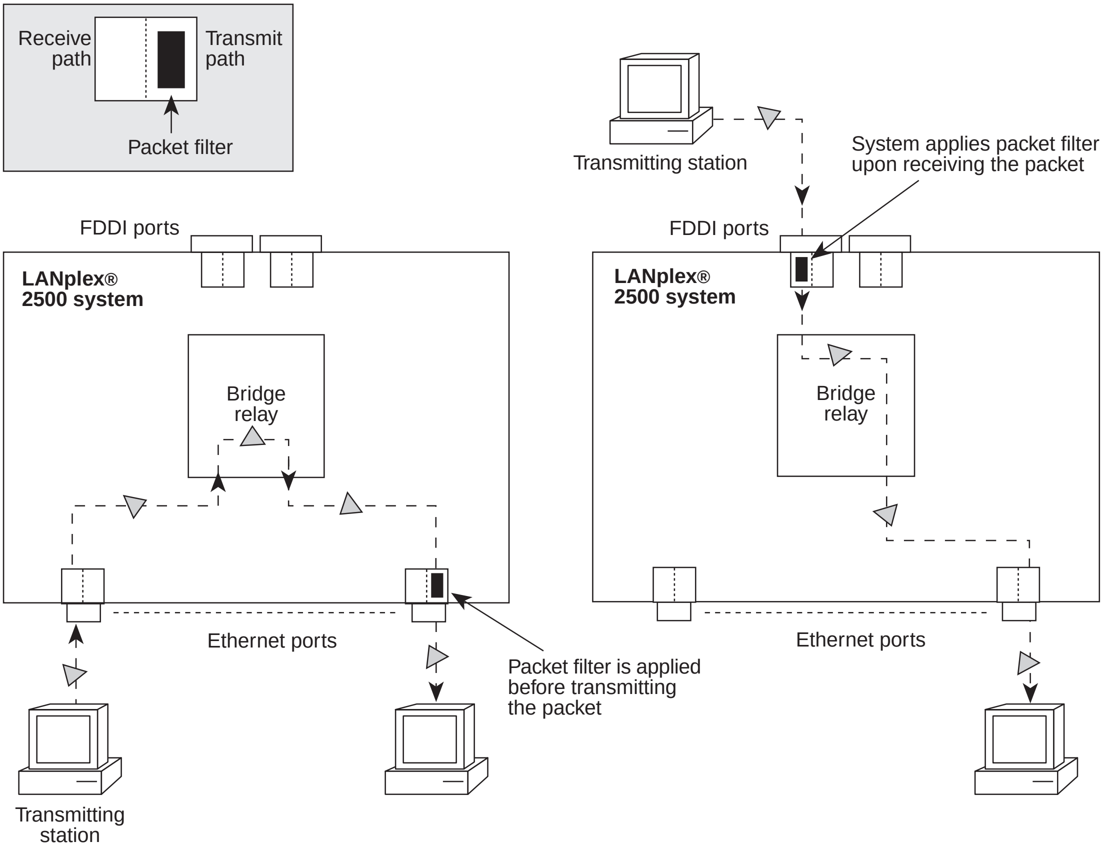

Figure 7-1 Assigning Filters to Paths

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 7-2 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

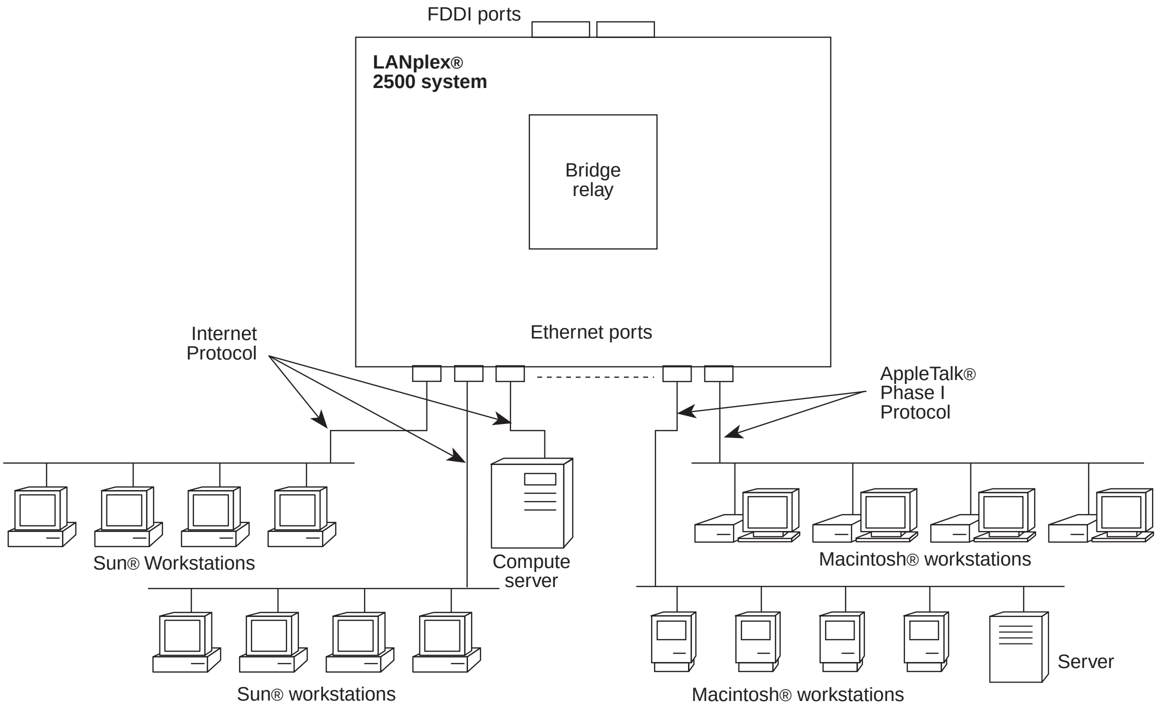

Figure 7-2 A Network with Two Protocols

Macintosh computers can use Internet Protocol (IP) to communicate with Sun workstations, but 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 has not been learned. Because the Sun workstations cannot interpret AppleTalk packets, you want to isolate these packets from the IP segments.

Solution: 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 protocol.
- Operator is *equal.*

The filter is installed as shown in Figure 7-3.

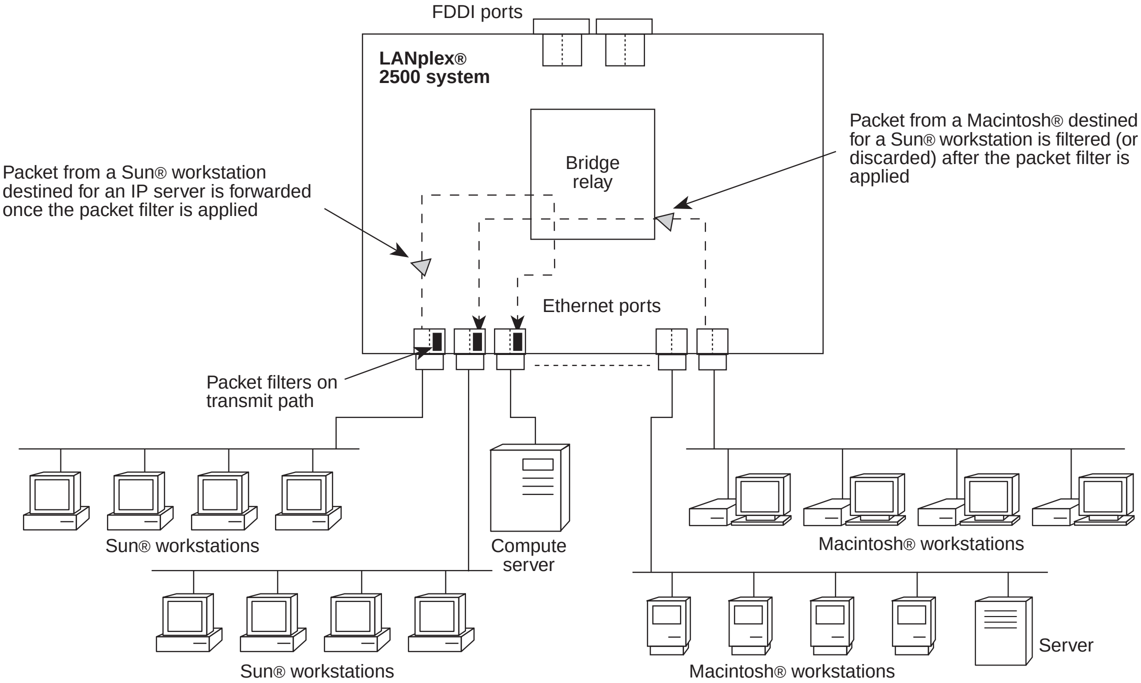
FDDI ports
LANplex®
2500 system
Bridge
relay
Packet from a Macintosh® destined for a Sun® workstation is filtered (or discarded) after the packet filter is applied
Packet from a Sun® workstation destined for an IP server is forwarded once the packet filter is applied
Ethernet ports
Packet filters on transmit path
Sun® workstations
Compute server
Macintosh® workstations
Sun® workstations
Macintosh® workstations
Server

Figure 7-3 Example of AppleTalk® Filter

Example 2: Filtering AppleTalk Phase II Packets

If your Macintosh computers use the AppleTalk Phase II protocol instead of the AppleTalk Phase I protocol (as shown in Example 1), then you must use a filter 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. The filter must first ensure that the packet is an 802.3 packet. If that field is:

- Greater than 1500, then the packet is an Ethernet packet and the value is interpreted as the type field.
- Less than or equal to 1500, then the packet is an 802.3 packet and the value is interpreted as the data length.

In an AppleTalk Phase II packet, a Subnetwork Access Protocol (SNAP) “header” follows the 802.3 header. See Figure 7-4. The filter must verify that the contents of this SNAP field match the AppleTalk packet’s SNAP field.

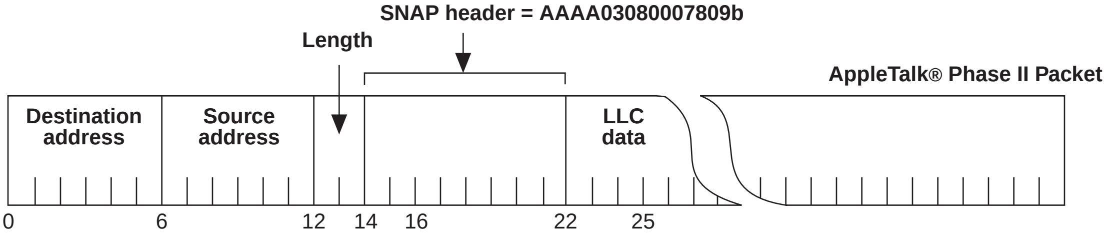

Figure 7-4 AppleTalk Phase II Packet Fields

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. Expressions can be differentiated 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 7-5 illustrates simple expressions form a complete packet filter definition.

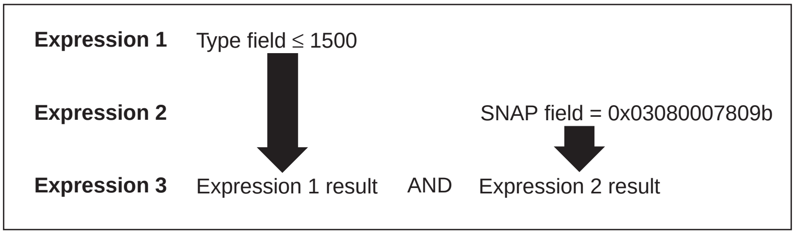

Figure 7-5 Packet Filter Expressions Combined

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 LANplex system 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
- Stations that have access only to a specific network resource
- Stations that have access only to a group of network segments
- Network segments whose attached stations can communicate only with each other

What Is an Address Group?

An address group is a list of MAC addresses. You can configure up to 32 address groups per LANplex 2500 system. You can associate the same address group 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 7-6.

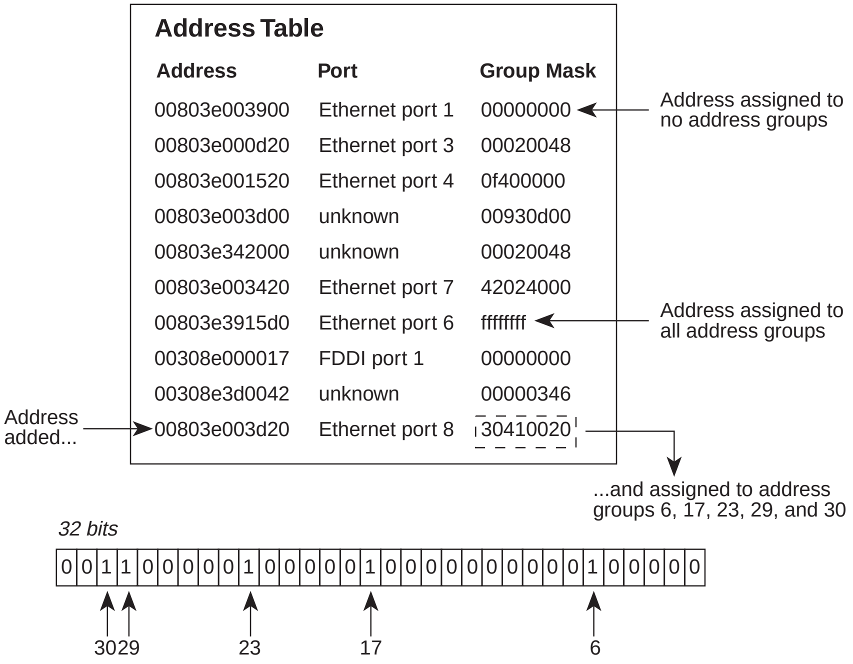

Address Table

Address	Port	Group Mask
00803e003900	Ethernet port 1	00000000
00803e000d20	Ethernet port 3	00020048
00803e001520	Ethernet port 4	0f400000
00803e003d00	unknown	00930d00
00803e342000	unknown	00020048
00803e003420	Ethernet port 7	42024000
00803e3915d0	Ethernet port 6	ffffffff
00308e000017	FDDI port 1	00000000
00308e3d0042	unknown	00000346
00803e003d20	Ethernet port 8	30410020

Address assigned to no address groups
Address assigned to all address groups
Address added...
...and assigned to address groups 6, 17, 23, 29, and 30
32 bits
0 0 1 1 0 0 0 0 0 1 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0
30
29
23
17
6

Figure 7-6 Adding an Address to an Address Group

If the address was never 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.

Because the address table can hold up to 8K address entries, the total number of addresses that can be configured into address groups is at most 8K.

What Is a Port Group?

Port groups are conceptually similar to address groups. A port group is a list of ports. Because the LANplex 2500 system has only 2 bridge ports (16 Ethernet ports and 2 FDDI), a group can have no more than 18 ports. You can configure up to 32 port groups per LANplex 2500 system.

When a port is added to a group, the port is inserted into a port table on the associated system. Each port table entry contains a 32-bit *group mask.* As with address group masks, each bit in the mask specifies *one* of the 32 groups. When the port is added to a group, the corresponding bit in the mask is set. See Figure 7-7

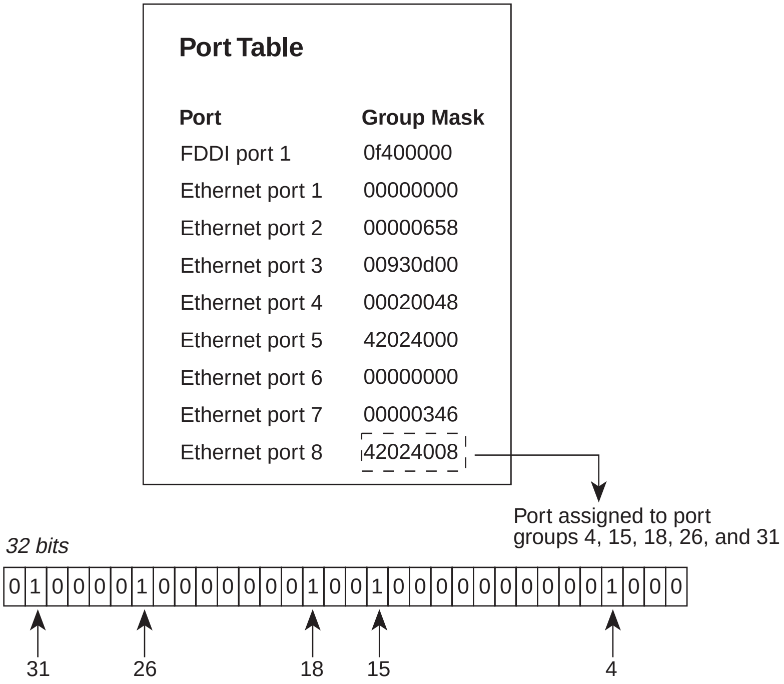

Port Table

Port	Group Mask
FDDI port 1	0f400000
Ethernet port 1	00000000
Ethernet port 2	00000658
Ethernet port 3	00930d00
Ethernet port 4	00020048
Ethernet port 5	42024000
Ethernet port 6	00000000
Ethernet port 7	00000346
Ethernet port 8	42024008

0	1	0	0	0	0	1	0	0	0	0	0	0	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	1	0	0	0

Figure 7-7 Adding a Port to a Port Group

Referencing Address Groups and Port Groups From a Packet Filter

After you configure address and port groups, you can refer to 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 7-3.

Table 7-3 Packet Filter Operands for Address Groups and Port Groups

Operand	Description
Source address group mask	Group mask associated with the source address in the packet
Destination address group mask	Group mask associated with the destination address in the packet
Source port group mask	Group mask associated with the bridge port on which this packet was received
Destination port group mask	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 refers to 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 7-8. 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 A and B 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 7-8 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 **00-01-02-03-04-01**.

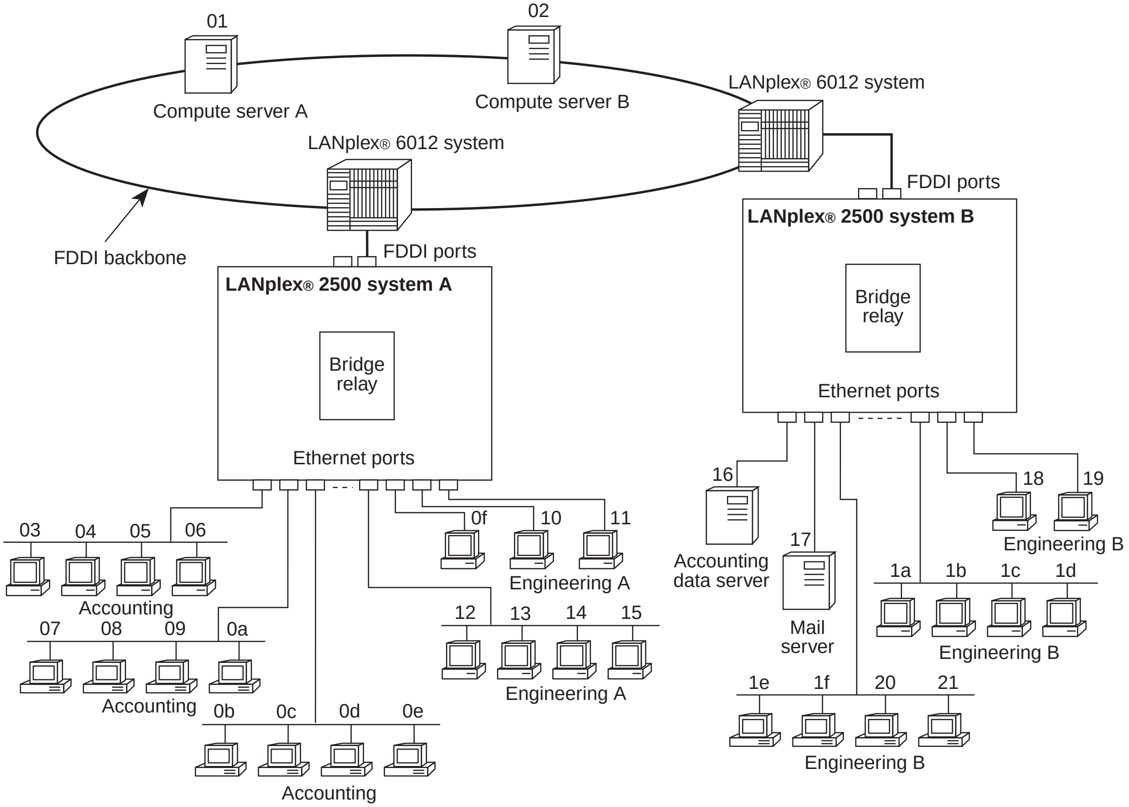

Figure 7-8 Network Needing Filtering to Restrict Server Access

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

- Users in the Accounting group must be able to communicate only with each other and with the Accounting data server.
- Users in Engineering A group must be able to communicate only with each other and Compute server A.
- Users in Engineering B group must be able to communicate only with each other and Compute server B.
- Users in all three groups must be able to communicate only with the Mail server.

To implement packet filtering in the designed scheme 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 MAC address (00-01-02-03-04-17) is included in each group. These address groups yield an address table as shown in Figure 7-9 on page 7-75.

The address table on each system contains the same address listing and group masks, but the port numbers are different numbers. The address table in Figure 7-9 on page 7-75 is for LANplex 2500 system Engineering Group A.

2 After setting up the address groups, 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 7-10. 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.

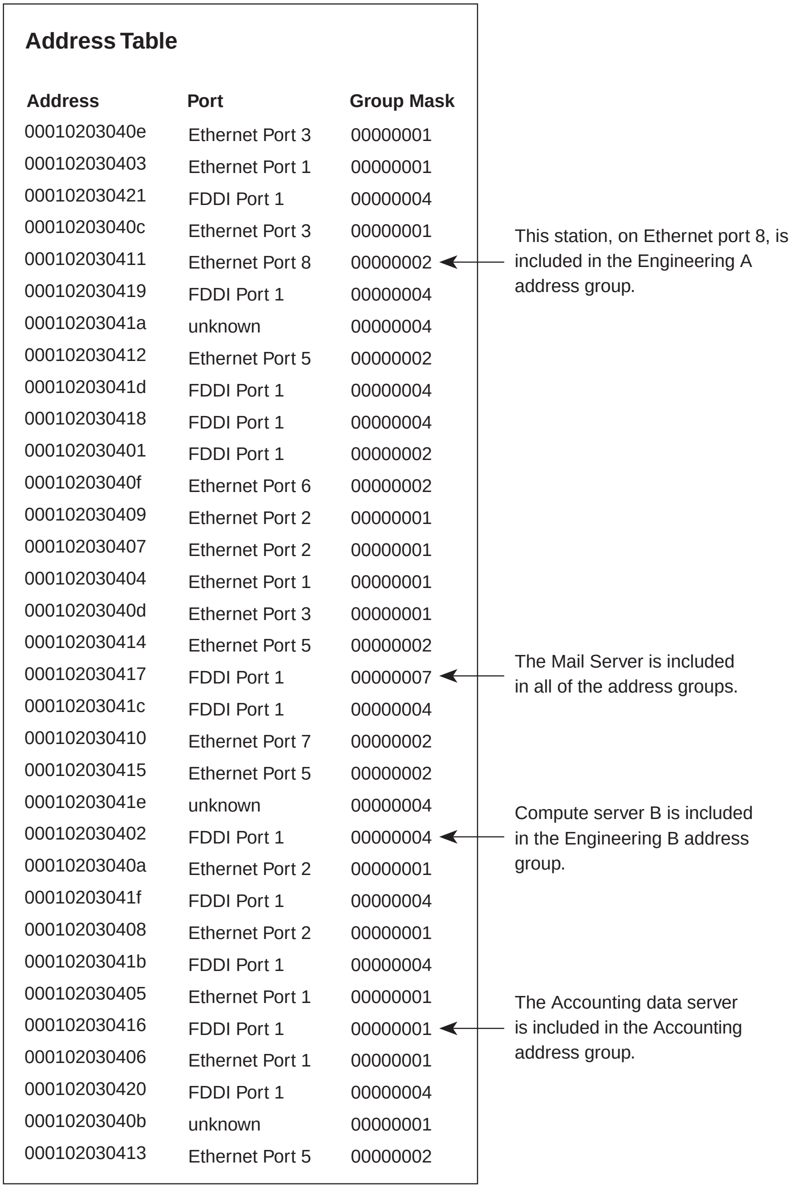

Address Table

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 7-9 Address Table for Restricting Server Access

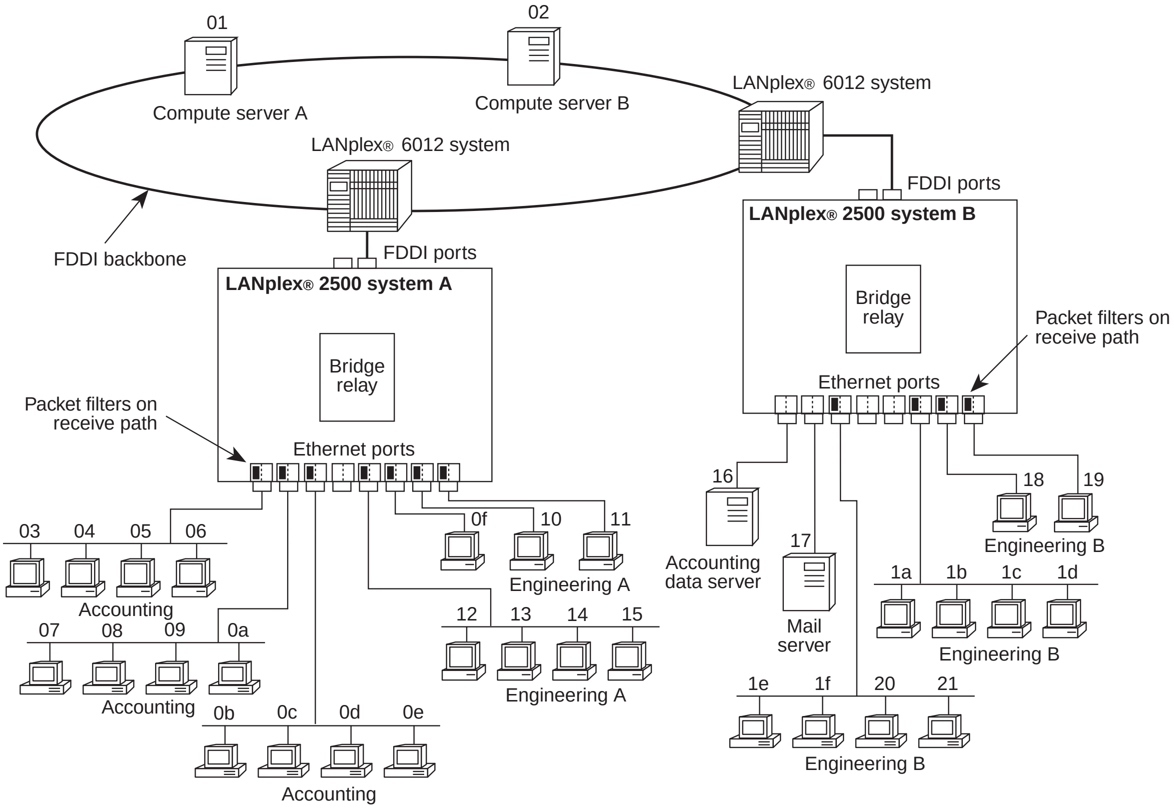

Figure 7-10 Address Group Filtering Example

Globally Administering Packet Filters

You can create packet filters and group definitions in two ways:

- Locally, using the Administration Console.
- On an external computer by creating files that contain the necessary information in the specified format and then loading them onto the LANplex system.

When you create the definitions externally, multiple LANplex systems 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 *LANplex® 2500 Administration Console User Guide*.

8

LANPLEX® BRIDGING EXTENSIONS

This chapter describes LANplex® 2500 bridging extensions, a functionality that enhances bridge 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 communication to 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 stop responding or fail. The LANplex 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.

Figure 8-1 illustrates the threshold mechanism for the multicast packet firewall. For information on setting this threshold, see the *LANplex® 2500 Administration Console User Guide*.

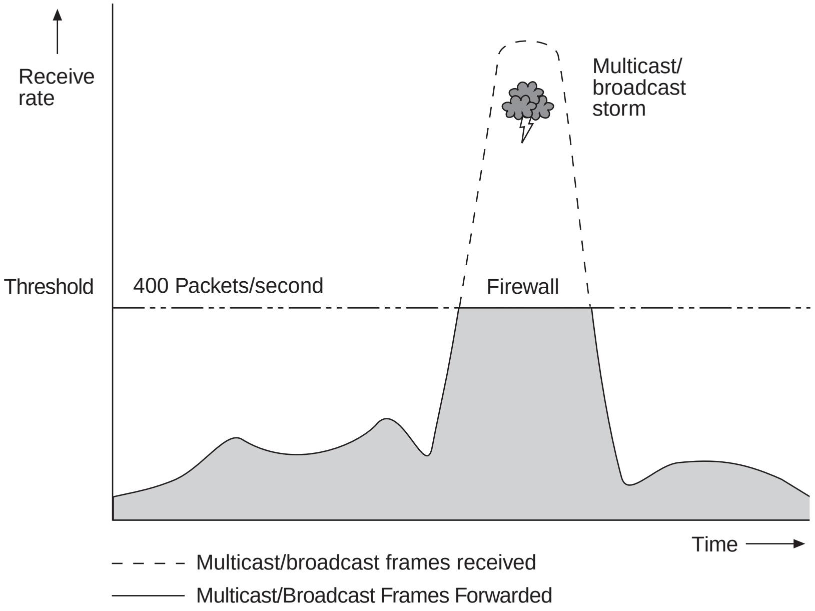

Figure 8-1 Multicast Packet Firewall Threshold Mechanism

IP Fragmentation

With IP fragmentation, FDDI and Ethernet stations connected to a LANplex 2500 system can communicate using IP even if the FDDI stations are transmitting packets that are too large to bridge.

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.

IP fragmentation is specified in RFC 791 (Internet Protocol) and RFC 1122. For information on enabling IP fragmentation, see the *LANplex® 2500 Administration Console User Guide*.

Reducing Packet Flooding

The LANplex 2500 system 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 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 causes increased packet flooding.

The LANplex 2500 system 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 LANplex system detects that a station has moved, it immediately re-assigns the associated port. This enhancement allows you to lengthen the LANplex 2500 system address aging interval, knowing that the LANplex 2500 system rapidly detects most station moves independent of the aging process. This ability results in reduced packet flooding and improved network performance.

For information on setting the address aging timer, see the *LANplex® 2500 Administration Console User Guide*.

Enhanced Network Security

In addition to allowing you to design and use packet filters to improve network security (as described in Chapter 7), the LANplex 2500 system 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. You have permanently tied this address to the specified port unless you manually remove it. 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 more convenient to let the bridge first learn all of the addresses on your network and then to convert these learned, or dynamic, addresses to static addresses to improve your network security.

For information, see the section on bridging in the *LANplex® 2500 Administration Console User Guide*.

III

FDDI TECHNOLOGY

Chapter 9 FDDI Overview and Implementation

Chapter 10 FDDI Networks

9

FDDI Overview and Implementation

This chapter discusses FDDI concepts and terms. The final section shows how FDDI is implemented in the LANplex® 2500 system.

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, 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 2 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.

The industry guideline for FDDI technology 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 always 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 9-1 illustrates the relationship of FDDI entities to the OSI Reference Model. Network attachments communicate with each other using predetermined protocols. The model divides these communication protocols into seven layers, defined so that each layer only requires services from the layer below it.

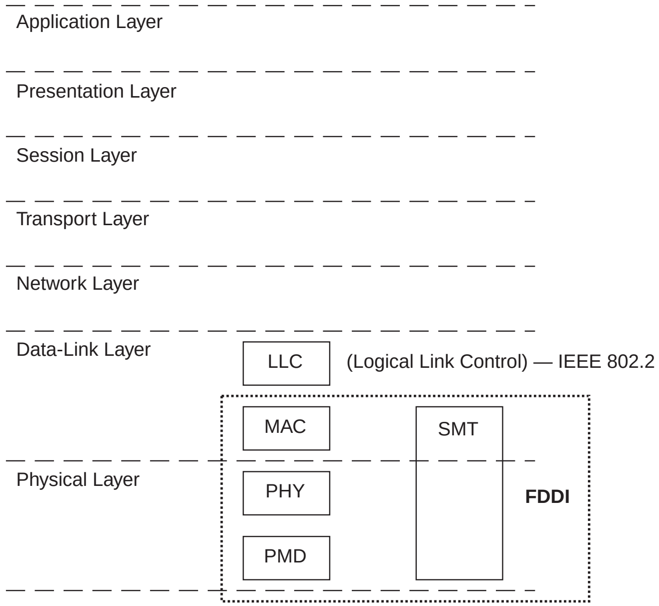

Figure 9-1 FDDI Relationship to OSI Reference Model

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 located at both ends of a physical connection 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 allows such a connection. A connection is rejected if each station has a policy that disallows the connection.

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

- **A 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** — Connects to the incoming fiber of the secondary ring and the outgoing fiber of the primary ring.
- **M port** — Also referred to as Master port; used by a concentrator station to provide connections within a concentrator tree.
- **S port** — Also referred to as Slave port; 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 destinations 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 destined for it by comparing the destination address to its own address. If it sees a match, 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:

- **Primary path** — The segment or segments of the primary ring that pass through a station. Conditions may exist in parts of the network that might 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** — One or more segments of the rings other than the primary ring and secondary ring that pass through the station.

Nodes and Attachments

As we have seen, 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. 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 used 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. The single attachment 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 no A, B, or S ports.

Node Types

Six station and concentrator types are used to describe station configurations and topologies. Figure 9-1 lists these node types and their abbreviations. Figure 9-2 shows how these six node types may connect to an FDDI dual ring.

The LANplex® 2500 system only supports SM-DAS and SAS topologies.

Table 9-1 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

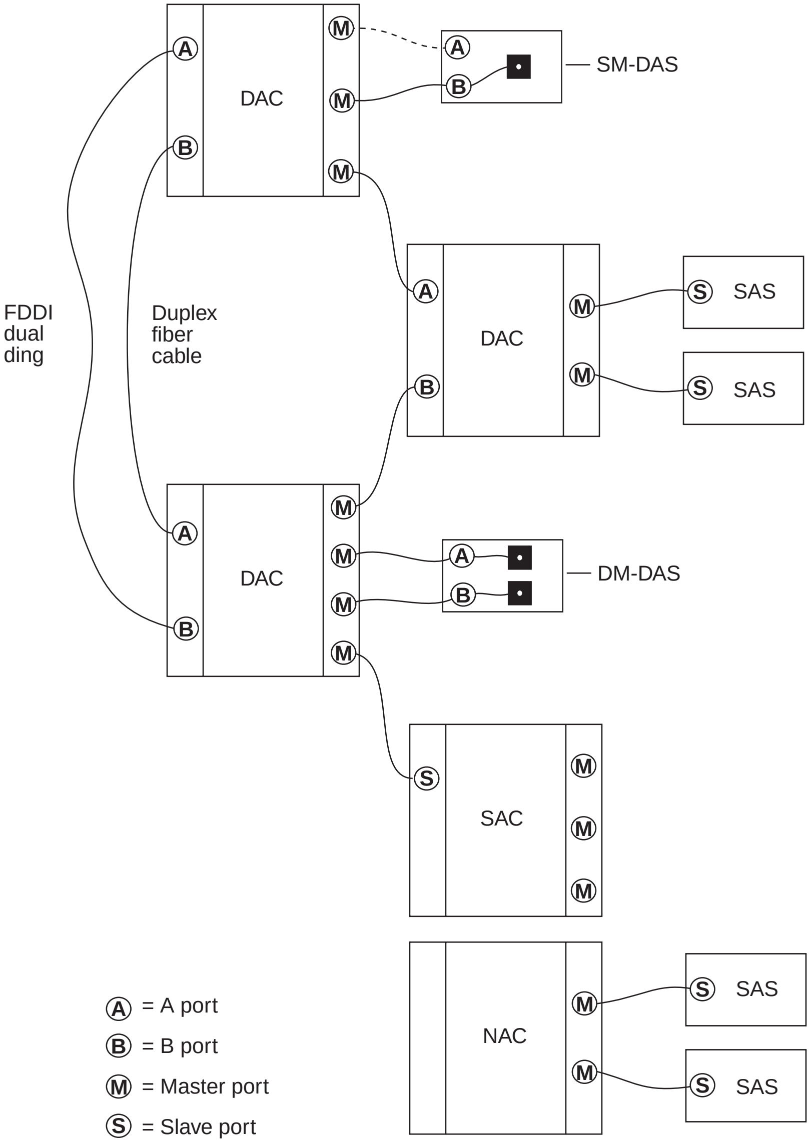

Figure 9-2 Examples of FDDI Configurations

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 falls into three broad categories:

- **Physical Connection Management (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)** — Interconnects PHYs and MACs on paths to achieve proper station configuration and network topology.
- **Ring Management (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 might prevent the ring from operating.

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 into manage stations on the network and to gather information about them. Frame-based protocols:

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

SMT has six key frame-based protocols:

- **Neighbor Notification** — 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** — Performs the remote management of station attributes. It operates on all SMT MIB attributes, attribute groups, and actions.
- **Status Reporting** — Allows a station to notify network managers about events such as station configuration changes and network errors.
- **Status Polling** — Provides a mechanism to obtain station status remotely through a request/response protocol.
- **Echo** — Performs loopback testing on the FDDI dual ring.
- **Synchronous Bandwidth Allocation** — Allocates synchronous bandwidth and monitors both synchronous and total bandwidth.

FDDI and the LANplex 2500 System

Your LANplex 2500 system combines into one system the power of FDDI and:

- Ethernet switching
- Ethernet-to-FDDI transparent bridging
- FDDI switching
- Fast Ethernet switching
- ATM switching.

This combination dramatically enhances LAN performance and increases the capacity of your existing Ethernet network.

With this FDDI 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. Figure 9-3 shows LANplex systems attached to an FDDI dual ring. The connection to the dual ring is made by the A and B ports on the LANplex 2500 system. DASs, excluding concentrators, may be attached to the dual ring, as shown. For additional FDDI configuration examples, see the *LANplex® 2500 Getting Started* guide.

3Com strongly recommends you connect equipment that can be turned on and off, such as workstations, only through concentrators. Connect intermediate systems that are seldom turned off, such as bridges and routers, 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 10: FDDI Networks, which discusses physical and logical topologies, FDDI connection rules, and dual homing.

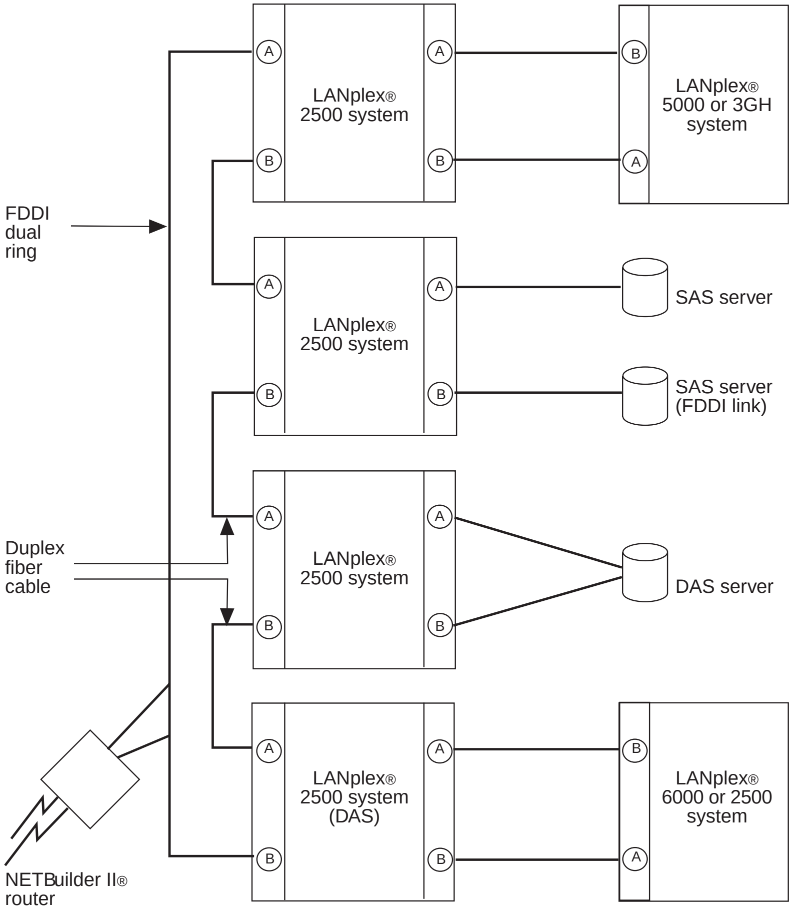

Figure 9-3 Sample FDDI Configuration for the LANplex 2500 system

10

FDDI Networks

This chapter provides general information about FDDI networks and describes the differences between physical and logical topologies. Also covered: FDDI connection rules and dual homing.

About FDDI Networks

FDDI networks have 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. An FDDI network can:

- Meet the networking needs of today's high-performance workstations that produce large quantities of data.
- Offer the speed, distance, and capacity required for the powerful workstations, applications, and super servers of the today.
- Handle 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 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 10-1.

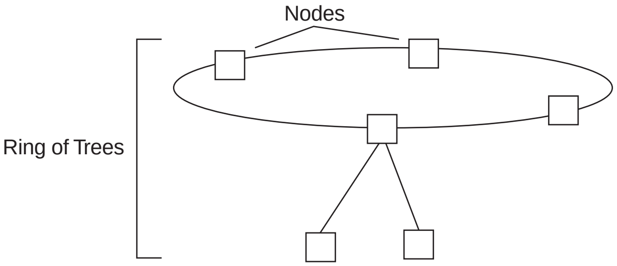

Figure 10-1 Physical Topology

- **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 10-2.

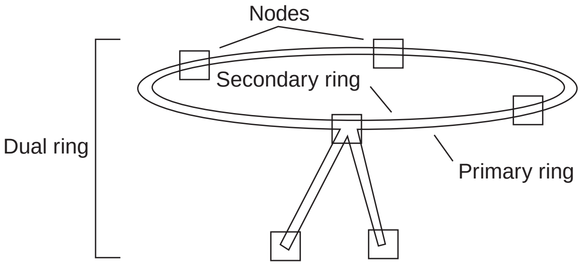

Figure 10-2 Logical Topology

Physical Topology: A Ring of Trees

The FDDI trunk ring consists of dual-attach stations (DASs) and dual-attach connectors (DACs). The DACs on the ring allow you to attach trees. The trees consist of branches of single-attach stations (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 10-3.

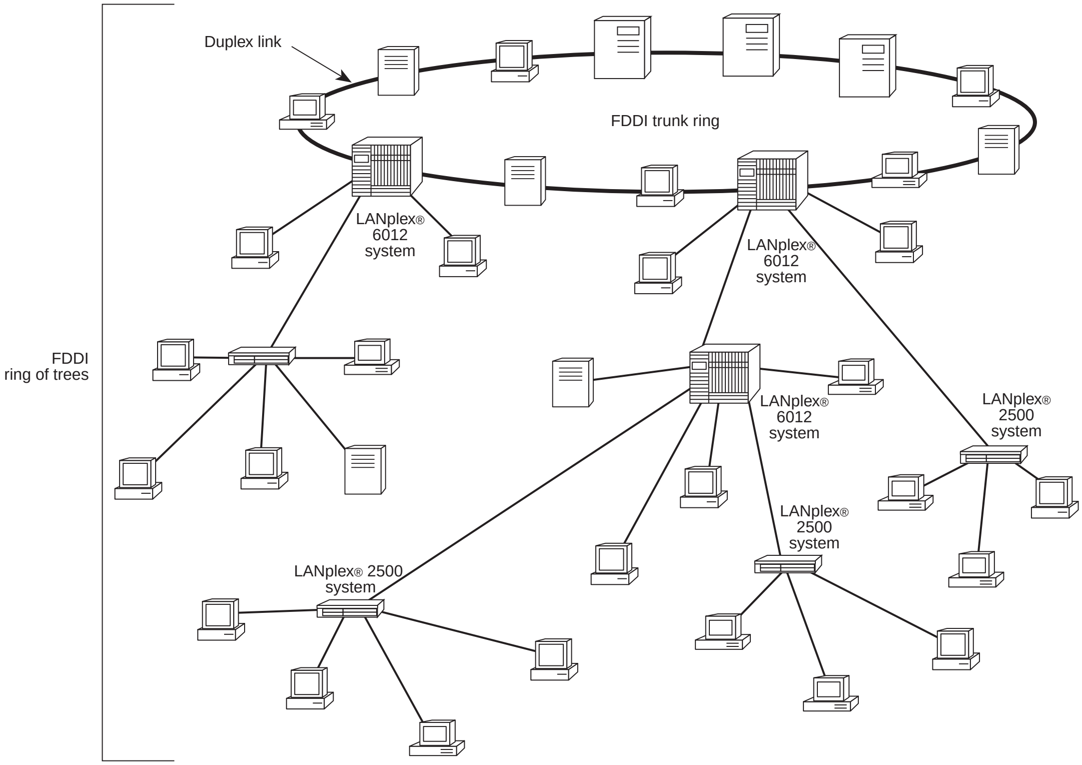

Figure 10-3 Ring of Trees

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. 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

Station Management Protocol (SMT) follows specific connection rules to ensure that only desired physical connection types are included in the LANplex 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 software 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 10-1.

Table 10-1 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) takes 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 does not go to THRU state in CFM. My B port takes 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 S 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 10-2 provides a connection rule matrix summarizing the validity of most types of connections.

Table 10-2 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 (requires notice to SMT)

Dual Homing

When the operation of a dual attachment node is crucial to your network, a configuration called dual homing can provide added reliability. A network administrator using dual homing can determine a station's operation by setting the appropriate configuration policy. The dual-homed station can be configured: with both links active or with one link active and one connection withheld as a backup. It becomes active only if the primary link fails. See Figure 10-4.

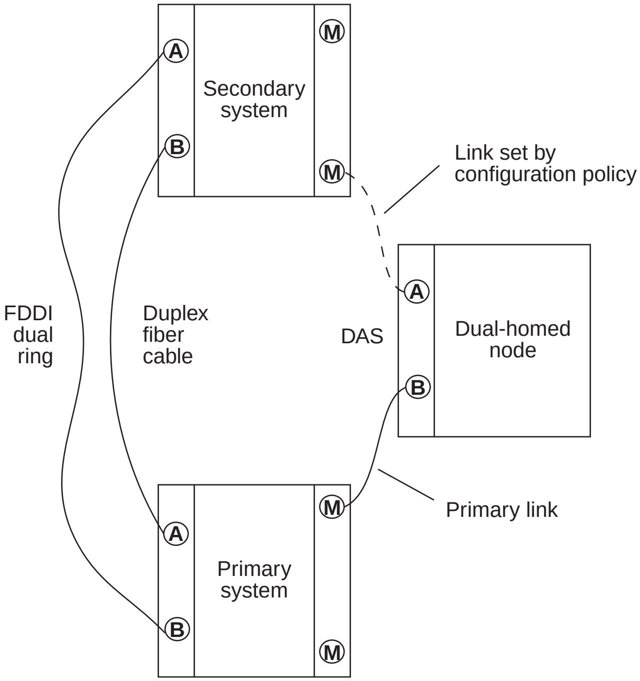

Figure 10-4 Dual Homing

For additional information about dual homing, see the *LANplex® 2500 Getting Started* guide.

IV

ATM Technology

Chapter 11 ATM Networks

11

ATM NETWORKS

This chapter provides general information about:

- The basic concepts of ATM and its architecture
- How ATM handles addressing, signaling and virtual connections
- Traffic management on the ATM network
- The basic concepts of LAN emulation (LANE)
- ATM and the LANplex system

About ATM

In 1986, the Comité Consultatif International Télégraphique et Téléphonique (CCITT), now known as the International Telecommunications Union (ITU), formed a study group to explore the concept of a high-speed, integrated network that could uniformly handle voice, data, and a variety of other services. The result of this study is the Broadband Integrated Services Digital Network (BISDN). BISDN services provide high-speed channels for transmitting digitized voice, data, video, and multimedia traffic. Asynchronous Transfer Mode (ATM) is a switching and multiplexing technology. ATM supports the BISDN services.

ATM is based on the specifications and standards developed by the ITU, the American National Standards Institute (ANSI), and the ATM Forum for transmitting a complete range of user traffic on any User-to-Network-Interface (UNI).

Some benefits of ATM technology:

- **Bandwidth efficiency** — Fixed-length cells allow cell-relay switches to process cells in parallel at high speeds.
- **Application transparency** — The ATM cell size is a compromise between the long frames generated by data communications applications and the

short, repetitive needs of voice transmission. ATM thus allows a free mixture of data and voice or video within the same application.

- **Scalable technology** — ATM accommodates a wide range of transmission rates and applications.
- **Seamless connectivity** — ATM is the first technology that can be deployed in local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WAN).

ATM Basics

ATM architecture differs fundamentally from IEEE 802-style LANs like Ethernet, FDDI, and token ring. LANs provide *connectionless* services, meaning that data is transmitted without the need for a prior connection setup between the sender and receiver. LAN data is transmitted using the Media Access Control (MAC) addresses in each packet to identify end-stations.

ATM is a *connection-oriented* transport service. The device attached to an ATM network must first establish a connection, called a Virtual Connection (VC), with another device attached on the network before information can be transmitted. See "Virtual Connections in ATM" on page 11-109 for more information.

ATM Network Interfaces

ATM standards specify two network interface types for ATM networks:

- The *user-to-network interface* (UNI). Typically, the UNI is the interface between the user and the user's (private) network ATM switch.
- The *network-to-network interface* (NNI). The NNI is an interface between ATM switches or networks.

Figure 11-1 shows ATM network UNIs and NNIs.

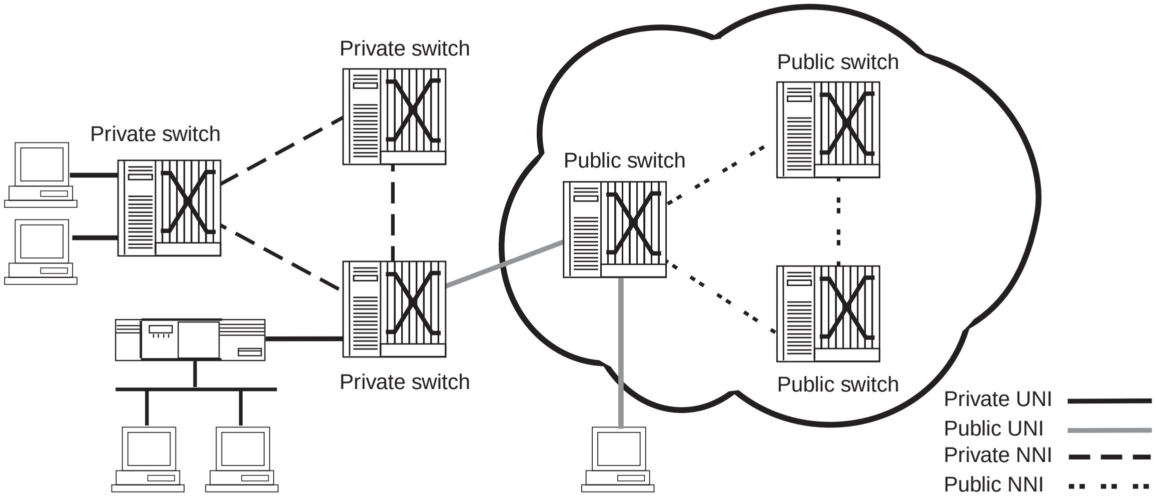

Figure 11-1 ATM Network UNIs and NNIs

ATM Cell Structure

One of the challenges in defining ATM was to determine a structure that could efficiently handle any type of traffic. Because ATM must carry voice, data, and video traffic simultaneously, the structure must accommodate a variety of bit rates and support bursty communications. ATM supports these needs by formatting all information into small fixed-length units called *cells*.

Each ATM cell comprises two parts: a 5-byte header and a 48-byte *information field (payload)*.

Header

Figure 11-2 shows the structure of the UNI ATM cell. The *header* contains the information necessary to deliver the cell to its destination. For networking purposes, only the header is significant. The NNI header eliminates the generic flow control (GFC) field, extending the virtual path identifier (VPI) field. This allows the NNI to support a much greater number of virtual paths.

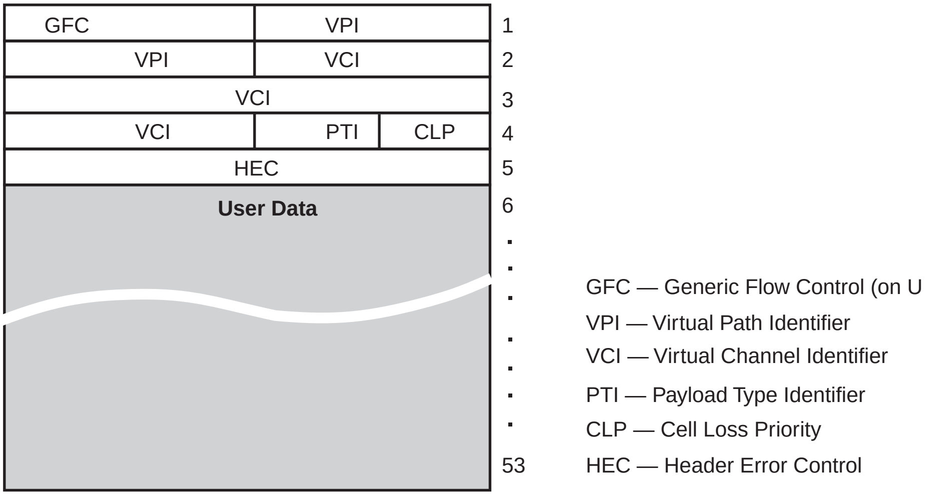
GFC
VPI
1
VPI
VCI
2
VCI
3
VCI
PTI
CLP
4
HEC
5
User Data
6
53
GFC — Generic Flow Control (on U
VPI — Virtual Path Identifier
VCI — Virtual Channel Identifier
PTI — Payload Type Identifier
CLP — Cell Loss Priority
HEC — Header Error Control

Figure 11-2 UNI ATM Cell Structure

These fields make up an ATM cell:

- **Generic Flow Control (GFC)**— Controls the flow of traffic across the user-network interface (UNI) and into the network. The mechanisms for using this field are still under development.
- **Virtual Path Identifier (VPI)** and **Virtual Channel Identifier (VCI)** — The VPI and VCI are addressing identifiers used to route cell traffic. Because of their routing significance, the VPI and VCI addressing identifiers are described further in the next section.
- **Payload Type Indicator (PTI)** — A 3-bit field that indicates if the cell contains user information or connection-associated layer management information.
- **Cell Loss Priority (CLP)** — Indicates that, under congested conditions, cells with this bit set are discarded by the network before cells with the CLP bit clear.
- **Header Error Control (HEC)** — Used by the physical layer for detection of bit errors in the cell header.

Following the HEC field is the 48-byte cell information field containing the user data.

Virtual Connections in ATM

ATM switching is performed at the ATM layer by defining virtual connections. Virtual connections are communication channels that provide for sequential, unidirectional transport of ATM cells. Multiple virtual connections can exist on a physical link.

A connection is identified by a circuit identifier, called a *virtual channel identifier* (VCI), to exchange data between two ATM stations over a previously established *virtual channel connection* (VCC).

Two levels of virtual connections are supported at the UNI:

- A VCC consists of a single connection established between two end-points. VCC switching uses both the VCI and the VPI in the cell header.
- A VPC consists of a bundle of VCCs carried transparently between two end-points. VPC switching uses only the VPI in the cell header.

Connections may be point-to-point or point-to-multipoint.

- A *point-to-point* connection simply connects two endpoints.
- A *point-to-multipoint* connection is a unidirectional connection between two or more endpoints. One ATM endpoint is designated as the *root* and serves as the source node in a tree topology. The other endpoints are *leaf* nodes. When the root receives information, it sends copies to all of the leaf nodes on the tree. Leaf nodes cannot communicate directly with each other or send information to the root. The initial connection is established through the root node and one leaf node, and then additional nodes are added one leaf at a time.

Virtual Circuits

Connections may be established on demand as switched virtual circuits (SVCs) or as pre-established permanent virtual circuits (PVC).

Switched Virtual Circuits

SVCs are established using UNI signaling ITU (Q.2931), which assigns the VPI and VCI to each link. Connections are established and released as needed. These connections remain open for an arbitrary amount of time but may not necessarily be automatically reestablished after a network failure. Both point-to-point and point-to-multipoint connections can be established.

Permanent Virtual Circuits

PVCs are established via network management and are left up indefinitely. Both point-to-point and point-to-multipoint connections can be established.

Virtual Paths and Virtual Channels

As in traditional LAN packets, the header of each ATM cell contains addressing information. However, rather than a specific destination address, each cell header contains two fields, the VPI and the VCI, that specify the virtual connection (VC) over which the cell should be forwarded. See Figure 11-3.

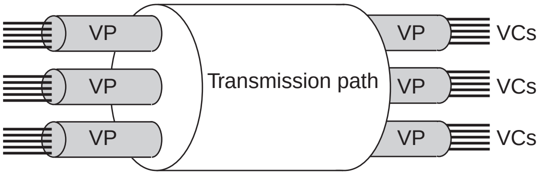

Figure 11-3 Virtual Channels and Virtual Paths

The VPI and VCI establish only the connections between two ATM entities, not the end-to-end connection of the user to the switch.

A virtual channel is a connection between two or more end-points. All communications proceed along the same VC, which preserves cell sequence and provides a specific quality of service.

Figure 11-4 shows the relationship among virtual paths (VPs), virtual channels (VCs), and virtual channel connections (VCCs).

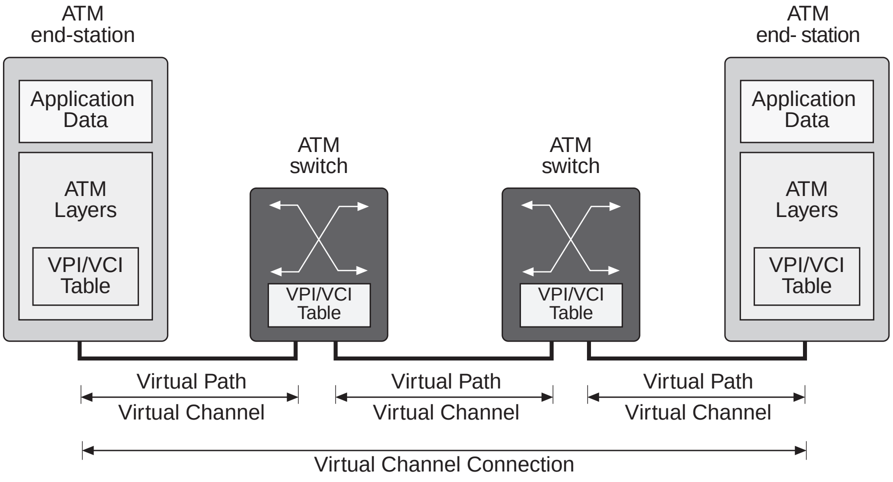

Figure 11-4 Virtual Connections

ATM Protocol Architecture

ATM layers do not map directly to the Open Systems Interconnect (OSI) Reference Model of traditional LAN networks. ATM layers follow the BISDN reference model, shown in Figure 11-5.

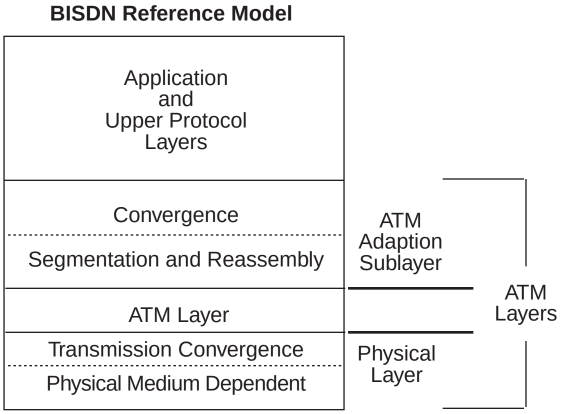

Figure 11-5 BISDN Protocol Reference Model

The primary layers of the BISDN reference model are:

- The Physical Layer
- The ATM Layer
- The ATM adaptation Layer

The Physical Layer

The Physical Layer defines how cells are transported over the network. This are of influence includes physical interfaces, media, and information rates.

The Physical Layer also defines how cells are converted to a line signal depending on the media type. ATM is media independent in that it is not tied to any particular physical layer. The ATM Forum UNI Specification Version 3.0 provides specifications for physical layer interfaces for both the public and private user-network interfaces.

SONET Physical Layer Interface

Synchronous Optional Network (SONET) is the physical layer most often associated with ATM. SONET provides, through a framing structure, the payload envelope necessary for the transport of ATM cells.

The LANplex system uses SONET OC-3 as the Physical Layer interface for the ATM network.

The ATM Layer

The ATM Layer provides a single mode of transport for many types of telecommunications services. Except for the quality of service (QOS) requested for the virtual circuit, the ATM Layer is unaware of the type of information (voice, video, or data) that it is carrying.

Quality of Service

Traffic management is concerned with ensuring that users get their desired quality of service. The ATM layer Quality of Service (QOS) is defined by a set of parameters such as *delay*, *delay variation*, and *cell loss ratio*. While setting up a connection on an ATM network, specific rate parameters related to the desired quality of service can be requested.

The applications that are transported by the ATM Layer are divided into four service classes: constant bit rate (CBR), variable bit rate (VBR), available bit rate (ABR), and unspecified bit rate (UBR).

The LANplex supports unspecified bit rate (UBR) only.

The ATM Adaptation Layer (AAL)

The ATM adaptation Layer (AAL) provides the user-oriented functions that are not included as part of the ATM Layer. These user-oriented functions allow the ATM Layer to support the transport of different types of higher layer protocols and services. In other words, the AAL adapts the pure transport function of the ATM Layer to meet the requirements of different service users. Figure 11-6 shows how the AAL maps to the OSI Reference Model.

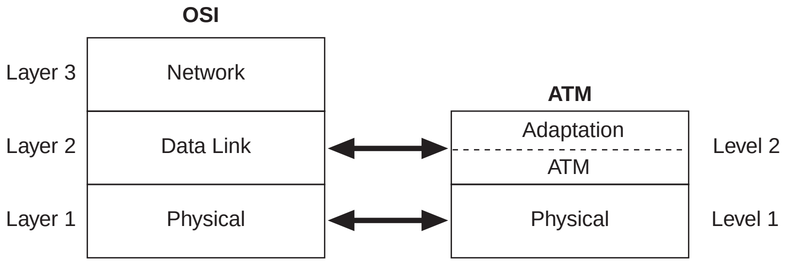

Figure 11-6 ATM adaptation Layer Mapped to the OSI Model

There are currently four types of AALs: AAL 1, AAL 2, AAL 3/4, and AAL5.

The LANplex uses AAL5 only.

AAL5 simplifies the segmentation and reassembly (SAR) portion of the Adaptation Layer to pack all 48-byes of the cell information field with data. This AAL makes ATM look like high speed frame relay. It also assumes that only one message is crossing the UNI at a time. That is, multiple end users at one location cannot interleave messages on the same virtual circuit, but must queue them for sequential transmission.

AAL Sublayers

The major function of the AAL is to provide an interface between user data and the ATM network. To do this, the AAL is subdivided into two sublayers:

- **The Convergence Sublayer (CS)** — The CS is service dependent. The CS detects or corrects bit errors and lost or misinserted cells, and it maintains the timing relationship between the source and destination stations.
- **The Segmentation and Reassembly (SAR)** Sublayer — When transmitting data, the SAR segments the higher-layer protocol data units (PDUs) into 48-byte units for placement in the information field of an ATM cell. When

receiving data, the SAR reassembles the contents of the ATM cell information fields into the higher-layer protocol data units.

Interim Local Management Interface (ILMI) Communication Protocol

The ILMI communication protocol is an open management protocol that supports the bi-directional exchange of management information between all end-stations and the switches to which they are connected.

The ILMI functions for a UNI provide status, configuration, statistics, and control information about virtual path and virtual channel connections available at the UNI. In addition, the ILMI provides for address registration across the UNI. ILMI communications use SNMP directly over AAL5.

ILMI Management Information Base (MIB)

Managed objects are accessed through the Management Information Base (MIB). The ILMI MIB provides the management application with the capability to control and monitor the ATM link and physical layer.

LANplex Switching software version 8.1.1 supports all items identified for the user-side device in a private UNI as defined in the UNI 3.0 specification, plus additional objects listed in UNI version 3.1 for the ILMI MIB.

ATM Address Registration

A User-to-Network Interface (UNI) Management Entity (UME) implements the management interface to the ATM network. Each ATM port has one UME, which manages the network prefix and address tables and provides access to the Interim Local Management Interface (ILMI) MIB.

To establish an ATM connection at the UNI, both the user and the network must know the ATM addresses that are in effect at that UNI. These ATM addresses are used in signaling messages to establish connections with ATM end-stations. The address registration procedure provides the means for the dynamic exchange of addressing information between the user and the network at the UNI.

To exchange address information, two MIB tables are defined. One contains network prefixes of the ATM switches, while the other contains registered ATM addresses for each end-station. The user side implements the network prefix table, and the network side creates and deletes entries in the table in order to register and de-register ATM addresses.

Address registration is performed in two phases:

1. The network side supplies the network portion of the address (network prefix).

2. The user side appends the end-system identifier (ESI) of the address and registers the complete address with the network side.

   Either side can dynamically register and de-register their respective parts of the address. This registration process provides the network with the intelligence it needs to automatically associate a virtual circuit with given source and destination addresses.

Integrating ATM in Legacy LANs

In a network that incorporates both ATM and existing LAN technologies, it is important to understand the significance of LAN-to-ATM conversion functions. Two methods that can be used to interconnect LANs over an ATM network are:

- The first method, *Classical IP over ATM*, implements IP "natively" over ATM, which enables direct communication among LANs operating under the same protocol. Classical IP uses logical IP subnets (LISs) to route between LAN-based and ATM-based subnets for IP routing connectivity. For more information on Classical IP over ATM, see Chapter 4 in the *LANplex® 2500 Extended Switching Guide* or Chapter 9 in the *LANplex® 2500 Administration Console Guide*.

Classical IP over ATM is available only in LANplex Extended Switching software.

- The second method, *LAN Emulation (LANE)*, interconnects LANs over an ATM network through bridging. LANE makes a connection-oriented ATM network look and behave like a shared connectionless LAN segment. For more information on LANE, see the next section.

Figure 11-7 shows how ATM operates in a LAN network.

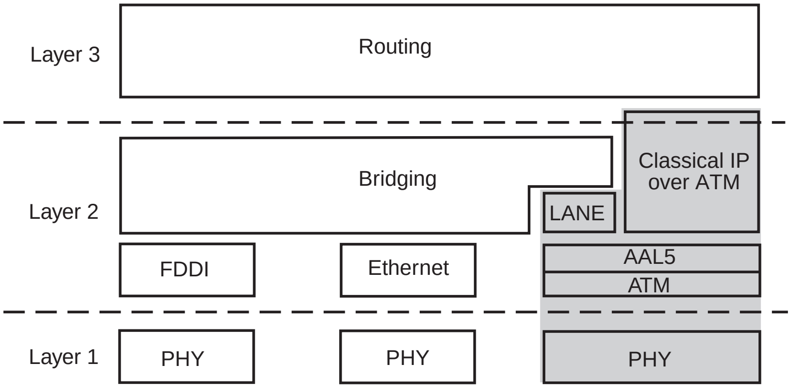

Figure 11-7 ATM in a LAN network

LAN Emulation (LANE)

LAN Emulation (LANE) enables existing applications to access an ATM network via protocol stacks as if they were operating over traditional LANs. Because LANE is implemented in edge devices and systems, it is transparent to the ATM network and to legacy LAN devices. In addition, it maps the MAC address-based networking protocols into ATM virtual connections.

To provide a connectionless MAC service on top of ATM, a LAN Emulation sublayer is placed on top of the ATM adaptation layer (AAL) and emulates the LAN service by making the switched point-to-point ATM network appear to the 802.2 logical link control (LLC) as an IEEE 802.x shared media LAN.

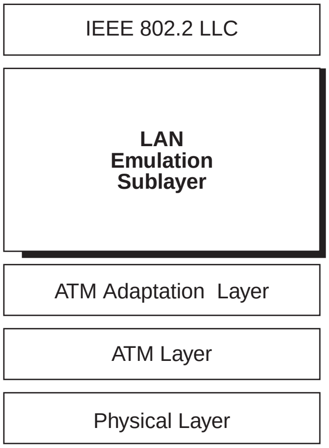

Figure 11-8 LAN Emulation Sublayer

In a traditional IEEE 802.x LAN, traffic is transmitted to all stations on the shared physical medium, with each station determining which packets it should receive and which packets should be discarded.

In an ATM network, each LAN segment can be emulated using the client/server model by allowing a selected group of ATM nodes (clients) to join a LAN emulation service (server). Each LAN is composed of all nodes that have joined a particular LAN emulation service. An end system is configured as a member of a single LAN emulation service, while an intermediate system, such as a router, belongs to multiple LAN emulation services and provides connectivity between each of the emulated LANs.

Figure 11-9 shows an ATM network is composed of three separate ATM LAN segments: A, B, and C. Each of a particular LAN segment's members is a member of the same LAN emulation service.

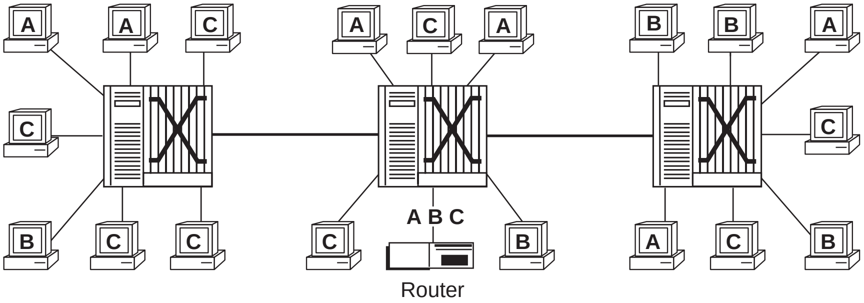

Figure 11-9 ATM Network Using LAN Emulation

Three important concepts are illustrated in Figure 11-9:

- Network nodes may be members of the same LAN emulation service (ATM LAN segment) even if they are connected to different switches in the network, as long as the switches are interconnected.
- Traffic destined for end nodes on the same LAN segment — for instance, traffic between two native ATM nodes on segment C — is directed by the LAN emulation service.
- Traffic destined for end nodes on different LAN segments must be bridged or routed. This means that each LAN emulation service is required to forward traffic between nodes on segments A and B to the router.

Membership in an ATM LAN segment is defined by the logical membership in a LAN emulation service rather than by a physical connection to a LAN segment.

LANE Components

Each emulated LAN contains of:

- One or more LAN Emulation Clients (LECs)
- A single LAN Emulation Service, consisting of
  - One LAN Emulation Server (LES)
  - One Broadcast and Unknown Server (BUS)

In addition, there is also a LAN Emulation Configuration Server (LECS) that services all of the emulated LANs (ELANs) that exist in the ATM network.

- **LAN Emulation Client (LEC)** — A set of functions implemented in an ATM endpoint which serves as an interface between it and the ATM network in support of LAN Emulation.
- **LAN Emulation Server (LES)** — The set of functions implemented in the ATM network support of LAN Emulation. The LES provides address resolution for the LECs within an Emulated LAN.
- **Broadcast and Unknown Server (BUS)** — The set of functions implemented in an ATM network that provides LAN-to-LAN transmission support while a LAN connection is being established. The BUS also supports Ethernet Broadcasts Mode by sending broadcast data to all LECs.
- **LAN Emulation Configuration Server (LECS)** — Assigns individual LAN Emulation clients to different emulated LANs. When a LEC initializes, it establishes a connection to the LECS. Based on its own programming, configuration database, and information provided by clients, the LECS assigns any client requesting configuration information to a specific Emulated LAN service by giving the client the LES's ATM address. This arrangement allows a client to be assigned to an emulation LAN based on either the physical location (ATM address) or the identity of a LAN destination which it is representing (ELAN name).

Virtual Channel Connections (VCC)

Communication among LECs and between LECs and the LES is performed over ATM virtual channel connections (VCCs). Each LEC must communicate with the LES over control and data VCCs. Emulated LANs operate in one of three environments: switched virtual circuit (SVC), permanent virtual circuit (PVC), or mixed SVC/PVC.

Emulated LANs (ELANs)

Using LAN emulation, ATM allows related users in separate physical segments to be effectively grouped into a common broadcast domain called an emulated LAN. Each emulated LAN is independent of the others and there is no direct communication across emulated LAN boundaries.

ATM and the LANplex System

Your LANplex 2500 system brings you the power of ATM by combining Ethernet switching, Ethernet-to-FDDI bridging, Fast Ethernet switching, FDDI switching, and ATM switching in an integrated system.

The LANplex uses two methods to adapt the existing network layer protocols of legacy LANs to the connection-oriented paradigm of ATM:

- *Classical IP over ATM* supports transparent translation of IP routing only over ATM and does not support broadcast or multicast addressing. Classical IP is supported in LANplex Extended Switching software only. For information on using Classical IP over ATM, see the *LANplex® 2500 Extended Switching Guide*.
- *LAN Emulation* supports transparent translation of higher level protocols, such as IP, IPX, and AppleTalk. LANE also supports broadcast and multicast addressing.

Configuring ATM

Before configuring ATM in your LANplex 2500 system, you should take these steps to confirm that your ATM connection is established:

- Check the ATM link status.
- Verify LANplex address registration is operational.
- Verify that signalling is operational.

For information on verifying your ATM connection and configuring ATM in the LANplex 2500 system, see *Chapter 9: Administering ATM* in the *LANplex® 2500 Administration Console User Guide*.

LAN Emulation

You can use LAN Emulation (LANE) in the LANplex 2500 system to provide unicast, multicast, and broadcast network behavior over connection-oriented ATM. LANE allows you to group physical segments into a common broadcast domain called an emulated LAN (ELAN). Each emulated LAN is independent of the others, and there is no direct communication across emulated LAN boundaries.

ELANs may be part of a virtual LAN (VLAN), spanning an entire network of devices connected by Ethernet, FDDI, and ATM.

In addition, LANplex 2500 Extended Switching software allows ELANs to be included in layer 3 protocol-sensitive VLANs. For information on configuring

VLANs, see *Chapter 2: VLANs in the LANplex System* in the *LANplex 2500 Extended Switching User Guide*.

Creating an Emulated LAN

You can create an 802.3 emulated LAN on ATM by defining the LAN's servers (the BUS and the LES) and then configuring each LAN Emulated Client (LEC). To create an emulated LAN, follow these steps:

1 Determine the location of the LES and BUS.

You can define the LES and BUS in any ATM end-point device that supports LES and BUS configuration, or on an ATM switch, such as 3Com's CELLplex™ 7000 system.

You can configure the LES and BUS on a LANplex system using LANplex Extended Software only.

2 Define the Broadcast and Unknown Server.

3 Define the LAN Emulation Server

For information on configuring ELANs in your LANplex 2500 system, See *Chapter 9: Administering ATM* in the *LANplex® 2500 Administration Console User Guide*.

Bridge Loops in ELAN Configurations

When configuring more than one LEC in a LANplex 2500 system, the potential exists for bridge loops, which may result in broadcast storms. This possibility exists because the LANplex can bridge between all ports, including logical bridge ports (LECs).

In the example illustrated in Figure 11-10:

- 2 LANplex 2500 systems are configured in the same ATM network, with 2 LECs defined on each system.
- 2 ELANs, called Marketing and Engineering, are configured in the ATM network, and both LANplex 2500 systems have connections to each of these ELANs.

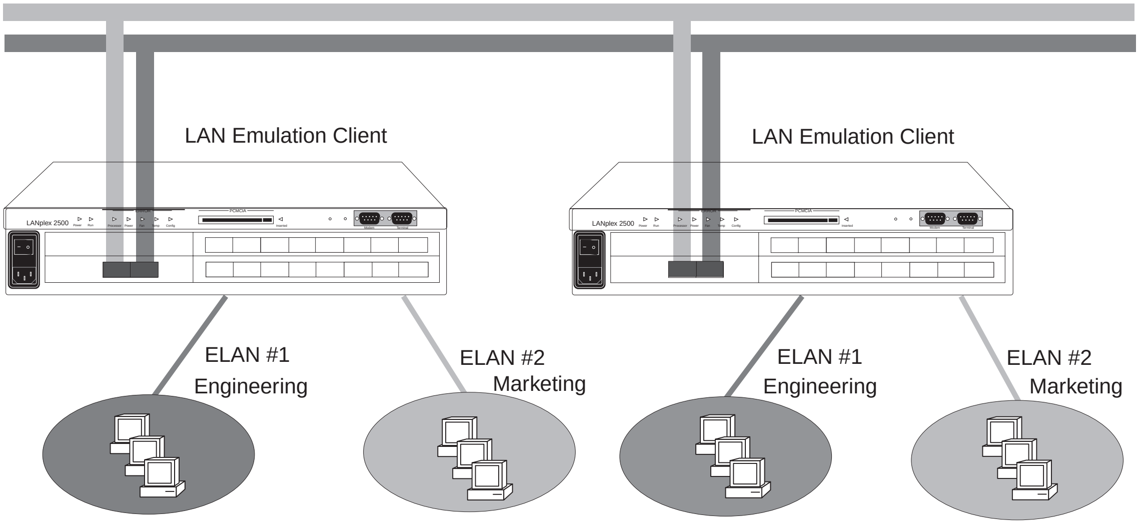

Figure 11-10 Multiple ELANs Configured on LANplex 2500 Systems

In this example, two bridge ports are connected to the same location, causing a bridge loop. A similar example would be the connection of two LANplex 2500 systems by means of two Ethernet segments.

Enabling Spanning Tree (STP) would block one of the ELANs, which may not be the desired behavior for your configuration. The procedure for eliminating bridge loops in a ELAN configuration differs depending on the software type installed. For information on eliminating bridge loops in an ELAN configuration using either LANplex Intelligent Switching software or LANplex Extended Switching software, see Chapter 9: *Administering ATM* in the *LANplex® 2500 Administration Console User Guide*.

For information on Spanning Tree, see the section on bridging in the *LANplex® 2500 Administration Console User Guide*.

V

APPENDIXES

Appendix A SNMP MIB Support

Appendix B Technical Support

A

SNMP MIB SUPPORT

This appendix lists the SNMP MIBs supported by the LANplex system software and describes the supported SNMP compilers.

SNMP MIBs

SNMP MIB files are shipped with the LANplex 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 that may occur from release to release are documented in the Installation and Release Notes.

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

- **atm.mib** — ATM MIB, RFC 1695

  The following ATM MIB attributes are not supported:

  atmInterfaceConfTable

  - atmInterfaceAdminAddress
  - atmInterfaceMyNeighborIpAddress
  - atmInterfaceMyNeighborIfName

  atmVplTable

  - atmVplCrossConnectIdentifier

  atmInterfaceDs3PlcpTable

  atmVpCrossConnectIndexNext

  atmVpCrossConnectTable

  atmVcCrossConnectIndexNext

  atmVcCrossConnectTable

  aal5VccTable

- **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

- **if.mib** — If MIB, RFC 1573
- **fddiSmt7.mib** — FDDI SMT 7.3 MIB, RFC 1512
- **lec.mib** — LEC MIB, af-lane-0044.00

  The following LEC MIB attributes are not supported:

  lecStatusTable

  - lecLastFailureRespCode
  - lecLastFailureState

  lecServerVccTable

  lecRouteDescrTable

  lecArpTable

  lecRDArpTable

- **les.mib** — ATM MIB, af-lane-1129.001

  The following LES MIB attributes are not supported:

  lesconfGroup

  - lesAtmAddrSpec
  - lesAtmAddrMask

  lesStatGroup

  lesLecStatGroup

  lesFaultGroup

- **lp.mib** — LANplex Systems MIB, version 1.3.0

  The following LANplex Systems MIB trap is not supported:

  lpsSystemFanFailure

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

  The following LANplex Optional FDDI MIB attributes are not supported:

  lpOptMAC Group

  - lpOptMACPriTable

  lpOptPATH Group

  - lpOptPATHSbaTable

- **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
  - 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 *lpOpFddi.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"
```

  These warning messages have no effect on the ability of SNM to use the schema file generated with SNM version 2.0 or later.

B

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 24 hours a day, seven days a week, through the following online systems:

- 3Com Bulletin Board Service (3ComBBS)
- World Wide Web site
- 3ComForum on CompuServe® online services
- 3ComFacts^{SM} 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 24 hours a day, seven days a week.

Access by Analog 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 Digital Modem

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® Online Service

3ComForum, based on CompuServe Online Service, contains 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^{SM} 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 at these international access numbers:

Country	Telephone Number
Hong Kong	(852) 2537 5610
U.K.	(44) (1442) 278279
U.S.	(1) (408) 727 7021

Local access numbers are available within the following countries:

Country	Telephone Number	Country	Telephone Number
Australia	800 123853	Netherlands	06 0228049
Belgium	0800 71279	Norway	800 11062
Denmark	800 17319	Portugal	0505 442607
Finland	98 001 4444	Russia (Moscow only)	956 0815
France	05 90 81 58	Spain	900 964445
Germany	0130 8180 63	Sweden	020 792954
Italy	1678 99085	U.K.	0800 626403

Support from Your Network Supplier

If additional assistance is required, contact your network supplier. Many suppliers are authorized 3Com service partners who are qualified to provide a variety of services, including network planning, installation, hardware maintenance, application training, and support services.

When you contact your network supplier for assistance, have the following information ready:

- Diagnostic error messages
- A list of system hardware and software, including revision levels
- Details about recent configuration changes, if applicable

If you are unable to contact your network supplier, see the following section on how to contact 3Com.

Support from 3Com

If you are unable to receive support from your network supplier, technical support contracts are available from 3Com.

In the U.S. and Canada, call **(800) 876-3266** for customer service.

If you are outside the U.S. and Canada, contact your local 3Com sales office to find your authorized service provider. Use one of these numbers:

Country	Telephone Number
Australia	1 800 678 515
Belgium*	0800 71429
Brazil	(55) (11) 546 0869
Canada	(416) 498 3266
Denmark*	800 17309
Finland*	0800 113153
France*	05 917959
Germany*	0130 821502
Hong Kong	(852) 2501 1111
Ireland*	1 800 553117
Italy*	1678 79489
Japan	(81) (3) 3345 7251

Country	Telephone Number
Mexico	(525) 532 0591
Netherlands*	06 0227788
Norway*	800 11376
Singapore	(65) 538 9368
South Africa	(27) (11) 803 7404
Spain*	900 983125
Sweden*	120 795482
Taiwan	(886) (2) 577 4352
United Arab Emirates	(971) (4) 349049
U.K.*	0800 966197
U.S.	(1) (408) 492 1790

* These numbers are toll-free.

Returning Products for Repair

Before you can send product directly to 3Com for repair, you must first obtain 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 In FDDI technology, each DAS contains two ports: A and 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 computer communication standards in the United States, developed the Fiber Distributed Data Interface (FDDI) standard.

AppleTalk® protocol 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 with which users can communicate directly. Users interact with the application layer through e-mail, and file transfer, and other services. 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.

ATM Adaption Layer (AAL) A set of protocols that translate higher-layer protocols into ATM format.

ATM Forum A consortium of vendors, carriers, and users formed to expedite industry agreement on ATM interfaces.

ATM Layer The part of the BISDN protocol stack that handles most of the ATM routing and processing.

backplane In the LANplex system, the “motherboard” that performs various logic and control functions. Located in the back of the chassis, it supports three 100Mbps FDDI paths, three 4 or 16 Mbps token ring paths, three 10Mbps Ethernet paths, and a VMEbus.

B port In FDDI technology, each DAS contains two ports: A and B. The B port is connected to the incoming fiber of the secondary ring and to 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*.

B to M link One of several detailed connection rules for a specific port relative to other ports. The B to M (master) port rule is a tree connection with possible redundancy. With this link, a station must not go to THRU state in Configuration Management (CFM). Port B on one station has precedence for connecting to an M port on a different station (single MAC station).

bridge Equipment that connects 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 that have different hardware and use different protocols. Examples: a bridge that connects an Ethernet network to an FDDI network allows the two networks to send signals to each other. The LANplex 2500 Ethernet/FDDI Switching Module (EFSM) can operate as a translation/transparent 802.1d bridge. See also *Spanning Tree Protocol*.

BUS Broadcast and Unknown Server. The set of functions implemented in an ATM network that provides LAN-to-LAN transmission support while a LAN connection is being established. The BUS also supports Ethernet Broadcasts Mode by sending broadcast data to all LECs.

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 stop responding or fail.

cell An ATM Layer protocol data-unit (PDU), characterized by fixed-length rather than variable-length payloads. The standard ATM cell is 48 bytes of payload with 5 bytes of header.

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.

CLP Cell loss priority. A bit in the ATM cell header indicating a need to discard the cell.

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

connection An ATM connection consists of the concatenation of ATM Layer links in order to provide an end-to-end information transfer capability to access points.

connectionless communications A form of packet switching that relies on global addresses in each packet rather than on predefined virtual circuits. Ethernet, FDDI, and token ring are connectionless.

connection-oriented communications A form of packet switching that requires a predefined circuit from source to destination to be established before data can be transferred.

concentrator An FDDI station having additional PHY/PMD entities beyond those required for its own attachment to an FDDI network. These additional PHY/PMD entities, called M (master) ports, are for the attachment of other FDDI stations (including other concentrators) in a tree topology.

DAC Dual attachment concentrator. A concentrator that offers two attachments to the FDDI network that are capable of accommodating a dual, counter-rotating ring. A DAC contains an A-B port pair and at least one M port.

DAS Dual attachment station. A station directly attached to FDDI's dual token rings. A DAS has four fiber attachments consisting of 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) and the Logical Link Control (LLC) sublayers. 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.

ELAN Emulated LAN. A virtual LAN created using the LANE specifications. See also *LANE*.

ESI End system identifier. A field in the ATM NSAP address that contains the MAC address (optional) of the end system.

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 LANplex system 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 FDDI, *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.

header Protocol control information located at the beginning of a protocol data unit. See also *PDU*.

hostname A meaningful, easy-to-remember name or title assigned by the user 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.

ILMI Interim Local Management Interface. ATM Forum-defined interim specifications for network management functions between a piece of ATM data terminal equipment (DTE) and an ATM switch over the user-to-network interface (UNI). Based on a limited subset of SNMP capabilities.

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 equipment from other vendors.

IP address Internet Protocol address. A unique identifier for a machine attached to a network that is 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. 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 peripheral devices. LANs are distinguished by their small geographical size, high data rate, and low error rate.

LANE LAN emulation. Set of ATM Forum-developed specifications for the operation of LAN-to-LAN bridged connectivity over an ATM network.

LCD Liquid crystal display. Display device consisting of a liquid crystal hermetically sealed between two glass planes. LCDs have low power requirements. The LANplex system control panel uses an LCD.

LEC LAN emulated client. The entity in end-stations, or clients, that performs data forwarding, address resolution, and other control functions.

LECS LAN emulation configuration server. The LECS assigns individual LECs to different Emulated LANs.

LES LAN Emulation Server. The LES implements the control coordination function for the Emulated LAN. It also provides a facility for registering and resolving MAC addresses and route descriptors to ATM addresses.

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. See also *data link layer*.

local path See *FDDI path*s.

M port Master port. Each PHY/PMD pair, designated a port, belongs to one of four types: A, B, M, or S. Concentrator stations (DAS and SAC) contain one or more M ports to provide connection within the concentrator tree.

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. See also *data link layer*.

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 LANplex system contains its own internal MIB.

multicast packet A single packet that is copied to a subset of addresses in a network. See also *broadcast packet*.

multicast storm See also *broadcast/multicast storms*.

multicast packet firewall A mechanism in the LANplex system that limits the rate at which multicast packets are forwarded through the system. This threshold is configurable.

NAC Null attachment concentrator. A system configured in such a way that it has no attachments to the FDDI dual ring. A no-attach station contains no A, B, or S ports.

NNI Network-to-network interface. The interface between two ATM network nodes.

nonvolatile memory Computer memory that is preserved when power is lost. Also called NVRAM.

OAM Operation and maintenance cell. This cell contains ATM maintenance and performance monitoring information. It does not form part of the upper-layer information transfer.

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 individual layer names*.

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.

PDU Protocol data unit. A unit of data specified in a layer protocol and consisting of protocol control information and layer user data.

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.

point-to-point call A two-way call or connection that has one source and one destination.

point-to-multipoint call A one-way call or connection that has only one source but may have many destinations.

primary path See *FDDI paths*.

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.

PVC Permanent virtual connection. A basic connection method that requires the user to define each connection manually.

remote management Usually refers to the management of workstations at distant locations linked to the main LAN by a network modem. Remote management can be performed in the LANplex system through a serial port connected to an exterior modem. See also *modem*.

repeater An FDDI node that serves as a two-way relay of the optical signals in an FDDI network. A repeater does not have MACs or concentrator functionality.

ring A series of stations across which information is passed sequentially, each station in turn examining or copying the information, finally returning it to the originating station. The ring has a predictable response time determined by the number of stations. See also *primary ring*, *secondary ring*, *FDDI dual ring*.

router A device that connects two remote networks by selectively forwarding messages between them. A router differs from a bridge and a gateway in that it selectively forwards information between the networks. Routers can be implemented in pairs, or a router can communicate directly with a computer. See also *bridge*.

roving MAC A MAC that is placed strategically within the LANplex system to pinpoint and diagnose network faults.

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.

S port Slave port. Each PHY/PMD pair, designated a port, belongs to one of four types: A, B, M, or S. A single attachment station (SAS or SAC) has an S port intended to be attached to an M port within a concentrator tree.

SAC Single attachment concentrator. A concentrator that offers one attachment to the FDDI network. A SAC has an S port to be attached to an M port within a concentrator tree.

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 path See *FDDI paths*.

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 implemented at present on a wide variety of computers and networking equipment. It may be used to manage many aspects of network and end-station operation. See also *protocol*.

SONET Synchronous Optical Network. ANSI standard for transmission over optical fiber. Used in the U.S. and Canada. A variation of the synchronous digital hierarchy (SDH) international standard.

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. See also *IEEE 802.1d*.

station An addressable logical and physical attachment in a ring that is capable of transmitting, receiving, and repeating information. An FDDI station has one or more PHY entities, one or more MAC entities, and only one SMT entity.

station ID The unique identifier for an FDDI station or concentrator.

SVC Switched virtual connection. An ATM standard signalling protocol that automatically establishes and releases connections as they are needed.

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 interconnection of or more LANs so they can communicate as if they were one LAN. The bridge listens to packets on the attached LANs and forwards packets from one LAN to another. See also *bridge*, *IEEE 802.1d*.

UME UNI management entity. The UMI is the code residing in the ATM devices at each end of a UNI circuit that implements the management interface to the ATM network.

UNI User-to-network interface. The UNI interconnects an ATM user with an ATM switch.

UNIX® A computer operating system, developed by AT&T, that is capable of multitasking.

VC Virtual circuit. A connection between end users that has defined end points and a route but does not have bandwidth dedicated to it. Bandwidth is allocated on demand by the network as users have traffic to transmit.

VCC Virtual channel connection. Virtual channels in two or more sequential physical circuits can be concatenated to create an end-to-end connection, called a VCC. A VCC Is a specific instance of a SVC or PVC. A VCC may traverse one end-to-end VPC or several sequential VPCs.

VCI Virtual channel identifier. Part of the identifier of a particular virtual circuit in the ATM fabric.

VPI Virtual path identifier. Part of the identifier of a particular virtual circuit in the ATM fabric.

INDEX

Numerics

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

A

AAL sublayer
  convergence (CS) 11-9
  segmentation and reassembly (SAR) 11-9
address
  aging 5-3, 8-3
    with Express switching 6-2
  destination 5-4
  learning 5-2, 6-3
  learning with Express switching 6-2
  network security 8-4
  source 5-3
  table 7-8
address groups
  defined 7-8
  example 7-11
  mask 7-8, 7-9
aging addresses 5-3, 8-3
AppleTalk
  Phase I 7-4
  Phase II 7-6
ATM
  adaption layer protocols 11-9
  address registration 11-10
  benefits 11-1
  cell information field 11-4
  cell structure 11-3
  ILMI MIB 11-10
  network interfaces 11-2
  physical layer 11-8
  service classes 11-8
attachments 9-6
authenticationFailure 3-5

B

BISDN 11-1
blocking port state 5-16
bridge
  aging timer 8-3
  designated 5-8
  least cost path 5-9
  root 5-8
Bridge Protocol Data Unit (BPDU) 5-18
bridging
  configuration messages 5-8
  extensions 8-1 to 8-4
  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
BUS 11-15

C

cabling 4-1 to 4-4
CBPDU
  best 5-12
  information 5-10
CCITT 11-1
cell
  ATM 11-3
cell information field 11-4
cell loss priority 11-4
cell structure
  ATM 11-3
circuit identifier 11-5
circuits
  permanent virtual circuits 11-6
  switched virtual circuits 11-5
CLP 11-4
coldStart 3-5
CompuServe® B-2
concentrator
  dual-attachment 9-10
Configuration Bridge Protocol Data Unit. *See* CBPDU
connection-oriented 11-2
control panel 2-3
conventions
  notice icons, About This Guide 2
  text, About This Guide 3
convergence sublayer 11-9

D

DAS 9-10
designated bridge 5-8
designated port 5-8, 5-9
disabled port state 5-17
dual homing 10-6
dual ring, FDDI 10-4
dual-attachment stations 9-10

E

emulated LAN (ELAN)
  creating 11-16
  described 11-17
  eliminating bridge loops 11-17
Express switching 6-1 to 6-4
  address aging 6-2
  address learning 6-2
  packet forwarding 6-2
external network management 2-4

F

fax service B-3
FDDI 3-6
  and OSI model 9-2
  capabilities 9-1
  connection rules 10-4
  dual homing 10-6
  dual ring 10-4
  events 3-5
  MIB 9-8
  networks 10-1 to 10-7
  nodes 9-5
  optical bypass switch 9-10
  overview 9-1
  paths 9-5
  ports 9-3
  proxy agents 3-7
  standards 9-2
  topologies 10-1
  trunk ring 10-3
filtering. *See* packet filter
flooding, packet 5-5, 8-3
forwarding
  packets 5-4, 6-3
  port state 5-16
fragmentation. *See* IP fragmentation
frame-based protocols 9-9

G

generic flow control 11-4
groups, address. *See* address group
groups, port. *See* port group

H

header error control 11-4
HEC 11-4

I

IEEE 802.1d
  bridging standards 5-1
in-band management 4-2
IP fragmentation
  defined 8-2
  RFC specified in 8-2
ITU 11-1

L

LAN emulation 11-12
  broadcast and unknown server (BUS) 11-15
  client (LEC) 11-15
  configuration server (LECS) 11-15
  server (LES) 11-15
LAN, FDDI 3-6
LANE 11-12
LAN-layer 3-6
LANplex
  Administration Console 2-3
  external network management 2-4
  managing from outside the network 2-4
  overview of operations 1-1
layers, OSI Reference Model 9-2
learning port state 5-16
LEC 11-15
LECS 11-15
LES 11-15
listening port state 5-16
looping, packet 5-6

M

M port 9-4
MAC (Media Access Control) 9-2, 9-4
management access 4-1 to 4-4

managing LANplex
- from outside the network 2-4
- protocols 3-1 to 3-4

managing LANplex systems
- protocols 3-1 to 3-4

Master port 9-4

MIB 3-5
- events 3-5
- FDDI 3-5, 9-8
- ILMI 11-10
- LANplex 3-5
- SNMP A-1 to A-3

MIB II
- traps 3-5

multicast packet firewall
- defined 8-1
- multicast storm 8-2
- receiveMulticastThreshold attribute 8-1
- threshold mechanism 8-2

multicast storm 8-1

multicast/broadcast traffic 8-1

N

network
- ATM 11-1 to 11-20
- FDDI 10-1 to 10-7
- reduced packet flooding 8-3
- security 8-4

network supplier support B-3

network topologies, FDDI 10-2

newRoot 3-5

nodes
- FDDI, on 9-5
- types 9-6

O

on-line technical services B-1

operand 7-1

operator 7-1

optical bypass switch
- on intermediate systems 9-10

OSI Reference Model
- and FDDI 9-2

out-of-band management 4-2

P

packet
- fields for operands 7-2
- flooding 5-5, 8-3
- forwarding 5-4, 6-2
- fragmentation 8-2
- looping 5-6
- reduced flooding 8-3

packet filter
- address groups 7-8, 7-11
- and an AppleTalk network 7-4
- and Express switching 6-3
- examples 7-3 to 7-7
- implementation on paths 7-1
- language 7-1
- purpose 7-1
- user-defined 7-1 to 7-16

paths
- FDDI, on 9-5
- receive 7-1
- transmit 7-1

payload type 11-4

PDU 11-9

permanent virtual circuits 11-6

PHY standard
- defined 9-2

physical access 4-1 to 4-4

physical layer
- ATM 11-8

PMD standard
- defined 9-2

port
- bridging states 5-16, 5-17
- designated 5-8
- FDDI 9-3
- identifier 5-10
- root 5-8, 5-9
- table 7-9

port groups
- defined 7-9
- mask 7-10

port state
- learning 5-16
- listening 5-16

protocol data unit 11-9

protocols 3-1 to 3-4
- frame-based 9-9
- SNMP 3-4 to 3-6
- virtual terminal 3-3

proxy agents 3-7

PT 11-4

PVC 11-6

Q

quality of service (QOS) 11-8

R

receive path 7-1
returning products for repair B-4
ring of trees 10-3
rlogin 3-3
root bridge 5-8
root port 5-8

S

S port 9-4
security
  manual address assignment 8-4
Slave port 9-4
SMT 9-2, 9-8
SNMP 3-4 to 3-6
  management applications 2-5
  MIBs A-1 to A-3
  traps 3-5
Spanning Tree
  algorithm 5-7
  and packet looping 5-6
  blocking paths 5-7
  BPDU 5-18
  CBPDU 5-8
  designated bridge 5-8
  designated port 5-8
  port identifier 5-10
  reconfiguring the topology 5-18
  root bridge 5-8
  root port 5-8
station
  FDDI 9-5
SunNet Manager™ 2-5
SVC 11-5
switched virtual circuits 11-5
switching. *See* Express switching

T

table
  address 7-8
  port 7-9
TCP/IP 3-3
technical support B-1
telnet 3-3
terminal
  emulation 3-3
  virtual 3-3
token 9-4
Topology Change Flag 5-18
topology, FDDI 10-2 to 10-4
Transcend® Enterprise Manager
  SNMP community strings, and 3-6
Transcend® Enterprise Manager software 3-4
transmit path 7-1
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
  tree topology 6-2
traps, SNMP 3-5

U

UDP/IP networks 3-4
User Datagram Protocol 3-4
user interfaces on the LANplex system 2-2
user-defined packet filtering 7-1 to 7-16

V

virtual channel 11-6
virtual channel connection (VCC) 11-5
virtual channel identifier (VCI) 11-5
virtual circuits (VC) 11-5
virtual LAN (VLAN) 11-16
virtual path (VP) 11-6
virtual path identifier (VPI) 11-4
virtual terminal protocols 3-3
virtual workgroup 7-1