




LANPLEX® 2500 ADMINISTRATION CONSOLE USER GUIDE



Part No. 801-00322-000
Published November 1996
Revision 03

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

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

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

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

UNITED STATES GOVERNMENT LEGENDS:

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

For units of the Department of Defense:

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

For civilian agencies:

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 Unpublished rights reserved under the copyright laws of the United States.

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

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

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

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

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

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

CONTENTS

ABOUT THIS GUIDE

- Introduction 1
- How to Use This Guide 2
- Conventions 3
- LANplex 2500 Documentation 4
- Documentation Comments 6

PART I INTRODUCTION

1 LANPLEX® 2500 ADMINISTRATION OVERVIEW

- About LANplex Administration 1-1
- Configuration Tasks 1-1

2 HOW TO USE THE ADMINISTRATION CONSOLE

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

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

PART II SYSTEM-LEVEL FUNCTIONS

3 CONFIGURING MANAGEMENT ACCESS TO THE SYSTEM

About Management Access	3-1
Using a Serial Connection	3-1
Using an IP Interface	3-1
In-band or Out-of-band?	3-2
Setting Up the Terminal Serial Port	3-2
Setting Up the Modem Serial Port	3-3
Setting the Port Speed	3-3
Configuring the External Modem	3-3
Setting Up an IP Interface for Management	3-4
General Setup Process	3-4
Administering Interfaces	3-4
Displaying Interfaces	3-5
Defining an IP Interface	3-6
Modifying an IP Interface	3-7
Removing an Interface	3-7
Administering Routes	3-7
Displaying the Routing Table	3-8
Defining a Static Route	3-9
Removing a Route	3-9
Flushing All Learned Routes	3-10
Setting the Default Route	3-10
Removing the Default Route	3-11
Administering the ARP Cache	3-11
Displaying the ARP Cache	3-11
Removing an ARP Cache Entry	3-11
Flushing ARP Cache Entries	3-12

Setting the RIP Mode	3-12
Pinging an IP Station	3-13
Displaying IP Statistics	3-14
Setting Up SNMP on Your System	3-15
Displaying SNMP Settings	3-15
Configuring Community Strings	3-15
Administering SNMP Trap Reporting	3-16
Displaying Trap Reporting Information	3-16
Configuring Trap Reporting	3-18
Removing Trap Destinations	3-19
Flushing All SNMP Trap Destinations	3-19
Setting Up SMT Event Proxying	3-19

4 ADMINISTERING YOUR SYSTEM ENVIRONMENT

Displaying the System Configuration	4-1
Setting Passwords	4-2
Setting the System Name	4-3
Changing the Date and Time	4-4
Rebooting the System	4-5
Displaying the System Up Time	4-5

5 BASELINING STATISTICS

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

6 SAVING, RESTORING, AND RESETTNG NONVOLATILE DATA

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

PART III ETHERNET, FDDI, AND ATM PARAMETERS

7 ADMINISTERING ETHERNET PORTS

- Displaying Ethernet Port Information 7-1
- Enabling or Disabling Full-Duplex Mode 7-7
- Labeling a Port 7-8
- Setting the Port State 7-9

8 ADMINISTERING FDDI RESOURCES

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

9 ADMINISTERING ATM

- ATM in Your Network 9-1
 - LAN Emulation
and Classical IP 9-1
- LAN Emulation 9-2
 - Before You Configure an ELAN 9-2
 - Checking Link Status 9-3
 - Verifying Address Registration 9-3

Verifying Signaling	9-4
Creating an Emulated LAN	9-4
Configuring Clients to Join an Existing Emulated LAN	9-5
Administering LECs (LAN Emulation Clients)	9-5
Displaying Information About LAN Emulation Clients	9-5
Modifying Information About LAN Emulation Clients	9-9
Defining LAN Emulation Clients	9-10
Preventing ATM Network Loops	9-11
Removing a LAN Emulation Client	9-12
Administering UNI Management Entities	9-12
Displaying UME Information	9-13
Listing Network Prefixes and Addresses	9-14
Setting the UME Connect State	9-15
Setting the Virtual Path Identifier	9-15
Setting the Virtual Channel Identifier	9-16
Administering ATM Ports	9-16
Displaying Port Information	9-16
Labeling a Port	9-19
Listing Virtual Channel Connection Information	9-19
Listing General VCC Information	9-19
Listing VCC Transmit Information	9-20
Listing VCC Receive Information	9-21

10 SETTING UP THE SYSTEM FOR ROVING ANALYSIS

About Roving Analysis	10-1
Displaying the Roving Analysis Configuration	10-3
Adding an Analyzer Port	10-3
Removing an Analyzer Port	10-4
Starting Port Monitoring	10-5
Stopping Port Monitoring	10-6

PART IV BRIDGING PARAMETERS

11 ADMINISTERING THE BRIDGE

Displaying Bridge Information	11-1
Setting the Bridging Mode	11-4
Enabling and Disabling IP Fragmentation	11-6
Enabling and Disabling IPX Snap Translation	11-6
Setting the Address Threshold	11-7
Setting the Aging Time	11-7

- Administering STP Bridge Parameters 11-8
 - Enabling and Disabling STP on a Bridge 11-8
 - Setting the Bridge Priority 11-8
 - Setting the Bridge Maximum Age 11-9
 - Setting the Bridge Hello Time 11-10
 - Setting the Bridge Forward Delay 11-10
 - Setting the STP Group Address 11-11
-

12 ADMINISTERING BRIDGE PORTS

- Displaying Bridge Port Information 12-1
 - Setting the Multicast Limit 12-7
 - Administering STP Bridge Port Parameters 12-8
 - Enabling and Disabling STP on a Port 12-8
 - Setting the Port Path Cost 12-9
 - Setting the Port Priority 12-10
 - Administering Port Addresses 12-11
 - Listing Addresses 12-11
 - Adding New Addresses 12-12
 - Removing Addresses 12-12
 - Flushing All Addresses 12-13
 - Flushing Dynamic Addresses 12-13
 - Freezing Dynamic Addresses 12-14
-

13 CREATING AND USING PACKET FILTERS

- About Packet Filtering 13-1
- Listing Packet Filters 13-2
- Displaying Packet Filters 13-3
- Creating Packet Filters 13-3
 - Concepts for Writing a Filter 13-4
 - How the Packet Filter Language Works 13-4
 - Basic Elements of a Packet Filter 13-6
 - Implementing Sequential Tests in a Packet Filter 13-8
 - Preprocessed and Run-time Storage 13-9
 - Procedure for Writing a Filter 13-10
 - Examples of Creating Filters 13-11
 - Filtering Problem 13-11
 - Packet Filter Solution 13-12
 - Tools for Writing a Filter 13-17
 - Using the Built-in Line Editor 13-18
 - Using an External Text Editor 13-20
- Deleting Packet Filters 13-20
- Editing, Checking, and Saving Packet Filters 13-20

Loading Packet Filters 13-22
Assigning Packet Filters to Ports 13-22
Unassigning Packet Filters from Ports 13-24

14 CONFIGURING ADDRESS AND PORT GROUPS TO USE IN PACKET FILTERS

Using Groups in Packet Filters 14-1
Listing Groups 14-2
Displaying Groups 14-3
Creating New Groups 14-4
Deleting Groups 14-6
Adding Addresses and Ports to Groups 14-7
Removing Addresses or Ports from a Group 14-9
Loading Groups 14-11

PART V APPENDICES

A PACKET FILTER OPCODES, EXAMPLES, AND SYNTAX ERRORS

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

B TECHNICAL SUPPORT

Online Technical Services B-1
 3Com Bulletin Board Service B-1
 Access by 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

INDEX

ABOUT THIS GUIDE

Introduction

The *LANplex® 2500 Administration Console User Guide* provides all the information you need to configure and manage your LANplex system once it is installed and the system is attached to the network. Prior to using this guide, you should have already installed and set up your system using the *LANplex 2500 Getting Started* guide.

Audience description

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



If the information in the Software Installation and Release Notes shipped with this product differs from the information in this guide, follow the release notes.

How to Use This Guide

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

Table 1 Description of Guide Parts

Refer to this part...	If you want to ...
I: Introduction	Learn about LANplex system administration Learn about the various system configurations and the quick commands to perform them Learn about password access to the Console Learn about the Administration Console menu structure and maneuver within the Console by using commands and moving between menus Set interface parameters (screen height and control keys) Run scripts of Console tasks Get help
II: System-Level Functions	Set up the system for management access through serial ports or using IP and setting up SNMP Configure SNMP community strings Set up trap reporting Configure system parameters, such as name, date/time, and passwords Baseline statistics Save, restore, and reset nonvolatile data
III: Ethernet, FDDI, and ATM Parameters	Display statistics for and labeling Ethernet ports Display statistics for and configuring various parameters for FDDI stations, ports, MACs, and paths Set up the system to monitor Ethernet port activity using roving analysis Display statistics for and configure various parameters for ATM ports

(continued)

Table 1 Description of Guide Parts (continued)

Refer to this part...	If you want to ...
IV: Bridging Parameters	Configure bridge and bridge port parameters Administer the Spanning Tree Protocol bridge and bridge port parameters Display and configure bridge port addresses Create and use packet filters Create address groups and port groups and use them as filtering criteria
V: Appendixes	Read additional information about packet filters: opcode descriptions, examples, and error messages Get Technical Support Return products for repair

Conventions

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

Table 2 Notice Icons




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

Table 3 Text Conventions

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

LANplex 2500 Documentation

The following documents comprise the LANplex 2500 documentation set. To order a document that you do not have or order additional documents, contact your sales representative for assistance.

- *LANplex® 2500 Unpacking Instructions*
Describes how to unpack your LANplex system. It also gives an inventory list of all the items that came with your system. (Shipped with system/Part No. 801-00353-000)
- *LANplex® 2500 Software Installation and Release Notes*
Provides information about the software release, including new features and bug fixes. It also provides information about any changes to the LANplex system's documentation. (Shipped with system)

- *LANplex® 2500 Getting Started*

Describes all the procedures necessary for planning your configuration and for installing, cabling, powering up, and troubleshooting your LANplex system. (Shipped with system/Part No. 801-00335-000)
- *LANplex® 2500 Operation Guide*

Provides information to help you understand system management and administration, FDDI technology, ATM technology, and bridging. It also describes how these concepts are implemented in the LANplex system. (Shipped with system/Part No. 801-00344-000)
- *LANplex® 2500 Administration Console User Guide* (this guide)

Provides information about using the Administration Console to configure and manage your LANplex system. (Shipped with system/Part No. 801-00322-000)
- *LANplex® 2500 Extended Switching User Guide*

Describes how the routing protocols, VLANs, and RMON are implemented in the LANplex system as well as providing information about using the Administration Console to configure and manage these features. (Order from 3Com/Part No. 801-00343-000)
- *LANplex® 2500 Intelligent Switching Administration Console Command Quick Reference* card

Contains all of the Administration Console intelligent switching commands for the LANplex system. (Folded card; shipped with system/Part No. 801-00318-000)
- *LANplex® 2500 Extended Switching Administration Console Command Quick Reference* card

Contains all of the Administration Console Extended Switching commands for the LANplex system. (Shipped with option package/Part No. 801-00319-000)
- *Module Installation Cards*

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

Documentation Comments


Your suggestions are very important to us. To help make LANplex documentation more useful to you, please send comments about this guide in an e-mail message to 3Com at:

sdtechpubs_comments@3Mail.3Com.com

Please include the following information when commenting:

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


*Example: LANplex® 2500 Operation Guide
Part No. 801-00344-000
Page 2-5 (chapter 2, page 5)*




INTRODUCTION

Chapter 1 LANplex® 2500 Administration Overview

Chapter 2 How to Use the Administration Console





LANPLEX® 2500 ADMINISTRATION OVERVIEW

This chapter introduces you to LANplex® 2500 system administration and briefly describes the switching system parameters that you can configure.



For information on LANplex Extended Switching functionality, refer to the LANplex® 2500 Extended Switching User Guide.

About LANplex Administration

The LANplex system switching software is installed at the factory in flash memory on the system processor. Because this software boots from flash memory automatically when you power on your system, the system is immediately ready for use in your network. However, you might need to configure certain parameters for the system to operate effectively in your networking environment. Additionally, when managing your LANplex 2500 system, you might want to view important MAC, port, bridge, and IP statistics. The LANplex 2500 Administration Console software allows you to configure your system and display these important statistics. For more complete network management, you can use an external application, such as 3Com's Transcend® Enterprise Manager suite of tools.

Configuration Tasks

This section uses tables to summarize the Intelligent Switching tasks and quick commands for the LANplex 2500 Administration Console.

- General System Commands (Table 1-1)
- System Management Setup Commands (Table 1-2)
- Bridging Commands (Table 1-3)
- Ethernet Commands (Table 1-4)
- FDDI Commands (Table 1-5)
- ATM Commands (Table 1-6)

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

Table 1-1 General System Commands

Task	Quick Command	For Details, See . . .
<p>Run a script of commands to set up a system Write a script of Console commands with the values you assign so that you can quickly configure one or more systems. You can run the same script on a number of systems to ensure consistent setup.</p>	script	
<p>Display the system configuration Display software and hardware revisions, module status information, and certain warning messages.</p>	system display	
<p>Install software into flash memory Update your system software. Software is initially installed at the factory. If you have purchased the LANplex® Extended Switching Software, you must install that software using this command.</p>	system softwareUpdate	<i>LANplex 2500 Software Installation and Release Notes</i>
<p>Display, set, enable, or disable a baseline for statistics Establish and use baselines for Ethernet, FDDI, and bridging statistics to evaluate recent activity in your system and on your network.</p>	system baseline	
<p>Configure timeout for remote sessions Configure the system to disconnect remote sessions after a specified time interval.</p>	system telnet	
<p>Control access to the Console Set passwords for levels of access (read, write, administer) and prohibit remote access during your session by locking the Console.</p>	system password system consoleLock	
<p>Name the system Assign a unique name to the system for management purposes. For example, you might name a system based on its location: <i>LANplex-Floor2</i>.</p>	system name	
<p>Set the system date and time Ensure that messages are accurately logged. The internal clock is set at the factory; change it for your time zone.</p>	system time	

(continued)

Table 1-1 General System Commands (continued)

Task	Quick Command	For Details, See. . .
Set screen height Adjust the Console screen height for your terminal.	<code>system screenHeight</code>	
Enable the [Control] keys when working in the Console Enable quick keys for the reboot (Ctrl+X) and abort (Ctrl+C) functions.	<code>system ctlKeys</code>	
Save, restore, or reset nonvolatile data in the system Provide a backup for nonvolatile data, restore nonvolatile data to the system, or reset nonvolatile data to defaults.	<code>system nvData</code>	
Reboot the system Restart the system. Disconnects rlogin and telnet sessions.	<code>system reboot</code>	
Display the system up time Display the amount of time the system has been running since the last reboot.	<code>system upTime</code>	page 4-5

Table 1-2 System Management Setup Commands

Task	Quick Command	For Details, See. . .
Configure the terminal serial port baud rate Change the factory default baud rate of the serial port, which allows you to connect a VT or tty type of terminal or terminal emulator to the system using a null modem cable.	<code>system serialPort terminalSpeed</code>	
Set up the system for an external modem Manage your system remotely with an external modem. You can set the modem serial port baud rate and configure an external modem.	<code>system serialPort connectModem</code> <code>system serialPort modemSpeed</code>	
Configure an IP address using an IP interface Communicate with the system using SNMP, rlogin, or telnet.	<code>ip interface display</code> <code>ip interface define</code> <code>ip interface modify</code> <code>ip interface remove</code>	
Define static routes Access a menu from which you can display, define, remove, and flush static routes for transmitting traffic through the system. Static routes override routes learned through RIP.	<code>ip route</code> <code>ip route default</code>	

(continued)

Table 1-2 System Management Setup Commands

Task	Quick Command	For Details, See. . .
Administer the ARP cache Display, remove, and flush the ARP cache (a table of known IP addresses and their corresponding MAC addresses).	ip arp display ip arp remove ip arp flush	
Set RIP's operational mode Define how Routing Information Protocol (RIP) messages are processed.	ip rip	
Ping an IP station or the system Find out if the system can reach an IP station or check that the system is on the network.	ip ping	
Display IP statistics Display datagram statistics and current RIP operational mode.	ip statistics	
Configure SNMP management Display current SNMP configurations and specify the type of authorization for SNMP management.	snmp display snmp community	
Configure SNMP trap reporting Display SNMP trap reporting information, add or modify trap reporting destination configurations, remove trap destinations, flush all SNMP trap reporting destinations, and set up SMT event proxying.	snmp trap display snmp trap addModify snmp trap remove snmp trap flush snmp trap smtProxyTraps	

Table 1-3 Bridging Commands

Task	Quick Command	For Details, See. . .
Display bridge information Display information about the bridge, such as statistics, bridge configurations, and spanning tree configurations.	bridge display	
Set the bridging mode Specify whether the bridge operates in IEEE 802.1d bridging mode or Express switching mode. The default is 802.1d.	bridge mode	
Enable or disable IP fragmentation Enable or disable the fragmenting of large FDDI packets to allow FDDI and Ethernet stations to communicate using IP.	bridge ipFragmentation	

(continued)

Table 1-3 Bridging Commands (continued)

Task	Quick Command	For Details, See . . .
<p>Enable or disable IPX snap translation</p> <p>Enable or disable the translation of 802.3_RAW IPX packets to FDDI_SNAP packets (when going from Ethernet to FDDI), and vice versa (when going from FDDI to Ethernet). The default is disabled.</p>	<code>bridge ipxSnapTranslation</code>	
<p>Set the bridge address threshold</p> <p>Specify the reporting threshold for the total number of Ethernet addresses known to the bridge. When the threshold is reached, the SNMP trap <i>addressThresholdEvent</i> is generated.</p>	<code>bridge addressThreshold</code>	
<p>Set the bridge address aging timer</p> <p>Specify how often dynamically learned addresses are aged by the bridge port. Appropriately configured aging prevents packet flooding.</p>	<code>bridge agingTime</code>	
<p>Configure Spanning Tree Protocol (STP) parameters for a bridge</p> <p>Enable or disable STP and set the bridge priority, the maximum age of stored configuration message information, the period between the generation of messages by a root bridge, the amount of time a bridge spends in the listening and learning states, and the group address.</p>	<code>bridge stpState</code> <code>bridge stpPriority</code> <code>bridge stpMaxAge</code> <code>bridge stpHelloTime</code> <code>bridge stpForwardDelay</code> <code>bridge stpGroupAddress</code>	
<p>Display bridge port information</p> <p>Display information about the bridge port, including STP configurations, in a summarized or detailed format.</p>	<code>bridge port summary</code> <code>bridge port detail</code>	
<p>Configure Spanning Tree Protocol (STP) parameters for a bridge port</p> <p>Enable or disable STP on a bridge port, and set the bridge port path cost and port priority.</p>	<code>bridge port stpState</code> <code>bridge port stpCost</code> <code>bridge port stpPriority</code>	
<p>Set the multicast packet firewall threshold</p> <p>Suppress multicast storms and limit the rate at which multicast packets are propagated by the system.</p>	<code>bridge port multicastLimit</code>	
<p>Administer bridge port addresses</p> <p>Administer the MAC address of stations connected to Ethernet and FDDI ports. This command accesses a menu from which you can list, add, remove, flush, and freeze bridge port addresses.</p>	<code>bridge port address</code>	

(continued)

Table 1-3 Bridging Commands (continued)

Task	Quick Command	For Details, See. . .
<p>Use packet filters to restrict which packets are forwarded through a bridge port</p> <p>Access a menu from which you can list packet filters, display a packet filter definition, create or edit a definition, load a definition onto the system, copy a definition, and assign or unassign a definition to a port.</p>	<code>bridge packetFilter</code>	
<p>Create address and port groups to use as filtering criteria</p> <p>Access a menu from which you can specify groups (either address groups or port groups) to use in a packet filter definition. From each menu, you can list, display, create, and delete groups. You can also add and remove address and ports to and from groups.</p>	<code>bridge packetFilter addressGroup</code> <code>bridge packetFilter portGroup</code>	

Table 1-4 Ethernet Commands

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

Table 1-5 FDDI Commands

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

Table 1-6 ATM Commands

Option	Quick Command	For Details, See . . .
Display LAN emulation client information Display information about emulated LAN clients in a summarized or detailed format.	atm lane lec summary atm lane lec detail	page 9-5
Define a LAN emulation client Provide information necessary for a client to join an emulated LAN.	atm lane lec define	page 9-10
Modify or remove a LAN emulation client Modify LAN emulation client parameters and remove a client from the emulated LAN.	atm lane lec modify atm lane lec remove	page 9-9 and page 9-12

(continued)

Table 1-6 ATM Commands (continued)

Option	Quick Command	For Details, See . . .
<p>Display UME information for ATM ports Display User-to-Network Interface (UNI) Management Entity (UME) information, including the connection state, virtual path identifier, and virtual channel identifier.</p>	<p><code>atm ume display</code></p>	<p>page 9-13</p>
<p>Administer UME information for ATM ports List registered network prefixes and addresses, set the connect state for management access and address registration, and set the identifier for the virtual path and virtual channel.</p>	<p><code>atm ume list</code> <code>atm ume state</code> <code>atm ume vpi</code> <code>atm ume vci</code></p>	<p>page 9-14 and following</p>
<p>Display ATM port information Display information about ATM port labels, status, activity, and errors in a summarized or detailed format.</p>	<p><code>atm ports summary</code> <code>atm ports detail</code></p>	<p>page 9-16</p>
<p>Label an ATM port Assign a unique name to an ATM port. Useful for port identification when managing the system.</p>	<p><code>atm ports label</code></p>	<p>page 9-19</p>
<p>List Virtual Channel Connection (VCC) information Access a menu from which you can list general information, as well as transmit and receive information, about the virtual channel connection.</p>	<p><code>atm ports vcc</code></p>	<p>page 9-19</p>

2

HOW TO USE THE ADMINISTRATION CONSOLE

This chapter familiarizes you with user access levels of the LANplex 2500® Administration Console and explains how to:

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

Initial User Access

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

Levels of User Access

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

Table 2-1 Password Access Levels

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

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

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

For information about setting passwords, see page 4-2. The passwords are stored in nonvolatile (NV) memory. You must enter the password correctly before you are allowed to continue.

The following examples show how the top-level menu structure changes based on the level of access.

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

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

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

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

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

```
Menu options: -----
display                - Display the system configuration
baseline               - Administer a statistics baseline
serialPort             - Administer the terminal and modem serial ports
name                   - Set the system name
screenHeight           - Set the console screen height
```

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

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

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

```

Menu options: -----
      display                - Display the system configuration
      baseline               - Administer a statistics baseline

Only the display option in the
baseline menu is available
Type 'q' to return to the previous menu or ? for help.
-----
Select a menu option (system):

```

Using Menus to Perform Tasks

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

```

                                     Option Descriptions
                                     |-----|
Menu options: -----
      system                   - Administer system-level functions
      ethernet                 - Administer Ethernet ports
      fddi                     - Administer FDDI resources
      atm                     - Administer ATM
      bridge                   - Administer bridging/VLANs
      ip                       - Administer IP
      snmp                     - Administer SNMP
      analyzer                 - Administer Roving Analysis
      script                   - Run a script of console commands
      logout                   - Logout of the Administration Console

The options you see vary with
levels of access
Type ? for help.
-----
Select a menu option:

```

Administration Console Menu Structure

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



The following menus display the options available for users with administer access. This access provides the most complete set of options.

System Menu

From the **system** menu, you can view the system configuration, set up your system for management, configure Administration Console interface parameters, work with nonvolatile data, and reboot the system. See Figure 2-1. For example, to configure an external modem from the Administration Console, enter **system** at the top-level menu, **serialPort** at the system menu, and then **connectModem** at the serialPort menu.

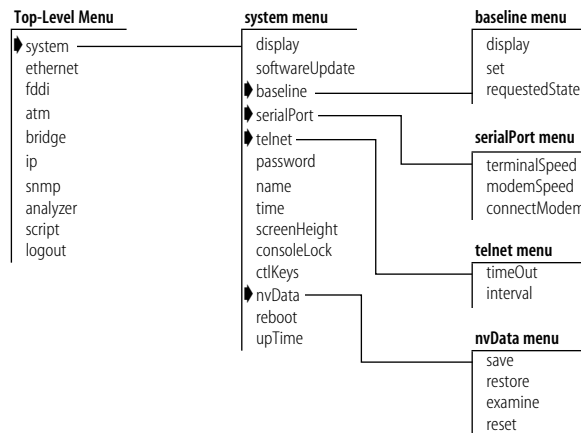


Figure 2-1 System Menu Hierarchy for Administer Access

Ethernet Menu

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




Figure 2-2 Ethernet Menu Hierarchy for Administer Access

FDDI Menu

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

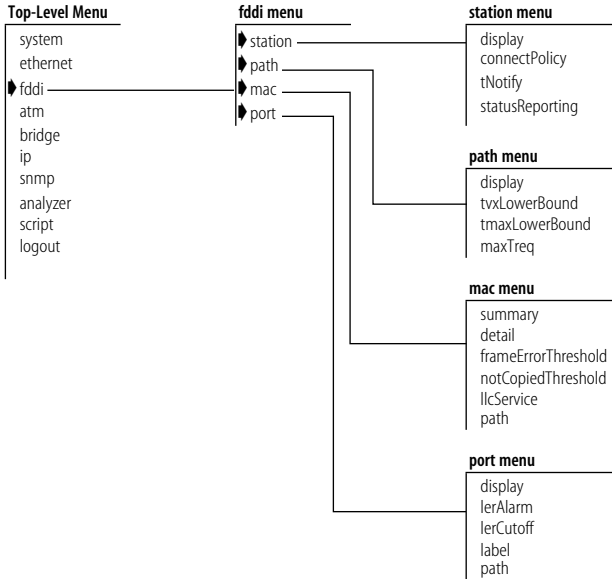


Figure 2-3 FDDI Menu Hierarchy for Administer Access

ATM Menu

From the **atm** menu, you can view information about and configure LAN Emulation (LANE), the UNI Management Entity (UME), and ATM ports. See Figure 2-4. For example, to set the UNI Management Entity state, you enter **atm** at the top-level menu, **ume** at the atm menu, and then **state** at the ume menu.

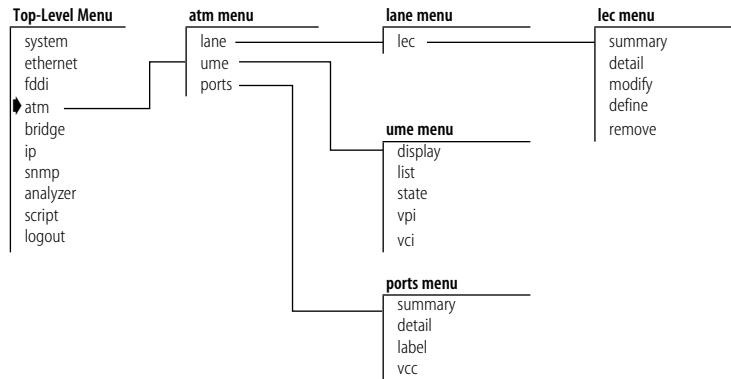


Figure 2-4 ATM Hierarchy for Administer Access

Bridge Menu

From the **bridge** menu, you can view information about and configure bridge-level parameters, including those for the Spanning Tree Protocol (STP). You can also configure the bridge at the port level and administer packet filters. See Figure 2-5. For example, to set the Spanning Tree state for a bridge port, enter **bridge** at the top-level menu, **port** at the bridge menu, and then **stpState** at the port menu.

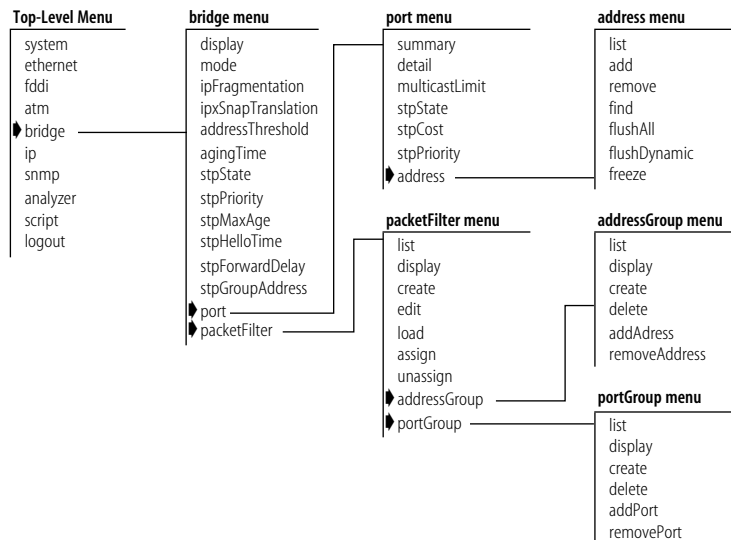


Figure 2-5 Bridge Menu Hierarchy for Administer Access

IP Menu

From the **ip** menu, you can view information about and configure Internet Protocol (IP) interfaces and routes. You can also administer the Address Resolution Protocol (ARP) and the Routing Information Protocol (RIP), and you can ping IP stations. See Figure 2-6. For example, to define a new IP interface, enter **ip** at the top-level menu, **interface** at the ip menu, and then **define** at the interface menu.

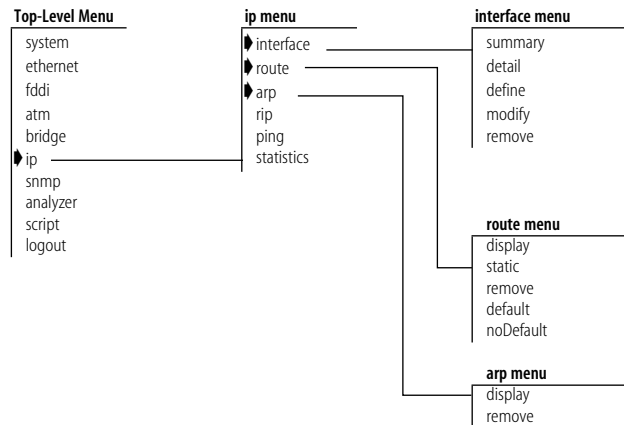


Figure 2-6 IP Menu Hierarchy for Administer Access

SNMP Menu

From the **snmp** menu, you can configure SNMP community strings and trap reporting. See Figure 2-7. For example, to flush all trap reporting destinations, enter **snmp** at the top-level menu, **trap** at the snmp menu, and then **flush** at the trap menu.




Figure 2-7 SNMP Menu Hierarchy for Administer Access

Analyzer Menu

From the **analyzer** menu, you can selectively choose any Ethernet network segment attached to a LANplex system and monitor its activity using a network analyzer. See Figure 2-8. For example, to add analyzer ports, enter **analyzer** at the top-level menu, and then **add** at the analyzer menu.




Figure 2-8 Analyzer Menu Hierarchy for Administer Access

Selecting Menu Options

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

Select a menu option: **system**

OR

Select a menu option: **sy**



Menu options are not case sensitive.

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

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

Select a menu option (system/baseline):

*Entering a
command string*

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

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

The most abbreviated version of the same command string is:

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

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

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

Entering Values

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

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

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

*Entering values in
command strings*

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

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

Getting Out

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

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

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

Administration Console Interface Parameters

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

Adjusting the Screen Height

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



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

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

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

Top-Level Menu

```

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

```

To set the screen height:

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

```
system screenHeight
```

You are prompted for a screen height value.

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

Example:

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

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

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

Example:

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

Disabling the Reboot and Abort Keys

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



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

To enable or disable the reboot and abort control keys:

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

```
system ctlKeys
```

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

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

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

Top-Level Menu

```

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

```

Remote Access Parameters

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

Preventing Disconnections

Because the Administration Console supports only a single shell at a time, you might be disconnected from your session if someone else remotely gains access to the Administration Console. The possible reasons for Console disconnections are listed in Table 2-2.

Table 2-2 Reasons for Console Disconnections

Access Method	Disconnected by...
Terminal through the serial port	Modem connection OR Telnet or rlogin connection
Telnet or rlogin	Modem connection

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



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

To lock the Administration Console:

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

```
system consoleLock
```

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

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

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

Top-Level Menu

```

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

```

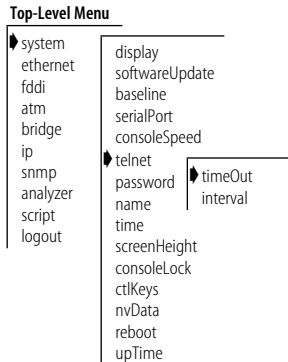
Enabling Timeout of Remote Sessions

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

To enable or disable the telnet timeout:

- 1 From the top level of the Administration Console, enter:
system telnet timeOut
- 2 Enter the telnet timeout state (**off** or **on**).

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

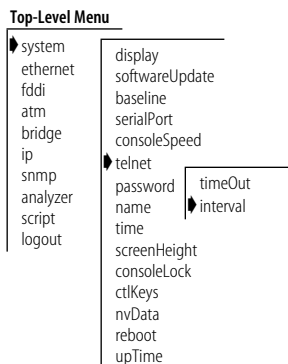


Setting Timeout Interval for Remote Sessions

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

To set the telnet timeout interval:

- 1 From the top level of the Administration Console, enter:
system telnet interval
- 2 Enter the telnet timeout interval (**30 minutes to 60 minutes**).



Running Scripts of Administration Console Tasks

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

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

Top-Level Menu

system
ethernet
fddi
atm
bridge
ip
snmp
analyzer
script
logout

To run a script:

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

script

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

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

The task you scripted is run in the Administration Console.

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

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


```
# This script performs some start-up configurations.
#
# Set the modem serial port baud rate.
#
system serialPort modemSpeed
300                # modem serial port baud rate
#
# Set the system name
#
system name
Engineering LANplex_4
#
# Assign an IP address to the LANplex.
#
ip interface define
158.101.112.99      # IP address for the system
255.255.0.0         # subnet mask
158.101.255.255    # broadcast address
1                   # cost
management         # type of interface
#
ip interface display
#
# Validate access to management workstation
#
ip ping
158.101.112.26     # management workstation address
#
# Enable the Spanning Tree Protocol
#
bridge stpState enabled
#
# Configure my node as an SNMP trap destination
#
snmp trap add
158.101.112.26     # management workstation address
all                # turn on all traps
q                  # no more trap destinations
#
snmp trap display
#
```

Getting Help in the Administration Console

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

Online Help

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

General online help

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

Help for specific menu options

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

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

```
? ethernet
```

Viewing More Levels of Menu Options

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

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

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

```
outline 2
```

Exiting from the Administration Console

If you are using an rlogin session to gain access to the system, exiting from it terminates the session. If you are accessing the system through the Console serial port, exiting returns you to the password prompt.

To exit from the Administration Console:


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

- 2 From the top-level menu, enter:

logout

Top-Level Menu

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


SYSTEM-LEVEL FUNCTIONS

Chapter 3 Configuring Management Access to the System

Chapter 4 Administering Your System Environment

Chapter 5 Baselineing Statistics

Chapter 6 Saving, Restoring, and Resetting Nonvolatile Data



3

CONFIGURING MANAGEMENT ACCESS TO THE SYSTEM

This chapter describes how to configure management access to the LANplex® 2500 system through a serial connection or an IP interface. It also describes how to configure the system so that you can manage it using the Simple Network Management Protocol (SNMP).

About Management Access

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

Using a Serial Connection

Direct access through the Console serial port is often preferred because you can stay attached during system reboots. You can also access the Administration Console through an external modem attached to the modem serial port.



See the LANplex 2500® Getting Started Guide for serial port pin-outs.

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

Using an IP Interface

An IP interface allows you to manage the system in-band through any Ethernet or FDDI port. Once an IP interface is configured, you can use rlogin or telnet to connect remotely to the Administration Console using the TCP/IP protocol from a host computer, or you can access the SNMP agent

from an external management application. The IP interface has a unique IP address.

In-band or Out-of-band?

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



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

Setting Up the Terminal Serial Port

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



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

To set the baud rate for the Console serial port:

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

```
system serialPort terminalSpeed
```

- 2 Enter the baud rate for the serial port.

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

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

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

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

Top-Level Menu

system	display	
ethernet	softwareUpdate	
fddi	baseline	
atm	serialPort	terminalSpeed
bridge	telnet	modemSpeed
ip	password	connectModem
snmp	name	
analyzer	time	
script	screenHeight	
logout	consoleLock	
	ctlKeys	
	nvData	
	reboot	
	upTime	

Setting Up the Modem Serial Port

For the modem serial port, you can set the port speed to match your external modem baud rate and then configure the external modem by establishing a connection between your current Console session and the modem serial port.

Setting the Port Speed

The default baud rate for the modem serial port is 9600. You might need to change the baud rate to match your external modem's baud rate.

To set the baud rate for the modem serial port:

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

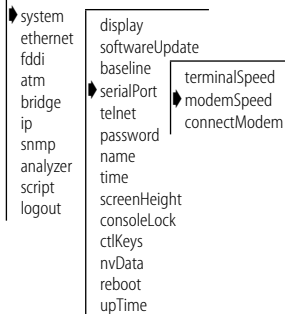
```
system serialPort modemSpeed
```

- 2 Enter the baud rate for the modem serial port.

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

The modem serial port baud rate is immediately changed.

Top-Level Menu



Configuring the External Modem

When you have set up the external modem from the Administration Console, characters you enter at the Console are transmitted as output on the modem port, and characters received as input on the modem port are echoed as output to the current Console session. Therefore, the Console appears to be directly connected to the external modem.

To configure the modem port:

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

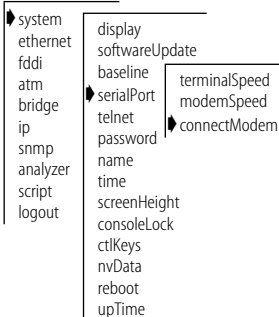
```
system serialPort connectModem
```

You can now issue the commands that support whatever communication parameters are appropriate to your installation. All characters entered in the Administration Console are transmitted to the modem port until you type the *escape sequence*.

- 2 When the modem is configured, enter the escape sequence (`~1`) with no intervening characters.

Entering the escape sequence breaks the connection to the modem serial port and returns to the previous menu.

Top-Level Menu



Setting Up an IP Interface for Management

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

General Setup Process

You must first define an interface, which includes assigning an IP address to that interface, and then ping your IP management station to be sure your system can reach it.

You finish your IP setup by checking that the following configurations are correct for your network and changing them as necessary:

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

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

Administering Interfaces

You define interfaces to establish the relationship between the ports on your system and the subnets in your IP network. You can have up to 32 IP interfaces for management per system.

An IP interface has the following associated information:

- **IP Address**

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

- **Subnet Mask**

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

■ Advertisement Address

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

■ Cost

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

■ State

The state indicates the availability of communications for this management interface. A value of "Up" indicates that one or more ports have link status. "Down" indicates no ports with link status.

Displaying Interfaces

You can display information about all IP interfaces configured for the system.

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

```
ip interface summary
```

OR

```
ip interface display
```


The system displays the current configuration. The display for both commands contains IP forwarding and RIP information as well as the IP interface information.

Example:

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

Index	Type	IP address	Subnet mask	Cost	State
1	management	158.101.1.1	255.255.255.0	1	Up
2	management	158.101.4.1	255.255.255.0	12	Up
3	management	158.101.6.1	255.255.255.0	15	Up
4	management	158.101.8.1	255.255.255.0	18	Up

Top-Level Menu



Defining an IP Interface

When you define an interface, you define the interface's IP address, subnet mask, advertisement address, cost, and the interface type.

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

Table 3-1 Recommended Settings for IP Management Access

Parameter	Recommended Setting
Type	management
IP address	<User defined>
Subnet mask	<User defined>
Advertisement address	Directed (all 1s in the host field)
Cost	1
State	Up or Down

To define an IP interface:

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

ip interface define

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

- 2 Enter the IP address of the interface.
- 3 Enter the subnet mask of the network to which you want to connect to the interface.
- 4 Enter the up to seven advertisement address to be used on the interface.
- 5 Enter the cost value of the interface.
- 6 Press [Return] to select management as the type of interface to configure.

Example:

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter advertisement address [158.101.1.255]:
```

Top-Level Menu

```
system
ethernet
fdi
atm
bridge
ip
snmp
analyzer
script
logout
  interface
    route
    arp
    rip
    ping
    statistic
    summary
    detail
    define
    modify
    remove
```

```
Enter cost [1]:
Enter interface type (management) [management]:
```

Modifying an IP Interface

To modify an IP interface that you have already defined:


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

```
ip interface modify
```

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

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

Top-Level Menu



Removing an Interface

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


To remove an IP interface definition:

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

```
ip interface remove
```

- 2 Enter the index numbers of the interfaces you want to remove.

Top-Level Menu



Administering Routes

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

Each routing table entry contains the following information:

- **Destination IP Address and Subnet Mask**

These elements define the address of the destination network, subnet, or host. A route matches a given IP address if the bits in the IP address that correspond to the bits set in the route subnet mask match the route

destination address. When it forwards a packet, if the system finds more than one routing table entry matching an address (for example, a route to the destination network and a route to the specific subnet within that network), it will use the most specific route (that is, the route with the most bits set in its subnet mask).

- **Routing Metric**

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

- **Gateway IP Address**

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

- **Status**

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


Table 3-2 Route Status

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

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

Displaying the Routing Table

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

Top-Level Menu

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

```
ip route display
```

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

```
IP routing is enabled, RIP is active,
ICMP router discovery is disabled.
Destination      Subnet mask      Metric      Gateway          Status
158.101.4.0      255.255.255.0    2           158.101.2.8      Static
158.101.3.0      255.255.255.0    2           158.101.1.2      Learned(RIP)
158.101.2.0      255.255.255.    1           --               Direct
158.101.1.0      255.255.255.0    1           --               Direct
Default Route    --                5           158.101.1.2      Learned (RIP)
```

Defining a Static Route

You might want to define a static route to transmit system traffic, such as system pings or SNMP response, through a consistent route. Before you define static routes, you must define at least one IP interface. (See “Defining an IP Interface” on page 3-6.) Static routes remain in the table until you remove them, or until you remove the corresponding interface. Static routes take precedence over dynamically learned routes to the same destination.

To define a static route:

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


```
ip route static
```

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

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

A static route is defined in the following example:

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

Top-Level Menu


Removing a Route

To remove a route:

- 1 From the top level of the Administration Console, enter:
ip route remove
- 2 Enter the destination IP address of the route.
- 3 Enter the subnet mask of the route.

The route is immediately deleted from the routing table.

Top-Level Menu



Flushing All Learned Routes


Flushing deletes all learned routes from the routing table.

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

ip route flush

All learned routes are immediately deleted from the routing table.

Top-Level Menu



Setting the Default Route

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


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

To statically configure the default route:

- 1 From the top level of the Administration Console, enter:
ip route default
- 2 Enter the gateway IP address of the route.

The default route is immediately added to the routing table.

Top-Level Menu




Removing the Default Route

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

```
ip route noDefault
```

The default route is immediately removed from the routing table.



Administering the ARP Cache

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

Displaying the ARP Cache

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

```
ip arp display
```

The system displays the contents of the ARP cache as shown in this example:

```
RIP is active.
```

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

Removing an ARP Cache Entry

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



To remove an entry from the ARP cache:

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

```
ip arp remove
```

- 2 Enter the IP address you want to remove.

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



Flushing ARP Cache Entries


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

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

```
ip arp flush
```

The ARP cache entries are immediately removed from the table.

Top-Level Menu



Setting the RIP Mode

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

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

RIP default mode

By default, RIP operates in *passive* mode.

To set the RIP operating mode:

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


```
ip rip
```

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

Example:

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

Top-Level Menu



Pinging an IP Station

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

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

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

To ping an IP station:

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

```
ip ping
```

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

```
IP Address: 192.9.200.40
```

There are two possible responses:

```
192.9.200.40 is alive
```

OR

```
no answer from 192.9.200.40
```

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

```
Network is unreachable
```

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

Top-Level Menu

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

Displaying IP Statistics

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

Table 3-3 IP Statistics

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

Top-Level Menu

```

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

```

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

ip statistics

Statistics are displayed, as shown in this example:

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

```

inReceives      forwDatagrams      inDelivers      outRequests
          51213              49743              3227              2285

outNoRoutes      inHdrErrors      inAddrErrors
          273                7                0

```

Setting Up SNMP on Your System

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

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

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

For more information about using SNMP to manage the LANplex system, see Chapter 3: *Management Access: Protocols* in the *LANplex® 2500 Operation Guide*.

Displaying SNMP Settings

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

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

```
snmp display
```

The community string settings are displayed as shown here:

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

Top-Level Menu

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

```
display
community
trap
```

Configuring Community Strings

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

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

When an SNMP agent receives an SNMP request, the agent compares the community string in the request with the community strings configured for the agent. SNMP *get*, *get-next*, and *set* requests are valid if the community string in the request matches the agent’s *read-write* community. Only the

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

Community string length

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

To set a community string:

Top-Level Menu

```

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

```

display
community
trap
  
```

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

```
snmp community
```

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

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

Administering SNMP Trap Reporting

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

Displaying Trap Reporting Information

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

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

```
snmp trap display
```

Top-Level Menu

```

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

```

display
community
trap
  
```

```

display
addModify
remove
flush
smtProxyTraps
  
```

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

Trap Descriptions:

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

Trap Destinations Configured:

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

Proxying of remote SMT events is disabled




Trap 8: Power Supply Failure and Trap 13: System Fan Failure appear in the trap list only if the LANplex® 2500 has a dual power supply.

Configuring Trap Reporting

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

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

Top-Level Menu



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

```
snmp trap addModify
```

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

- 2 Enter an IP address of the SNMP manager (destination address).
You are prompted for the trap numbers to enable for that destination.

- 3 Enter the trap number(s).

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



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

This example shows a trap configuration:

```
Enter the trap destination address: 158.101.222.3
Enter the trap numbers to enable (1-8,12-24|all)
[1-8,12-24|all]: all
```

Address Error

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

```
Trap address invalid or unreachable
```

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

Removing Trap Destinations

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


To remove a destination:

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

```
snmp trap remove
```

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

- 2 Enter the SNMP trap reporting destination address you want to remove. The destination address is removed and you return to the previous menu.



Flushing All SNMP Trap Destinations

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

To flush all SNMP trap reporting destinations:

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


```
snmp trap flush
```

You receive the following prompt:

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

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

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



Setting Up SMT Event Proxying

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



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

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

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

- On only one LANplex system, 1) enable local SNMP traps as described in the “Configuring Trap Reporting” on page 3-18 and 2) enable proxying of remote SMT events. On all other LANplex systems in your network, 1) disable proxying of remote SMT events and 2) enable only SNMP traps that are *not* SMT-related. SMT-related traps include all of those in the LANplex Optional FDDI MIB. This configuration provides access to the events occurring locally on the one LANplex system and to those reported by other stations on the FDDI ring (including other LANplex systems).
- Enable local SNMP traps and disable the proxying of remote SMT events on every LANplex system in your network. Local traps will be reported to the management station (which will cover all your LANplex systems), but SMT events from systems other than LANplex systems in your network will not be reported.


To enable or disable the proxying of remote SMT events:

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

```
snmp trap smtProxyTraps
```

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

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



4

ADMINISTERING YOUR SYSTEM ENVIRONMENT

This chapter focuses on the administration of your LANplex® system environment, which involves:

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

Displaying the System Configuration

The system configuration display provides software and hardware revisions, module status information, and warning messages for certain system conditions.

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

```
system display
```

Example of a system configuration display:

```
LANplex 2500 (rev 8.4) - System ID 0f2b00
Intelligent Switching Software
Version 8.1.0 - Built 8/22/96 06:26:55 PM

  Slot      Contents
-----
  1         FDDI Module (rev. 4.2)
  2         ATM Module (rev. 2.3)
  3         Ethernet Module (rev. 5.1)
  4         Ethernet Module (rev. 5.1)
```

The display contains the following general system information:

- The system type (LANplex 2500)
- System ID

Top-Level Menu

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

- Software version
- Software build date and time

Module entries The display also has one entry for each module in the system, listing:

- Type of module (Ethernet, FDDI, ATM)
- Hardware revision number of each module

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

- System temperature over the maximum level for normal operation
- Fan failure
- Power supply failure (if redundant power supply is present)

Setting Passwords

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

Initial passwords Because the initial passwords stored in the nonvolatile memory of the system are null for all levels, press [Return] at the password prompt.

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

To set a password:

- 1 From the top level of the Administration Console, enter:
system password
- 2 At the prompt that requests you to enter a password access level to change, enter one of the following:
read
write
administer
- 3 At the prompt for your old password, enter the old password.
- 4 Enter the new password.

Top-Level Menu

```

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

```

The password can have up to 32 characters and is case-sensitive. To enter a null password, press [Return].

- 5 Retype the new password for verification. The system does not display the password as you type.

Example:

```
Select menu option (system): password
Password access level (read,write,administer): read
Old password:
New password:
Retype new password:
The administration console password has been successfully
changed.
```

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

Setting the System Name

You should give the LANplex system an easily recognizable and unique name to help you manage the system. For example, you might want to name the system according to its physical location (example: LP2500-ENGLAB).

To name the system:

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

system name

You are prompted for the name of the system:

```
Enter new string (no spaces) [LANplex]:
```

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

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

Top-Level Menu

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

Changing the Date and Time

Top-Level Menu

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

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

To change either the date or the time:

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

```
system time
```

The system displays the current date and time, along with a prompt asking whether you want to change the time. Example:

```
The current system time is 08/24/96 04:37:57 PM.
Do you want to change the system time (n,y) [y]:
```

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

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

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

Table 4-1 Date and Time Variables

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

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

Example:

```
Enter the new system time (mm/dd/yy hh:mm:ss xM): 09/30/96
10:00:00 AM
Press RETURN at the exact time:
```

Rebooting the System

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

To reboot the system:

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

```
system reboot
```

The following message appears:

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

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

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

Top-Level Menu

```

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

```

Displaying the System Up Time

You can display the amount of time the system has been running since the last reboot.

To display the system up time, enter the following command from the top level of the Administration Console:

```
system upTime
```

The administration console displays the amount of time the system has been running in days and hours. Example:

```
System up time -
                Hours: 22
                Minutes: 26
```

Top-Level Menu

```

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

```



5

BASELINING STATISTICS

This chapter describes how baselining statistics work in the LANplex® system, and how to set, display, enable, or disable a baseline statistic.

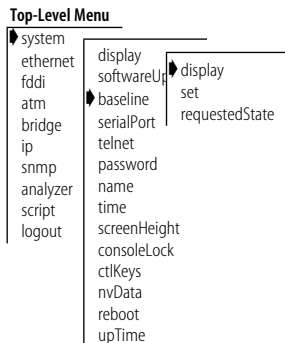
About Setting Baselines

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

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

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

Displaying the Current Baseline



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

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

```
system baseline display
```

Example:

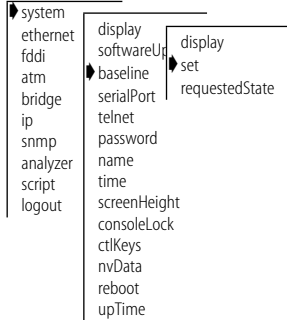
```
Baseline set at 08/28/96 10:42:52 AM is currently enabled.
```

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

```
A baseline has not yet been set.
```

Setting Baselines

Top-Level Menu



Setting a baseline resets the counters to 0. The accumulated totals since power up are maintained by the system. The baseline is time-stamped.

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

```
system baseline set
```

A message similar to the following display appears:

```
Baseline set at 08/28/96 10:42:52 AM.
```

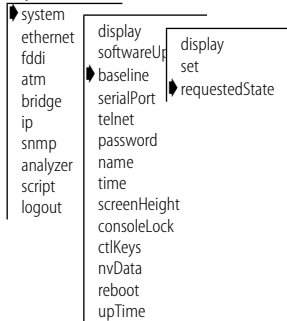
Baselining is automatically enabled when a baseline is set.

Enabling or Disabling Baselines

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

To enable the current baseline:

Top-Level Menu



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

```
system baseline requestedState
```

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

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

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

The new value is confirmed as shown here:

```
Baseline set at 08/28/96 10:42:52 AM has been disabled.
```

6

SAVING, RESTORING, AND RESETTING NONVOLATILE DATA

This chapter describes the nonvolatile (NV) data in the LANplex® system and how to save, restore, and reset the data.

Working with Nonvolatile Data

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

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

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

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

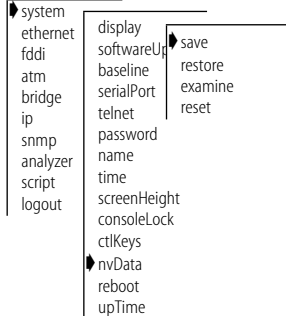
- Software version number
- System ID
- Date and time of creation
- Data checksums

Saving NV Data

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

To save NV data:

Top-Level Menu



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

```
system nvData save
```

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

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

Example:

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

If the information is incorrect or if a connection cannot be made with the specified host, a message similar to this one appears:

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

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

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

Failure Error - Configuration not stored.

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

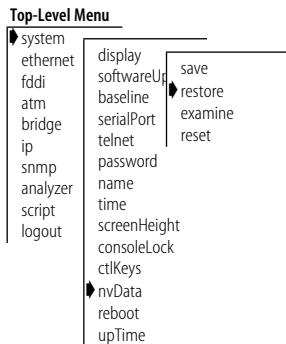
At the end of the save, the system displays to the previous menu.

Restoring NV Data

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

- Rule 1 Exact Match* — The system IDs, module types, and module revisions (if applicable) all match between the saved configuration and the system on which you are restoring the image.
- Rule 2 System ID Mismatch* — System IDs do not match between the saved NV file and the target system. Mismatches in system IDs are allowed. Before restoring the NV data to a system with a different system ID, however, be aware of the following NV data that might cause problems when restored:
- Management IP addresses (defined in IP interface configurations) are saved as NV data and restored. To avoid duplicate IP address problems, before connecting the restored system to the network, you might need to change the IP address of defined interfaces. Modifying IP interface definitions is described on page 3-7.
 - Statically configured Ethernet addresses are saved as NV data. Be sure not to have duplicate addresses when you restore the NV data. Listing statically configured addresses is described on page 12-11.

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



To restore the NV data:

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

```
system nvData restore
```

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

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

If the information is incorrect, or if a connection cannot be made with the specified host, a message similar to this one appears:

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

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

You are prompted to load the proposal.

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

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

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

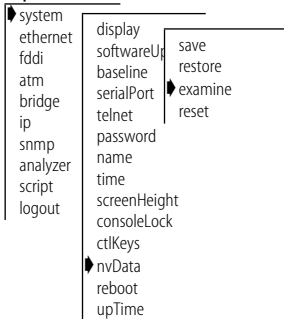
If you enter **y**, the system NV data is restored as proposed.

If you enter **n**, the restoration fails and the previous menu appears on the screen.

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

Examining a Saved NV Data File

Top-Level Menu



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

To examine the file:

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

```
system nvData examine
```

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

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

If the information is incorrect, or if a connection cannot be made with the specified host, a message similar to this one appears:

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

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

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

The system then displays the NV data menu options.

Resetting NV Data to Defaults

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



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

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

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

```
system nvData reset
```

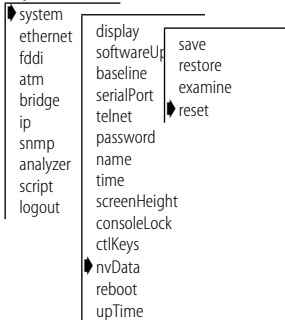
You see the following prompt:

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

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

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

Top-Level Menu






ETHERNET, FDDI, AND ATM PARAMETERS


Chapter 7 Administering Ethernet Ports

Chapter 8 Administering FDDI Resources

Chapter 9 Administering ATM

Chapter 10 Setting Up the System for Roving Analysis





ADMINISTERING ETHERNET PORTS

This chapter describes how to:

- View Ethernet port information
- Use full-duplex mode with Fast Ethernet ports
- Configure Ethernet port labels
- Enable or disable an Ethernet port

Displaying Ethernet Port Information

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

To display Ethernet port statistics relative to a baseline, see Chapter 5.

To display information about the Ethernet ports:

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


```
ethernet summary
```

OR

```
ethernet detail
```

- 2 Enter the port numbers for which you want to view information.

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



```

port          rxFrames          rxBytes          rxFrameRate      rxByteRate
  1             406430           36336795         0                 0
 12            242400           29275605         0                 0

port  rxPeakByteRate  rxPeakFrameRate  noRxBuffers      alignmentErrs
  1             90484            163              0                 0
 12            58438            394              0                 0

port          fcsErrs          lengthErrs      rxInternalErrs    rxDiscards
  1              0                0               0                 0
 12              0                0               0                 0

port          rxUnicasts      rxMulticasts      txFrames          txBytes
  1             365811         40619             1422085           234636091
 12            242033          367               1256455           300242671

port          txFrameRate      txByteRate      txPeakFrameRate  txPeakByteRate
  1              3                345             208               271724
 12              3                345             402               321722

port          txQOverflows  excessCollision  excessDeferrals    txInternalErrs
  1              0                0               0                 0
 12              0                0               0                 0

port  carrierSenseErr  txDiscards      txUnicasts      txMulticasts
  1              0                0             528268          893836
 12              0                0             322389          934076

port          collisions  lateCollisions  requestedState    portState
  1              0                0             enabled          on-line
 12              0                0             enabled          on-line

port          portType          linkStatus          macAddress
  1    10BaseT(RJ45)      enabled            00-80-3e-0b-48-02
 12    10BaseT(RJ45)      enabled            00-80-3e-0b-48-0d

port          portLabel          duplexMode
  1             Office113_SPARCstation5      n/a
 12             Office322_Quadra900          n/a

```

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

```

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

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

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

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

Table 7-1 Fields for Ethernet Port Attributes

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

(continued)

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

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

(continued)

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

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

*Frame processing and
Ethernet statistics*

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

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

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

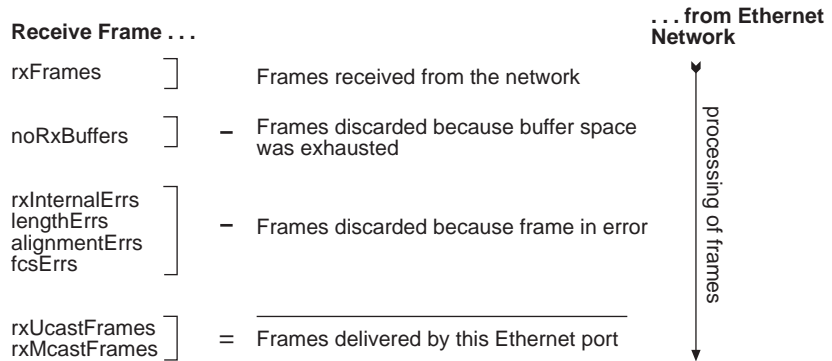


Figure 7-1 How Frame Processing Affects Ethernet Receive Frame Statistics

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

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

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

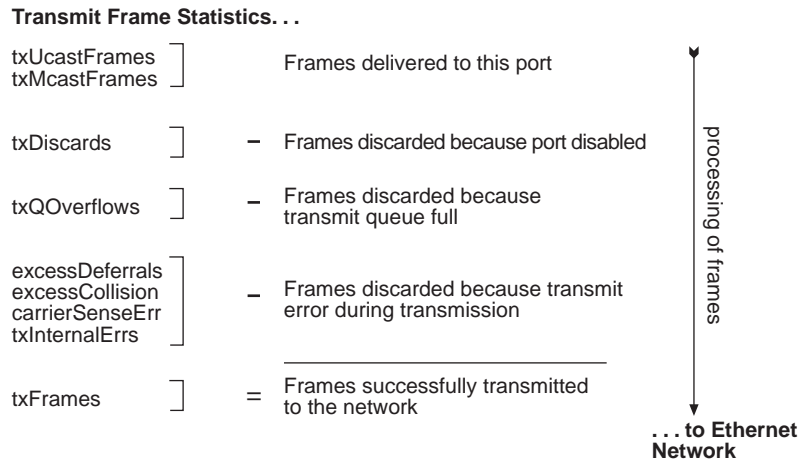


Figure 7-2 How Frame Processing Affects Ethernet Transmit Frame Statistics

Enabling or Disabling Full-Duplex Mode

The LANplex® 2500 Fast Ethernet 100BASE-TX and 100BASE-FX port is capable of operating in full-duplex mode. Full-duplex mode can enhance the throughput of the LANplex Fast Ethernet connection by allowing data to be transmitted and received simultaneously.



CAUTION: *When configuring full-duplex mode, you must configure both the sending and the receiving devices.*

To enable or disable full-duplex mode:

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

ethernet duplexMode

Top-Level Menu

- system
- ethernet
 - summary
 - detail
 - duplexMode
 - label
 - portState
- fdi
- atm
- ip
- snmp
- analyzer
- script
- logout

You are prompted for the port number(s).

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

After you have selected the port(s), the following warning message is displayed on the screen:

Warning – Changing mode to full duplex disables collision detection. The device connected to this port must be configured for the same duplex mode.

You are prompted to confirm that you want to change the duplex mode.

- 3 Enter **y** (yes) or **n** (no) at the prompt. The default is **y**.

If you enter **n**, you are returned to the main menu. If you enter **y**, you are prompted for the new duplex mode value.

- 4 Enter **full** to enable full-duplex mode. The default is *half*.

If you have not previously configured connected devices, follow steps 1 through 4 to enable full-duplex.

Labeling a Port

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

To label an Ethernet port:

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

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

Top-Level Menu

system	
ethernet	summary
fdi	detail
atm	duplexMode
bridge	label
ip	portState
snmp	
analyzer	
script	
logout	

Setting the Port State

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

To enable or disable an Ethernet port:


- 1 From the top level of the Administration Console, enter:
ethernet portState
- 2 Enter the numbers of the ports you want to enable or disable.
- 3 Enter **enabled** or **disabled** for each Ethernet port.

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

Top-Level Menu

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





ADMINISTERING FDDI RESOURCES

This chapter describes how to display information about and configure the LANplex® 2500 system and its:

- FDDI stations
- FDDI paths
- Media Access Controls (MACs)
- FDDI ports



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

For more information about FDDI in the LANplex system, see the *LANplex® 2500 Operation Guide*.

Administering FDDI Stations

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


You can display station information and set the following parameters:

- Connection policies
- Neighbor notification timer
- Status reporting

Displaying Station Information

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

Top-Level Menu



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

```
fddi station display
```

You are prompted for a station.

- 2 Enter the ID or the station about which you want to view information.

Example station information:

```

configuration          tNotify statusReporting  connectPolicy
      isolated                30          enabled          0x8000

      ecmState remoteDisconnect  traceMaxExp
              in                false          87500000

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

Table 8-1 describes these statistics.

Table 8-1 Fields for FDDI Station Attributes

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

(continued)

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

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

Setting the Connection Policies

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

The LANplex 2500 FDDI ports can be of type A or type B. By default, all connections to the LANplex 2500 FDDI ports are valid, except for M-M connections. The possible connections to reject and their corresponding bits are listed in Table 8-2.

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


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

(continued)

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

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

To set the connection policies of an FDDI station:

Top-Level Menu

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

```
fddi station connectPolicy
```

You are prompted for a station.

- 2 Enter the ID of the station for which you want to set the connection policies.
- 3 Enter the value of the connection policy for that station.

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

Example:


```
Select Fddi station [1]:
Enter new value [8000]:
```


Setting Neighbor Notification Timer

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

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

To set the *T-notify* timer:



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

```
fddi station tNotify
```

You are prompted for a station.

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

Valid values are 2–30 seconds.


Example:

```
Select Fddi station [1]:
Enter new value [30]:
```

Enabling and Disabling Status Reporting

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

To enable or disable status reporting for a station:



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

```
fddi station statusReporting
```

You are prompted for a station.

- 2 Enter the station ID number.

- 3 Enter the new statusReporting value (**enabled** or **disabled**).

See the following example:

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

Administering FDDI Paths

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

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

- tvxLowerBound
- tmaxLowerBound
- maxTreq

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

Displaying Path Information

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

To display FDDI path information:


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

```
fddi path display
```

You are prompted for a station and path.

- 2 Enter the station about which you want to view information.

Top-Level Menu



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

Example of path information:

```

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

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

```

Table 8-3 describes these statistics.

Table 8-3 Description of Fields for FDDI Path Attributes


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

Setting tvxLowerBound

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

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

Top-Level Menu



To set tvxLowerBound:

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

```
fddi path tvxLowerBound
```

You are prompted for a station, path, and value.

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

The default is 2500 microseconds (μ s).

See the following example:


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

Setting tmaxLowerBound

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

To set tmaxLowerBound:

Top-Level Menu



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

```
fddi path tmaxLowerBound
```

You are prompted for a station, path, and value.

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

The default is 165000 microseconds (μ s).

See the example below:

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


Setting maxT-Req

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

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

To set maxT-Req:

Top-Level Menu



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

```
fddi path maxTreq
```

You are prompted for a station number, path, and value.

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

The default value is 165000 microseconds (μ s)

Example:

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

Administering FDDI MACs

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


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

Displaying MAC Information

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

To view the FDDI MAC summary or detailed statistics:

Top-Level Menu



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

```
fddi mac summary
```

OR

```
fddi mac detail
```

You are prompted for a MAC number.

- 2 Enter the MAC number.

Example summary display of FDDI MAC information:

```

location          currentPath      station
      n/a                primary          n/a

rxFrames          txFrames        rxBytes
      19                  17              1619

txBytes          Errors          noRxBuffers
      1437                0                0

rxErrors          txQOverflows    smtAddress
      0                    0              08-00-02-10-b4-1e

upstream          downstream
08-00-02-10-9e-0a 08-00-02-10-9e-0a
  
```

Example detail display of FDDI MAC information:

rxFrames	rxBytes	rxFrameRate
18	2102	0
0	0	0
rxByteRate	rxPeakFrameRate	rxPeakByteRate
0	5	526
0	0	0
lostCount	lateCount	notCopiedCount
0	0	0
0	0	0
notCopiedThresh	notCopiedRatio	notCopiedCond
6550	0	inactive
6550	0	inactive
errorCount	frameErrThresh	frameErrorRatio
0	655	0
frameErrCond	noRxBuffers	rxErrors
inactive	1051	1051
inactive	389	389
tvxExpiredCount	rxInternalErrs	rxDiscards
0	0	10
0	0	0
rxUnicasts	rxMulticasts	txFrames
4	4	11
0	0	5
txBytes	txFrameRate	txByteRate
1051	0	0
389	0	0
txPeakFrameRate	txPeakByteRate	txInternalErrs
3	263	0
1	81	0
txQOverflows	txDiscards	txUnicasts
0	0	4
0	0	0
txMulticasts	frameCount	tokenCount
7	1746	66092634
6	34	518145081
ringOpCount	currentPath	dupAddrTest
3	primary	passed
2	primary	none
duplicateAddr	upstreamDupAddr	llcAvailable
false	false	true
false	false	true
llcService	smtAddress	upstream
enabled	00-80-3e-10-b7-16	00-80-3e-10-b7-01
enabled	00-80-3e-10-b7-17	unknown
downstream	oldUpstream	oldDownstream
00-80-3e-10-b7-40	unknown	unknown
unknown	unknown	unknown
downstreamType	rmtState	tMaxCapab
unknown	ring op	1342200 us
unknown	ring op	1342200 us
tvxCapab	tReq	tNeg
1342200 us	164986 us	164986 us
1342200 us	164986 us	164986 us
tMax	tvxValue	
167770 us	2621 us	
167770 us	2621 us	

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

Table 8-4 Fields for FDDI MAC Attributes

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

(continued)

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

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

(continued)

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

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

*Frame Processing and
FDDI MAC Statistics*

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

- There is no buffer space available.
- The frame is in error.
- LLC service is disabled.
- This is an NSA Frame and the A-bit is set.

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

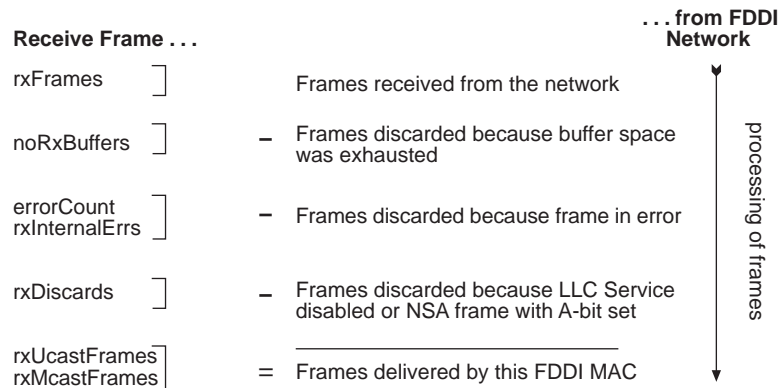


Figure 8-1 How Frame Processing Affects FDDI MAC Receive Frame Statistics

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

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

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

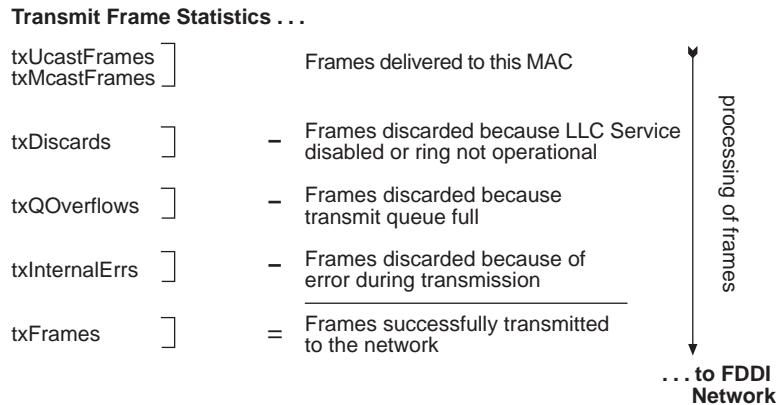



Figure 8-2 How Frame Processing Affects FDDI MAC Transmit Frame Statistics

Setting the Frame Error Threshold

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

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

To set the *FrameErrorThreshold* attribute:



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

```
fddi mac frameErrorThreshold
```

You are prompted for a MAC number and new threshold value.

- 2 Enter the MAC number.
- 3 Enter the new threshold value.

See the following example:

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

Setting the Not Copied Threshold

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

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

To set the *NotCopiedThreshold*:

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


```
fddi mac NotCopiedThreshold
```

You are prompted for a MAC number and new threshold value.

- 2 Enter the MAC number.
- 3 Enter the new threshold value.

Example:

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




Enabling and Disabling LLC Service

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

To enable or disable LLC service for a MAC:

Top-Level Menu



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

```
fddi mac llcService
```

You are prompted for a MAC number and to enable or disable LLC service.

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

Example:

```
Select Fddi MAC [1]:
```


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

Setting the MAC Paths

The possible backplane path assignments for MACs include primary and secondary.

To assign MACs to paths:

Top-Level Menu



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

```
fddi mac path
```

You are prompted for a MAC number and a path assignment for the MAC.

- 2 Enter the MAC number.
- 3 Enter the path.

Administering FDDI Ports

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

- lerAlarm
- lerCutoff
- port labels
- port paths

Displaying Port Information

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

To view FDDI port information:

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

```
fddi port display
```

You are prompted for a port.

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

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

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

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

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

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

Top-Level Menu




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

Table 8-5 Fields for FDDI Port Attributes

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

Setting lerAlarm

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


receiving alarms if your network is in poor health. The SMT Standard recommended value is 8.



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

To set the `lerAlarm` value:

Top-Level Menu



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

fddi port lerAlarm

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

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

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


Setting `lerCutoff`

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



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

Top-Level Menu



To set the *lerCutoff* value:

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

```
fddi port lerCutoff
```

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

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


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

Setting Port Labels

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

To label an FDDI port:

Top-Level Menu



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

```
fddi port label
```

You are prompted for a port number and a label string of up to 32 characters.

- 2 Enter the port number.
- 3 Enter the label character string (no spaces).

Setting the Port Paths

In the LANplex 2500 system, you can assign the A and B ports to either the primary or the secondary path.

To assign ports to paths:


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

```
fddi port path
```

You are prompted for a port number.

- 2 Enter the number of the port you want to configure.
- 3 Select the DAS configuration **isol** or **thru** for peer mode at the prompt.
- 4 Select the DAS configuration **isol**, **wrapAB**, or **dualHome** for tree mode at the prompt.

Top-Level Menu



9

ADMINISTERING ATM

This chapter describes how to administer asynchronous transfer mode (ATM) communications on the LANplex® 2500 system. It includes information about:

- LAN Emulation (LANE)
- UNI Management Entity (UME). UNI is a User to Network Interface.
- ATM ports

ATM in Your Network

ATM architecture differs fundamentally from IEEE 802.x technology. IEEE 802.x LANs, such as FDDI, Ethernet, and token ring, are connectionless and use the Media Access Control (MAC) addresses in each packet to communicate to end-stations. ATM is connection-oriented and uses a circuit identifier, called a virtual channel identifier, to exchange data between two ATM stations over a previously established virtual channel connections (VCC).

LAN Emulation and Classical IP

To forward data over an ATM interface in an existing 802.x LAN network, two methods are provided to adapt existing data link layer protocols to the connection-oriented paradigm of ATM. These methods are **LAN Emulation (LANE)** and **Classical IP over ATM**.

- The LANE method of adapting networks to ATM supports transparent translation of higher level protocols, such as IP, IPX, and AppleTalk. LANE also supports broadcast and multicast addressing. For more information on LANE on the LANplex 2500 system, see page 9-2.
- The Classical IP method supports transparent translation of IP only over ATM, and does not support broadcast or multicast addressing. Classical IP is supported in Extended Switching software only. For more information, see the *LANplex® 2500 Extended Switching Guide*.

LAN Emulation

LAN Emulation (LANE) emulates connectionless network technologies by providing unicast, multicast, and broadcast services over connection-oriented ATM. An emulated LAN (ELAN) can consist of many LECs. An ELAN consists of the following components:

- **A Broadcast and Unknown Server (BUS)**

The BUS is responsible for handling broadcast, multicast, and initial unicast frames sent from a LAN Emulation Client. Each ELAN contains only one BUS.

- **A LAN Emulation Server (LES)**

The LES is responsible for guaranteeing MAC address uniqueness and resolving MAC addresses to ATM addresses for LECs. Each ELAN contains only one LES.

- **LAN Emulation Clients (LECs)**

The LEC is the end node from the perspective of the ATM network. It performs data forwarding, address resolution and other control functions. Additionally, it maintains the LAN emulation software.

A LES and BUS must be defined as part of an ELAN before a LEC may be defined.

- **A LAN Emulation Configuration Server (LECS)**

The LECS provides configuration information about the ATM and LAN networks. It also provides the address of the LES to the LEC.

The LES and the BUS can be configured on the same LANplex or on different LANplex.

Before You Configure an ELAN

Before you configure an LAN emulation network, complete the following procedures:

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

Checking Link Status

To check link status:

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

```
atm ume display
```

The interface prompts you to enter the numbers of the UME you want to display.

- 2 Enter the numbers of the UMEs, or **all**.

Example:

```
    reqState      state      reqVpi
connected      connected      0
    vpi          reqVci      vci
    0            16          16
    connCount    discCount    rxPdus
    2            2            232
    dropRxDpus   txPdus
    0            209
```

The status for the state field should be **connected**. If the state is **disconnected**, a problem exists on the link. If this occurs, check the cabling to ensure that the system is correctly connected to the ATM switch.

Verifying Address Registration

When you are sure the link is connected, verify that the LANplex system can register addresses at the network side and the user side for the UME Network Interface (UNI). Two addresses are registered on the network side: one for LAN emulation and one for classical IP (CLIP) or ATM.

To verify registered addresses:

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

```
atm ume list
```

The interface prompts you for the numbers of the ATM ports you want to display.

Example:

```
Select Atm port(s) (1-2|all) [1]:

Port 1:
Prefixes registered at the user side UNI:
47-0005-80-ffe100-0000-f21a-200e

Addresses registered at the network side UNI:
47-0005-80-ffe100-0000-f21a-200e-00803e1f6712-00
47-0005-80-ffe100-0000-f21a-200e-00803e1f6713-00
```

Verifying Signaling

The ATM signaling protocol allows end stations to establish, maintain, and clear ATM connections between endpoints through virtual connections (VCCs). One or more connections can exist on a physical link. Two types of virtual connections exist: virtual path identifiers (VPIs) and virtual channel identifiers (VCIs).

To verify that UNI signaling is operational:

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

```
atm ports vcc list
```

In addition to newly defined VCCs, two additional VCCs appear in the display. One standard VPI and VCI are reserved for UNI signaling and ILMI registration. UNI signaling uses VPI 0/VCI 5. ILMI registration uses VPI 0/VCI 16.

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 LEC on the network, or on an ATM switch, such as 3Com's CELLplex™ 7000 system.
- 2 Define the Broadcast and Unknown Server.
- 3 Define the LAN Emulation Server.
- 4 Define the LAN Emulation Clients.



The LANplex 2500 system supports Unspecified Bit Rate (UBR) only.

Configuring Clients to Join an Existing Emulated LAN

You can configure a LAN Emulation Client (LEC) to join an existing 802.3 emulated LAN by providing information about the LAN Emulation Server (LES) and the Broadcast and Unknown Server (BUS). See “Defining LAN Emulation Clients” on page 9-10 for information on defining a LEC for inclusion in the emulated LAN.

Administering LECs (LAN Emulation Clients)


From the Administration Console you can:

- Display summary or detailed information about emulated LAN clients
- Modify information for emulated LAN clients
- Define a client for inclusion in an emulated LAN
- Remove a client from the network

Displaying Information About LAN Emulation Clients

You can display a summary report or a detailed report of information about LAN Emulation Clients. The summary displays information about the location and state and the most important statistics about the client’s general activity. The detailed includes the summary information plus additional statistics.

To display information about a LAN Emulation Client:



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

```
atm lane lec summary
```

or

```
atm lane lec detail
```

You are prompted for the number of a LEC.

- 2 Specify the number of the LAN Emulation Client about which you want information.

The information is displayed in the format you specified.

Summary display example:

```

location          bridgePort        lecState
  Port 1          Port 18        operational
  enabled          configSource    cfgMaxArpRtryCt
  enabled          manual          1
cfgArpRespTime    cfgMaxUnkFrmCt    arpAfterCt
  1                1                0
  inArpReqs        inArpRsps         outBusUcDscrds
  0                0                0
                  elanName          elanType
                  system1          802.3
elanMtu
  1516
                  lesAddr
47-0000-00-000000-0000-0000-00aa-08000210b37c-40
                  lecsAddr
00-0000-00-000000-0000-0000-0000-000000000000-00

```

Detail display example (see following page):

```

location                bridgePort              lecState
  Port 1                 Port 18                  operational

  enabled                configSource            cfgMaxArpRtryCt
  enabled                 manual                    1

  maxArpRtryCt           cfgArpRespTime         arpRespTime
    1                      1                          1

  cfgMaxUnkFrmCt         arpAfterCt              inArpReqs
    1                       0                          0

  inArpRsps              outBusUcdscrds
    0                        0

                                cfgElanName

                                elanName                  cfgElanType
                                system1                      802.3

  elanType                cfgElanMtu              elanMtu
  802.3                    1516                      1516

  cfgLesAccessType        lesAccessType           cfgConnCompTime
  manual                    manual                        4

  connCompTime           cfgCtrlTimer            controlTime
    4                       120                          120

  cfgFlushRspTime        flushRspTime            maxUnkFrmTime
    4                       4                          1

                                maxUnkFrmTim          cfgPathSwDlyTime          pthSwtDlyTim
    1                       6                          6

  cfgVccTime              vccTime                 topologyChgFlag
  1200                     1200                        0

  cfgFwdDlyTime           fwdDelayTime            cfgAgeTime
    15                      15                          300

  ageTime                 inCtrls                 outCtrls
    300                     1                          1

  inDds                   outDds                  inBus
    0                       0                          0

  outBus                  outArpReqs              outArpRsps
    0                       0                          0

  macAddress
00-80-3e-20-98-12

                                busAddress
                                47-0000-00-000000-0000-0000-00aa-08000210b37c-80

                                cfgLesAddr
                                47-0000-00-000000-0000-0000-00aa-08000210b37c-40

                                lesAddr
07-0000-00-000000-0000-0000-00aa-08000210b37c-40

                                lecAddress
                                47-0000-00-000000-0000-0000-00aa-08000210b37c-00

```


Table 9-1 describes the attributes of the LAN Emulation Client. References to the ATM Forum Specification for LAN Emulation are listed in parentheses after the parameter description.

Table 9-1 Fields for LAN Emulation Client Attributes

Parameter	Description
location	ATM port to which the client is connected
cfgElanType	Requested type of LAN to emulate (802.3)
ElanType	Type of LAN currently emulated (802.3)
cfgElanMtu	Requested maximum frame size
oprElanMtu	Current maximum frame size
lecState	Current client machine state
cfgLesAccessType	Requested method for determining how to access the LAN Emulation Server
oprLesAccessType	Current type of access to the LAN Emulation Server
enabled	Management state of the client (disabled or enabled)
cfgMaxArpRtryCt	Requested maximum number of LE_ARP attempts
cfgArpRespTime	Requested maximum time (in seconds) that the client expects between an LE_ARP request and an LE_ARP response (C20)
cfgConnCompTime	Requested time during which data or a READY_IND message is expected from a calling party (C28)
cfgCtrlTimer	Requested timeout period for request/response control frame interactions (C7)
cfgflushRspTime	Requested time limit for receiving an LE_FLUSH_RESPONSE after sending an LE_FLUSH_REQUEST (C21)
cfgMaxUnkFrmCt	Requested maximum unknown frame count (C10)
maxUnkFrmTime	Requested time period during which the client will send no more than the maximum number of unknown frames to the Broadcast Unknown Server (BUS) for a given destination (C11)
cfgPathSwDlyTime	Requested time since sending a frame to the BUS after which the client assumes that the frame has been discarded or delivered (C22)
cfgVccTime	Requested time after which the client should release any data direct VCC that did not transmit or receive data frames (C12)
topologyChgFlag	Boolean value indicating that the client is using Forward Delay Time rather than Aging Time to age nonlocal entries in its LE_ARP cache (C19)

Table 9-1 Fields for LAN Emulation Client Attributes (continued)

Parameter	Description
cfgFwdDlyTime	Requested maximum time that the client will maintain an entry for a nonlocal MAC address in its LE_ARP cache as long as the topologyChgFlag is true (C18)
cfgAgeTime	Requested maximum time that the client will maintain an entry in its LE_ARP cache (C17)
arpAfterCt	Number of unknown frames after which the client will send an LE_ARP request
inArpReqs	Number of LE_ARP requests received
inArpRsps	Number of LE_ARP responses received
inCtrls	Number of control frames received
outCtrls	Number of control frames transmitted
inDds	Number of frames received on data direct circuits
outDds	Number of frames transmitted on data direct circuits
inBus	Number of frames transmitted to the Broadcast and Unknown Server
outBus	Number of frames received from the Broadcast and Unknown Server
outArpReqs	Number of LE_ARP requests transmitted
outArpRsps	Number of LE_ARP responses transmitted
outBusUcdscrds	Number of unknown unicast frames thrown away by the Broadcast and Unknown Server
cfgElanName	Requested emulated LAN name
oprElanName	Current emulated LAN name
busAddress	ATM address of Broadcast and Unknown Server
lecAddress	ATM address of LAN client
cfgLesAddr	ATM address of requested LAN Emulation Server
oprLesAddr	ATM address of current LAN Emulation Server

Top-Level Menu**Modifying Information About LAN Emulation Clients**

You may configure LAN Emulation Client parameters. To modify these parameters, follow the steps in this section:

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

```
atm lane lec modify
```

You are prompted for the parameter you want to modify.

- 2 Enter the parameter.
You are prompted for the number(s) of the LEC(s) you want to modify.
- 3 Select the number(s) of the LEC(s) or **a11**.
You are prompted for the new value of the option you selected.
- 4 Enter the new value.



If you modify the ELAN name in the LEC configuration, you must disable, and then enable the LEC for the change to take effect.


Defining LAN Emulation Clients

When you define a LAN Emulation Client, you give information necessary for a client to be included in the emulated LAN. A BUS and a LES must already be defined as part of the emulated LAN before you can define a client.

When the client attempts to join the emulated LAN, some of the configured information is carried along with the join request sent to the LES. The LES can alter this information. The client then has the option of accepting or rejecting any changes made by the LES. If the changes are accepted or if there are no changes, the client successfully joins the emulated LAN. Otherwise, the join fails.

To define a client:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:
atm lane lec define
- 2 Enter the number of the ATM port to which the client is attached.
You can attach only one client per port to the same emulated LAN.
- 3 Enter the LAN emulation Configuration Server (LECS) access type (**lecs, manual**) [**LECS**].



If you choose LECs as your means of joining an ELAN, and another LANplex is providing LES and BUS services, you must update the ATM switch with the ELAN name and the corresponding LES address.

- 4 Enter the name of the emulated LAN to which you are adding the client, or press [Return] to choose the default ELAN (shown in brackets). The ELAN name may contain a maximum of 32 characters.
- 5 Enter the maximum frame size (MTU) (**Unspecified**, **1516**, or **4544**).
All LECs within the same emulated LAN must have the same MTU size. Enter **1516** for Ethernet ports, for FDDI passing Ethernet traffic, or for FDDI using IP fragmentation. Enter **4544** for FDDI-to-FDDI traffic. Enter **Unspecified** to allow the LES to select an appropriate MTU size.
- 6 Enter the state for the next reboot.
- 7 Reboot the system to activate the LEC ports.

Example:

```
Select menu option (atm/lane/lec): define
Select ATM port [1]:
Enter LEC access type (lecs, manual) [lecs]:
Enter Elan Name []: Elan_1
Enter Elan MTU Size (Unspecified,1516,4544) [Unspecified]:
Enter Enable State for next reboot (disable, enable)
[disable]: enable
```

You must REBOOT to be able to enable a newly created LEC

Preventing ATM Network Loops

If you have defined LEC ports that share more than one ELAN on two or more LANplex 2500 systems, you need to follow certain steps when you define LECs to prevent network loops. These steps differ, depending on whether you use Intelligent Switching software or Extended Switching software:

- **Intelligent switching software** — Follow these steps to prevent ATM network loops:
 - Before you define an ATM LEC port, be sure that the Spanning Tree protocol is enabled. Enter the command **bridge stpState enable** to enable this protocol.
 - Enter **enable** at the Enter enable state for next reboot prompt.
 - When you finish defining all LEC ports, reboot the LANplex system to activate the ports.

- **Extended switching software** — Follow these steps to prevent ATM network loops:
 - Enter **disable** at the `Enter enable state for next reboot` prompt.
 - When you finish defining all LEC ports, reboot the LANplex system to activate the ports.
 - After you reboot the system, log in to the Administration Console and create a new VLAN for each LEC port beyond the first one. Associate each LEC port with a different VLAN. Be sure that you associate only one LEC with each VLAN. You can use several default VLANs if you are not sure which protocol to specify.
 - Set the state of each LEC port to **enable** through the LEC modify menu.

Removing a LAN Emulation Client

To remove a client from the emulated LAN:

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



```
atm lane lec remove
```

 You are prompted for the number of the LEC to remove.
- 2 Specify the number of the LEC to remove.
 You are prompted to confirm the deletion.
- 3 Enter **y** to confirm or **n** to cancel. If you enter **y**, the system reboots and deletes the specified LEC.

Example:

```
Select menu option (atm/lane/lec): remove
Select LEC (1-3|all) [1]: 1
Delete lec/s resulting in system reboot?" (n,y) [y]: y
```

Top-Level Menu



Administering UNI Management Entities


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.

Each LANplex UME registers one address for LAN Emulation and one address for Classical IP over ATM with the switch port to which it is attached. You can display information about each UME and configure its attributes.

Displaying UME Information

When you display UME information for an ATM port, you display values for the connection state, the VPI (virtual path identifier), the VCI (virtual channel identifier), and other attributes.

To display UME information:



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

atm ume display

You are prompted to select an ATM port.

- 2 Enter the ATM port for which you want to display the UME information.

The UME information appears on the screen.

Example:

```

reqState      state      reqVpi
connected    connected  0

      vpi      reqVci      vci
      0       16       16

connCount     discCount   rxPdus
0             0             0

dropRxDpus   txPdus
0            0
    
```

Table 9-2 describes the types of information provided about the UNI Management Entity.

Table 9-2 UNI Management Entity Attributes

Attribute	Description
connCount	Number of times port has successfully connected since last reboot
discCoun	Number of times port has disconnected since last reboot
dropRxDpus	Number of protocol data units (frames) received but not processed

continued

Table 9-2 UNI Management Entity Attributes (continued)


Attribute	Description
reqState	Requested connection state for management access and address registration (connected or disconnected)
reqVci	Requested virtual channel identifier. This value takes effect after you disable and then enable the port.
reqVpi	Requested Virtual Path Identifier. This value takes effect after you disable and then enable the port.
state	Current connection state for management access and address registration. Possible values: <ul style="list-style-type: none"> ■ Connected: address registration has been successfully completed. ■ Connecting: the circuit is up, but the port hasn't successfully completed address registration. ■ Disconnected: either the reqState has been set to disconnected or the circuit is down. ■ Disconnecting: address registration has been terminated but the circuit is still up.
vpi	Current virtual path identifier
vci	Current virtual channel identifier
rxPdus	Number of protocol data units (frames) received
txPdus	Number of protocol data units (frames) transmitted.

Listing Network Prefixes and Addresses

You can list the registered network prefixes and addresses of any ATM port.

To list prefixes and addresses:

Top-Level Menu



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

```
atm ume list
```

You are prompted for an ATM port.


- 2 Enter the number of the ATM port that you want to register.

The registered network prefixes and addresses are listed for the port you requested.

Setting the Virtual Channel Identifier

You can set the virtual channel identifier (VCI) to be used by UME for the Interim Local Management Interface (ILMI). To set the VCI:

Top-Level Menu



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

```
atm ume vci
```

You are prompted for the number of an ATM port.

- 2 Enter the number of the ATM port for which you want to set the virtual channel identifier.

- 3 Enter the new value for the VCI.

The new VCI takes effect after the port has been disabled and then enabled.

Administering ATM Ports


You can display summary or detailed reports, create labels, and list virtual channel connection (VCC) information for ATM ports.

Displaying Port Information

You can display a summary of ATM port information or a detailed report. When you display a summary, you receive information about the port, including its label, status, and the most important statistics about general port activity and port errors. The detailed report includes the information in the summary plus additional port statistics, such as number of cells discarded.

To display information about ATM ports:

Top-Level Menu



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

```
atm ports summary
```

or

```
atm ports detail
```

The port information is displayed in the format you specified.

Example summary display:

```

      reqStatus      operStatus      linkStatus
outOfService      outOfService      down

      cfgVPCs        cfgVCCs        bandwidth
           1          1          148608000

      rxCells        txCells
           0          3144

```

Example detailed display:

```

      status          cfgVPCs          cfgVCCs
inService           1                21

bandWidth           rxCells          txCells
148608000          1279104701         563694141

ocdEvents           hecErrors          txCellDiscards
           1                0          4288896944

transType           mediaType          uniType
sonet              multiMode          private

uniVersion          tcAlarm          maxVPCs
           3.0          noTcAlarm          16

      maxVCCs          maxVPIbits          maxVCIBits
           1024          4                10

      lastChange
           11 secs

      portLabel

```

Table 9-3 describes these statistics.


Table 9-3 Fields for ATM Ports

Field	Description
bandWidth	Total bandwidth available for port (in bits per second)
cfgVCCs	Current number of virtual channel connections
cfgVPCs	Current number of virtual path connections
hecErrors	Header Error Checksum count
lastChange	Last time the link state of the port changed since the last reboot
maxVCCs	Maximum number of possible virtual channel connections
maxVCIdbits	Maximum number of bits available to represent a Virtual Channel Identifier
maxVPCs	Maximum number of possible virtual path connections
maxVPIbits	Maximum number of bits available to represent a Virtual Path Identifier
mediaType	Type of physical connection media
ocdEvents	Number of times an out-of-cell delineation was detected
portLabel	Label for the physical port
rxCells	Number of cells received
rxCellsDropped	Number of cells received but thrown away
status	Current state of port (inService or outOfService)
tcAlarm	Transmission convergence alarm
transType	Type of transmission sublayer
txCellDiscards	Number of cells thrown out after attempting to transmit
txCells	Number of cells transmitted
uniType	Type of User-Network Interface (public or private)
uniVersion	UNI specification implemented

Labeling a Port

Port labels serve as a useful reference point and as an accurate way to identify your ports for management. You might want to label your ATM ports so that you can easily identify the device specifically attached to each port (for example, LAN, workstation, or server).

To label an ATM port:



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

```
atm ports label
```

2 Enter the numbers of the port(s) you want to label.

3 Enter the label of each ATM port.

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


Listing Virtual Channel Connection Information

You can list general virtual channel connection information as well as specific transmit and receive information.

Listing General VCC Information

To list general virtual channel connection information, enter the following command from the top level of the Administration Console:

```
atm ports vcc list
```



Example:

vpi/vci	reqStatus	operStatus	lastChange	aalType
0/5	other	localDown	711	5

Table 9-4 describes VCC general information.


Table 9-4 Fields for Virtual Channel Connection Information

Field	Description
aalType	Type of ATM Adaptation Layer (aal5)
lastChange	Last time the circuit changed state
operStatus	Current status for the circuit
reqStatus	Requested status for the circuit
vpi/vci	The VPI and VCI numbers for the circuit listed in the present entry

Listing VCC Transmit Information

To list virtual channel connection transmit information, enter the following command from the top level of the Administration Console:

Top-Level Menu



```
atm ports vcc listxmt
```

Example:

```

vpi/vci    transmit
           trafficDescriptor    parameter 1    parameter 2    parameter 3
0/5        atmNoClpNoScr        CLP 0+1 pcr
                                   350490
  
```


Table 9-5 describes VCC transmit information.


Table 9-5 Fields for Virtual Channel Connection Transmit Information

Field	Description
transmit trafficDescriptor	Transmit characteristics of the circuit
parameter 1 parameter 2 parameter 3	Values associated with the transmit trafficDescriptor
vpi/vci	The VPI and VCI numbers for the circuit listed in the present entry

Listing VCC Receive Information

To list virtual channel connection receive information, enter the following command from the top level of the Administration Console:

```
atm ports vcc listrcv
```



Example:

```

vpi/vci   receive
          trafficDescriptor   parameter 1   parameter 2   parameter 3

0/5       atmNoClpNoScr      CLP 0+1 pcr
                               350490
    
```

Table 9-6 describes VCC receive information.

Table 9-6 Fields for Virtual Channel Connection Receive Information

Field	Description
vpi/vci	The VPI and VCI numbers for the circuit listed in the present entry
receive trafficDescriptor	Receive characteristics of the circuit
parameter 1 parameter 2 parameter 3	Values associated with the receive trafficDescriptor

10

SETTING UP THE SYSTEM FOR ROVING ANALYSIS

This chapter describes how to set up the LANplex® 2500 system for roving analysis. With roving analysis, you can monitor Ethernet port activity either locally or remotely using a network analyzer attached to the system.

About Roving Analysis

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

You can monitor a port to:

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

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



Roving analysis frames over a remote Fast Ethernet connection are truncated if greater than 1495 bytes.

To enable the monitoring of ports on a LANplex, take these general steps, explained further in this chapter:

- 1 Select an Ethernet port to which you want to attach the network analyzer.
- 2 Select the Ethernet port that you want to monitor (either local or remote). If the port is remote, you must configure it from the LANplex system on which the remote port is located. The remote system must be located on the same FDDI ring as the system to which the analyzer is attached.

Figure 10-1 shows the process for establishing local and remote monitoring of ports.

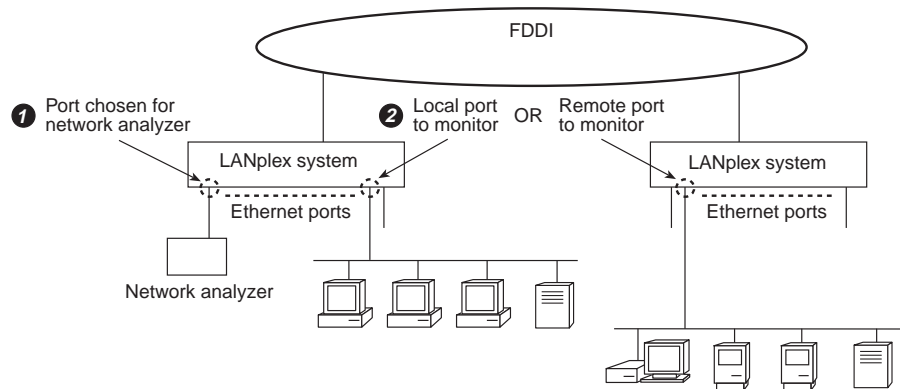


Figure 10-1 Roving Analysis of Local and Remote Ethernet Ports

Configuration rules

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

Displaying the Roving Analysis Configuration

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

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

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

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

analyzer display

Example configuration display:

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

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

Top-Level Menu

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

Adding an Analyzer Port

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

To add analyzer ports:

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

Top-Level Menu

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

The MAC address of the analyzer port is displayed. Record this information for setting up the port you want to monitor. Example:

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

Port selection errors

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

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

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



If the physical port configuration changes in the system (that is, if you remove or rearrange modules), the MAC address of the analyzer port remains fixed. If the module with the analyzer port is moved to another slot, however, then the NVRAM is cleared.

Removing an Analyzer Port

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

To remove analyzer ports:

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

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

Top-Level Menu

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

Starting Port Monitoring

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



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

To start monitoring a new port:

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



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

Example for starting port monitoring:

```
Select port type [Ethernet]:
Select port (1-16): 16
Address: 00-80-3e-0a-3b-02
```

Port selection errors

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

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

MAC address error

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

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



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

Top-Level Menu

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

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

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

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

Stopping Port Monitoring

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

To remove a port configured for monitoring:

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

```
analyzer stop
```


- 2 Press [Return] to select Ethernet as the port type.

- 3 Enter the number of the Ethernet port currently being monitored.

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


Top-Level Menu

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

BRIDGING PARAMETERS

- Chapter 11** Administering the Bridge
- Chapter 12** Administering Bridge Ports
- Chapter 13** Creating and Using Packet Filters
- Chapter 14** Configuring Address and Port Groups to Use in Packet Filters



11

ADMINISTERING THE BRIDGE

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

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

For information about configuring the bridge port, see Chapter 12. For information about creating packet filters for a bridge, see Chapter 13.

Displaying Bridge Information

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

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

```
bridge display
```

Information about the bridge is displayed.

Top-Level Menu

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

Example display of bridge information.

```

    stpState           timeSinceLastTopologyChange
    enabled           1 day 22 hrs 8 mins 31 secs

    topologyChangeCount
    1

    topologyChangeFlag   BridgeIdentifier
    false                8000 08003e0b4800

    designatedRoot      stpGroupAddress      bridgeMaxAge
    7fff 00803e028e02    01-80-c2-00-00-00    20

    maxAge              bridgeHelloTime      helloTime
    20                  2                    2

    bridgeFwdDelay      forwardDelay          holdTime
    15                  15                   1

    rootCost            rootPort              priority
    10                  1                    0x8000

    agingTime           mode                addrTableSize
    300                 802.1d              8191

    addressCount        peakAddrCount        addrThreshold
    214                 214                  8000

    ipFragmentation     ipxTranslation        trFddiMode
    enabled             disabled              n/a

    SRBridgeNumber
    n/a

```

Each item in the bridge parameter display is described in Table 11-1.

Table 11-1 Bridge Attributes

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

(continued)

Table 11-1 Bridge Attributes (continued)

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

Setting the Bridging Mode

You can run a bridge in either Transparent bridging mode or in Express Switching mode. The advantages and disadvantages of using Express Switching are defined in Chapter 6: *Express Switching* in the *LANplex 2500 Operation Guide*.



You cannot configure an ATM LEC as a backbone port for Express Switching mode.

Default value

By default, the bridge is set to run in Transparent bridging mode.



Prior to setting a bridge for Express Switching mode, you must flush any statically configured addresses. See Chapter 12: Administering Bridge Ports for more information about flushing addresses.

To set the bridging mode:

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

bridge mode

- 2 To turn on Express Switching mode, enter:

express

To turn on transparent bridging mode, enter:

transparent

If you are enabling Express Switching, you are prompted for information about a backbone port through which to forward the packets received on Ethernet ports.

- 3 Enter the backbone port type: **Ethernet** or **FDDI** (You cannot enter ATM LEC as the backbone port type.)
- 4 Enter the port number **1** or **2** through which Express packets will be forwarded.

Top-Level Menu

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

Enabling and Disabling IP Fragmentation

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

Default value The default value is *enabled*.

To enable or disable IP fragmentation for a bridge:

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

bridge ipfragmentation

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

enabled

To disable IP fragmentation on a bridge, enter:

disabled

Top-Level Menu

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

Enabling and Disabling IPX Snap Translation

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

Default value The default value is *enabled*.

To enable or disable IPX snap translation for a bridge:

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

bridge ipxSnapTranslation

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

enabled

To disable IPX snap translation on a bridge, enter:

disabled

Top-Level Menu

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

Setting the Address Threshold

Address threshold values

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

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

To set the address threshold:

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

Top-Level Menu

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

Setting the Aging Time

Aging time values

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

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

To set the bridge aging time:

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

Top-Level Menu

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

Administering STP Bridge Parameters

You can enable or disable Spanning Tree Protocol in the system and set the following STP bridge parameters: priority, maximum age, hello time, forward delay, and group address. For more information about how the Spanning Tree parameters interact at the bridge level to create a loopless network, see Chapter 5: *Transparent Bridging* in the *LANplex® 2500 Operation Guide*.

Enabling and Disabling STP on a Bridge

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

The default value is *disabled*.

To enable or disable Spanning Tree Protocol:

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

Top-Level Menu

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

Setting the Bridge Priority

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

Bridge priority values

The bridge priority value is appended as the most significant portion of a bridge identifier (for example: 8000 00803e003dc0). It is a 2-octet value.

Top-Level Menu

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

To configure the STP bridge priority:

- 1 From the top level of the Administration Console, enter:
bridge stpPriority
- 2 Enter the priority value at the prompt.

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

Setting the Bridge Maximum Age

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

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

*Maximum Age
recommended value*

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

To configure the bridge maximum age:

- 1 From the top level of the Administration Console, enter:
bridge stpMaxAge
- 2 Enter the STP bridge max age value.

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

Top-Level Menu

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

Setting the Bridge Hello Time

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

Hello time recommended value

The recommended time is 2 seconds.

To configure the bridge hello time:

Top-Level Menu

```

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

```

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

```
bridge stpHelloTime
```

- 2 Enter the bridge hello time value.

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

Setting the Bridge Forward Delay

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

Setting the value too low could result in temporary loops as the Spanning Tree algorithm reconfigures the topology. However, setting the value too high can lead to a longer wait as the Spanning Tree Protocol reconfigures.

Forward delay recommended value

The recommended value is 15 seconds.

Top-Level Menu

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

Setting the STP Group Address

To configure the forward delay value:

- 1 From the top level of the Administration Console, enter:
bridge stpForwardDelay
- 2 Enter the forward delay value.

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

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

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

To set the STP group address:

- 1 From the top level of the Administration Console, enter:
bridge stpGroupAddress
- 2 Enter the group address.

For IBM Spanning Tree Protocol, the group address must be C0:00:00:00:01:00

Top-Level Menu

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



12

ADMINISTERING BRIDGE PORTS

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

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

Displaying Bridge Port Information

Bridge port information includes the STP configurations for the bridge port. You can display this information in either summary or detail format.

To display bridge information:

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

```
bridge port summary
```

OR


```
bridge port detail
```

You are prompted for the port type.

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

You are prompted for port number(s).

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



The following example shows a bridge port summary display.

```

rxFrames      rxDiscards      txFrames
1680326      1095            654715

                PortID              Stp
                0x8004              Enabled

        State      fwdTransitions
Forwarding          0

```

The following example shows a bridge port detail display.

```

rxFrames      rxBockedDiscs      rxSameSEgDiscs
rxErrorDiscs      rxMcastLimit      rxMcastExcDiscs
rxMcastExceeds      rxSecurityDiscs      rxOtherDiscs
rxAllFilters      rxMcastFilters      rxForwardUcasts
rxFloodUcasts      rxForwardMcasts      txBlockedDiscs
txMtuExcDiscs      txAllFilters      txMcastFilters

txFrames      portID      stp
              0x8001      Enabled

        State      fwdTransition      priority
forwarding          0      0x80

        pathCost      designatedCost      designatedPort
          10              0              0x0

SRRingNumber      SRHopLimit      designatedRoot

designatedBridge

```


Table 12-1 describes the type of information provided for the bridge port.

Table 12-1 Bridge Port Attributes

Parameter	Description
designatedBridge	Identification of the designated bridge of the LAN to which the port is attached
designatedCost	Cost through this port to get to the root bridge. The designated cost of the root port is the same as the cost received in incoming BPDUs from the designated bridge for that LAN.
designatedPort	Identification of the designated port on the designated bridge
designatedRoot	Identification of the bridge designated as root
fwdTransitions	Number of times the port has entered forwarding state. This value is useful for checking the stability of a bridged topology. The more transitions in and out of the forwarding state, the more unstable is the topology.
pathCost	Cost to be added to the total path cost when this port is the root port
port	Ethernet, ATM-Lec, or FDDI (maximum count: 1, 2 = FDDI, 3–18 = Ethernet, and 1-14 for ATM-Lec)
portId	Identification of the port, which includes the port priority and the port number (for example: 8002)
priority	First factor to determine if a port is to be the designated port when more than one bridge port is attached to the same LAN. If all ports in a bridge have the same priority, then the port number is used as the determining factor.
rxAllFilters	Number of frames discarded because of a user-defined packet filter on the receive all path of this bridge port
rxBlockedDiscs	Number of frames discarded by this port because the receiving bridge port was not in the forwarding state
rxDiscards	Total number of received frames discarded (summary report only)
rxErrorDiscs	Number of frames discarded by this port because of internal bridge system errors (such as hardware and software address table discrepancies)
rxFloodUcasts	Number of unicast frames received on this port that were flooded to one or more ports
rxForwardMcasts	Number of multicast frames received on this bridge port and were forwarded to another bridge port

continued

Table 12-1 Bridge Port Attributes (continued)

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

(continued)

Table 12-1 Bridge Port Attributes (continued)

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

Frame processing and bridge port statistics

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

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

Figure 12-1 shows the order in which discard decisions are made.

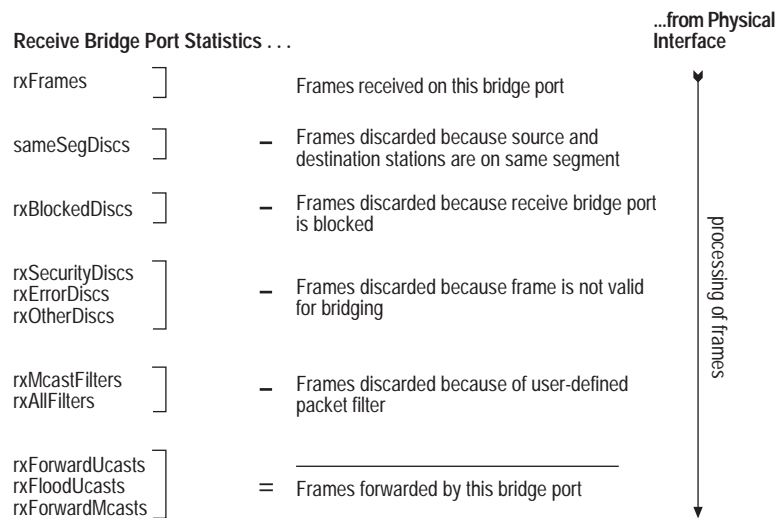


Figure 12-1 How Frame Processing Affects Receive Bridge Port Statistics

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

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

Figure 12-2 shows the order in which the discard decisions are made.

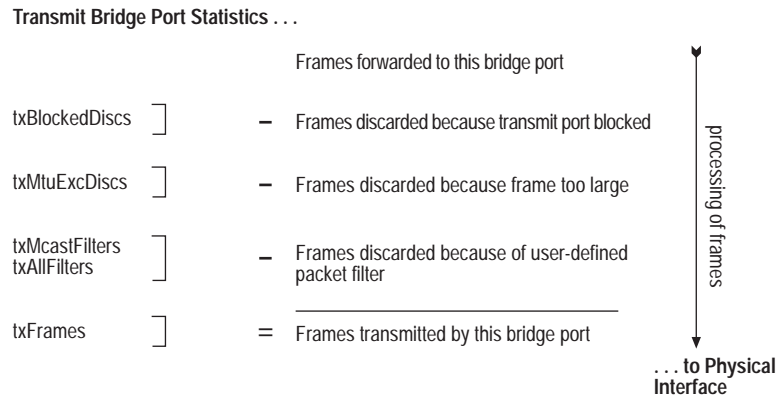


Figure 12-2 How Frame Processing Affects Transmit Bridge Port Statistics

Setting the Multicast Limit

You can assign a multicast packet firewall threshold to a bridge port on the LANplex 2500 system to limit the forwarding rate of multicast traffic originating on the Ethernet segment connected to the port. For more information about the multicast packet firewall, see Chapter 8: *LANplex Bridging Extensions* in the *LANplex 2500 Operation Guide*.

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

To set the multicast limit:

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

```
bridge port multicastLimit
```

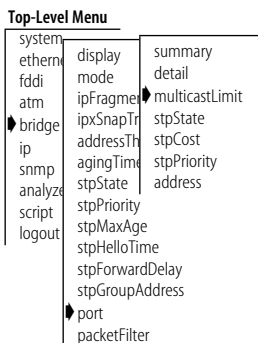
You are prompted for the port type.

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

You are prompted for port number(s).

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

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



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

Example:

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

Administering STP Bridge Port Parameters

You can enable or disable the Spanning Tree Protocol for one or more ports on the system. This only affects the operation of the port if the Spanning Tree Protocol is enabled. You can also set the following STP port parameters: path cost and priority. For more information about how Spanning Tree parameters interact at the bridge-port level, see Chapter 5: *Transparent Bridging* in the *LANplex 2500 Operation Guide*.

Enabling and Disabling STP on a Port

You can enable and disable the Spanning Tree Protocol for any port in the system. When STP is disabled for a port but enabled for the entire bridge, a port does not forward frames or participate in the Spanning Tree algorithm. (See page 11-8 for instructions on enabling STP for the entire bridge.) When STP is disabled for a port as well as for the entire bridge, the port will continue to forward frames.

Default value

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

To enable or disable STP on a port:

Top-Level Menu

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

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

```
bridge port stpState
```

You are prompted for the port type.

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

You are prompted for the port number(s).

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

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

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

Example showing values being set for more than one port:

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

Setting the Port Path Cost

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

Path cost value

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

To configure the path cost:

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

```
bridge port stpCost
```

You are prompted for the port type.

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

You are prompted for the port number(s).

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


You are prompted for the path cost for each port you specified.

- 4 Enter the path cost for the port(s).

Example showing values being set for more than one port:

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

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



Setting the Port Priority

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

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

To configure the port priority:

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

```
bridge port stpPriority
```

You are prompted for the port type.

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

You are prompted for the port number(s).

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

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

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

Example showing values being set for more than one port:

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

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

Top-Level Menu

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

Administering Port Addresses

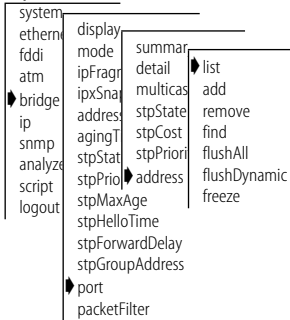
You can administer the MAC addresses of stations connected to Ethernet, FDDI, and ATM ports on the LANplex system.

Listing Addresses

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

To list currently defined MAC addresses:

Top-Level Menu



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

```
bridge port address list
```

You are prompted for the port type.

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

You are prompted for the port number(s).

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

Example address.

Addresses for Ethernet port 1:

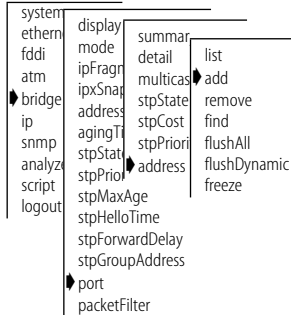
Ethernet address	Type	Age (secs.)
08-00-20-1d-67-e2	Dynamic	219
00-80-3e-02-68-00	Dynamic	219
00-20-af-29-7b-74	Dynamic	219
08-00-02-05-91-c1	Dynamic	219
00-80-3e-02-6d-00	Dynamic	219
00-80-3e-08-5f-00	Dynamic	219
00-80-3e-00-3d-00	Dynamic	219

Adding New Addresses

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

To add a MAC address:

Top-Level Menu



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

```
bridge port address add
```

You are prompted for the port type.

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

You are prompted for the port number.

- 3 Enter the number of the port.

You are prompted for one or more addresses to add.

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

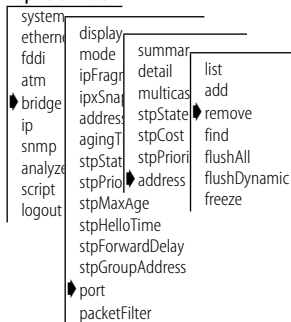
- 5 Enter **q** to return to the previous menu when you finish entering addresses.

Removing Addresses

You can remove individual MAC addresses from selected ports.

To remove an address:

Top-Level Menu



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

```
bridge port address remove
```

You are prompted for the port type.

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

You are prompted for the port number.

- 3 Enter the number of the port.

You are prompted for one or more addresses to remove.

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

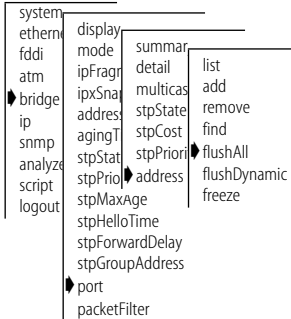
- 5 Enter **q** to return to the previous menu, once you have entered all of the addresses to be removed.

Flushing All Addresses

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

To flush *all* addresses:

Top-Level Menu



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

bridge port address flushAll

You are prompted for the port type.

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

You are prompted for the port number(s).

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

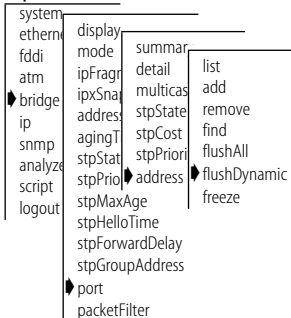
All addresses are flushed from the ports you specified.

Flushing Dynamic Addresses

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

To flush dynamic addresses:

Top-Level Menu



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

bridge port address flushDynamic

You are prompted for the port type.

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

You are prompted for the port number(s).

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

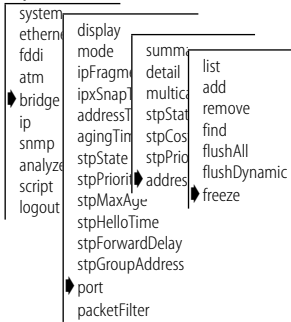
The addresses are flushed from the address table.

Freezing Dynamic Addresses

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

To freeze all dynamic addresses:

Top-Level Menu



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

bridge port address freeze

You are prompted for the port type.

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

You are prompted for the port number(s).

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

The dynamic addresses become static.

13

CREATING AND USING PACKET FILTERS

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

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

About Packet Filtering

Independently configurable packet filtering is provided for the various packet processing paths on each Ethernet and FDDI port of a LANplex system. The packet processing paths are defined in Table 13-1.

Table 13-1 Packet Processing Paths

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

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



For additional detailed explanations of packet filter concepts, see Chapter 7: User-defined Packet Filtering in the LANplex 2500 Operation Guide.

Listing Packet Filters

Top-Level Menu

```

system
ethernet
fddi
atm
bridge
ip
snmp
analyze
script
logout
  display
  mode
  ipFragme
  ipxSnapT
  addressTf
  agingTim
  stpState
  stpPriority
  stpMaxAg
  stpHelloT
  stpForwardDelay
  stpGroupAddress
  port
  packetFilter
  list
  display
  create
  delete
  edit
  load
  assign
  unassign
  addressGroup
  portGroup

```

When you list the packet filters for the system, the filter identification, filter name (if any), and filter assignments are displayed.

To list the currently defined packet filters, enter the following from the top level of the Administration Console:

```
bridge packetFilter list
```

The listing of packet filters is displayed. Example of the output:

```

Ethernet Packet Filters
  Packet Filter 1 - Receive OUI 08-00-1E
    Port 4, Transmit Multicast
    Port 3, Transmit Multicast
    Port 3, Receive Multicast
    Port 5, Receive Multicast
  Packet Filter 2 - Type > 900 or Multicast
    Port 6, Receive All
    Port 8, Transmit All
    Port 8, Receive All
  Packet Filter 3 - Forward IP packets only
  No port assignments

```

In this example, the system has two packet filters. The first packet filter has a filter id of 1 and a user-defined name of “Receive OUI 08-00-1E.” This filter is loaded onto ports 4, 3, and 5. On port 3, the filter is assigned to both the *transmit multicast* and the *receive multicast* paths.

The second filter (filter id 2, user name “Type > 900 or Multicast”) is assigned to ports 6 and 8. The filter is assigned to both the *receive all* and the *transmit all* paths of port 8.

Displaying Packet Filters

When displaying the contents of a single packet filter, you select the packet filter using the filter id number, which you can obtain by listing the packet filters as described in the previous section. The packet filter instructions are displayed; however, any comments in the original packet filter definition file are not displayed because they are not saved with the packet filter.

To display the contents of a packet filter:

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

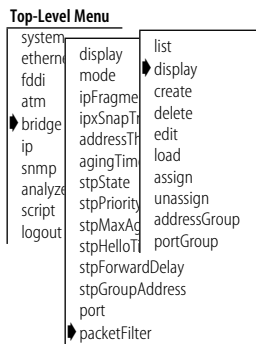
```
bridge packetFilter display
```

You are prompted for the number of the packet filter you want to display.

- 2 Enter the packet filter id number.

The packet filter id and name are displayed, followed by a listing of the packet filter instructions. Example:

```
Select packet filter to be displayed [1-n]: 2
Packet filter 2 - Type > 900 or Multicast
  name "Type > 900 or Multicast"
  pushLiteral.w                0x900
  pushField.w                  12
  gt
  reject
  pushField.b                   0
  pushLiteral.b                0x01
  and
  not
```



Creating Packet Filters

You create custom packet filters by writing a *packet filter definition*. Each packet-processing path on a port may have a unique packet filter definition or may share a definition with other ports. Packet filter definitions are written in the *packet filter language*. This language allows you to construct complex logical expressions.

After writing a packet filter definition, you load it into a LANplex system and the corresponding port assignments are preserved in the nonvolatile memory (NVRAM) of the system. This process ensures that the packet filter configuration for each system is saved across system reboots and power failures.

Concepts for Writing a Filter

Before writing a packet filter, you should understand these basic concepts:

- How the packet filter language works
- The basic elements of a packet filter
- How to implement sequential tests in a packet filter
- The pre-processed and run-time storage requirements

How the Packet Filter Language Works

You define packet filters using a simple, *stack-oriented* language. Stack-oriented means that the language uses a LIFO (last in, first out) queue when the packet filter is running. The program places values (called operands) on the stack and tests them with various logical expressions (called operators), such as *and*, *or*, *equal*, and *not equal* (see Table 13-3 and Table 13-4). These expressions typically test the values of various fields in the received packet, which include MAC addresses, type fields, IP addresses, and Service Access Points (SAPs).

A program in the packet filter language consists of a series of one or more instructions that results in the top of the stack containing a byte value after execution of the last instruction in the program. This byte value determines whether to forward or discard the packet.

In this stack-oriented language, instructions:

- *push* operands onto the stack
- *pop* the operands from the stack for comparison purposes
- *push* the results back onto the stack

Therefore, with the exception of the push instructions, instructions (such as logical operators) locate their operands implicitly and do not require additional operand specifiers in the instruction stream.

Opcodes are the variables used to identify the type of operands and operators you are specifying in the packet filter instructions.

Table 13-2 describes the instructions and stacks of a packet filter.

Table 13-2 Packet Filter Instructions and Stacks — Descriptions and Guidelines

Element	Descriptions and Guidelines
Instructions	<p>Each instruction in a packet filter definition must be on a separate line in the packet filter definition file.</p>
<i>Instruction format</i>	<p>An instruction consists of an <i>opcode</i> followed by explicit <i>operands</i> and a <i>comment</i>. Although comments are optional, it is recommended that you use them throughout the packet filter to make it easier for yourself and others to administer the filters. The opcode includes an explicit operand size specification.</p> <p>The general syntax of an instruction is:</p> <pre data-bbox="469 626 1172 649"><opcode>[.<size>] [<operand>...] [# <comment>]</pre> <p>Example:</p> <pre data-bbox="469 713 1300 736">pushliteral.1 0xffffffff00 #load the type field mask</pre> <p>Use any combination of uppercase and lowercase letters for the opcode and size.</p> <p>The contents of a line following the first # outside a quoted string are ignored, so use the # to begin your comments.</p>
<i>Operand sizes</i>	<p>The following operand sizes are supported:</p> <ul style="list-style-type: none"> ■ 1 byte = .b ■ 2 bytes = .w ■ 4 bytes = .l ■ 6 bytes = .a (Included primarily for use with 48-bit, IEEE, globally assigned MAC addresses)
<i>Maximum length</i>	<p>The maximum length for a filter definition is 4096 bytes.</p>
Stack	<p>The packet filter language uses a <i>stack</i> to store the operands that will be used by an instruction and the results of the instruction.</p> <p>Operands are “popped” from the stack as required by the instructions. An instruction using two or more operands takes the first operand from the top of the stack, with subsequent operands taken in order from succeeding levels of the stack.</p> <p>The stack is a maximum of 64 bytes long, with space within the stack allocated in multiples of 4 bytes. Thus you can have a maximum of 16 operands on the stack.</p> <p>The address size operand .a consumes 8 bytes on the stack, decreasing the maximum number of operands on the stack for a 48-bit address.</p>

Basic Elements of a Packet Filter

Before creating a packet filter, you must decide which part of the packet you want to filter. You can filter Ethernet packets by the destination address, source address, type/length, or some part of the data. You can filter FDDI packets by the destination address, source address, or some part of the data. A packet filter operates on these fields to make filtering decisions. Ethernet and FDDI packet fields are shown in Figure 13-1.

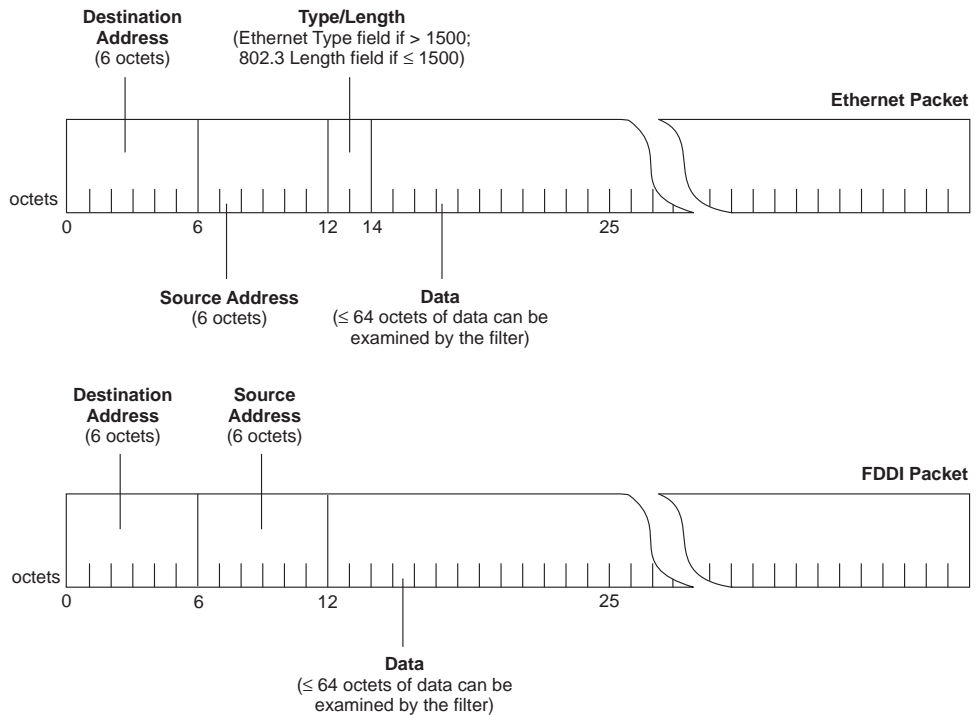


Figure 13-1 Ethernet and FDDI Packet Fields

The Ethernet and FDDI packet fields in Figure 13-1 are used as *operands* in the packet filter. The two simplest operands are described in Table 13-3.

Table 13-3 Two Packet Filter Operands

Operand	Description	Opcode
packet field	A field in the packet that can reside at any offset. The size of the field can be 1, 2, 4, or 6 bytes. Typically, you only specify a 6-byte field when you want the filter to examine a 48-bit address.	pushField
constant	A literal value to which you are comparing a packet field. As with a field, a constant can be 1, 2, 4, or 6 bytes long.	pushLiteral

The *operators* that you specify in the packet filter allow the filter to make a logical decision about whether the packet should be forwarded or discarded. These operators are described in Table 13-4.

Table 13-4 Packet Filter Operators

Operator	Result	Opcode
equal	true if operand 1 = operand 2	eq
not equal	true if operand 1 \neq operand 2	ne
less than	true if operand 1 < operand 2	lt
less than or equal	true if operand 1 \leq operand 2	le
greater than	true if operand 1 > operand 2	gt
greater than or equal	true if operand 1 \geq operand 2	ge
and	operand 1 bit-wise AND operand 2	and
or	operand 1 bit-wise OR operand 2	or
exclusive or	operand 1 bit-wise XOR operand 2	xor
not	true if operand 1 = false	not
shift left	operand 1 SHIFT LEFT operand 2	shiftl
shift right	operand 1 SHIFT RIGHT operand 2	shiftr



The operators **and**, **or**, and **exclusive or** are bit-wise operators. Each bit of the operands is logically compared to produce the resulting bit.

Implementing Sequential Tests in a Packet Filter

Filter language expressions are normally evaluated to completion — a packet is accepted if the value remaining on the top of the stack is nonzero. Frequently, however, a single test is insufficient to filter packets effectively. When more tests are warranted, you want to accept a packet that either:

- Satisfies at least one criterion specified in two or more tests (that is, ORs the results of the tests)
- or
- Satisfies all criteria specified in two or more tests (that is, ANDs the results of the tests)

The *accept* and *reject* instructions are used to implement sequential tests, as shown in Figure 13-2. When using *accept* or *reject*, construct the packet filter so that the tests more likely to be satisfied are performed *before* tests that are less likely to be satisfied.

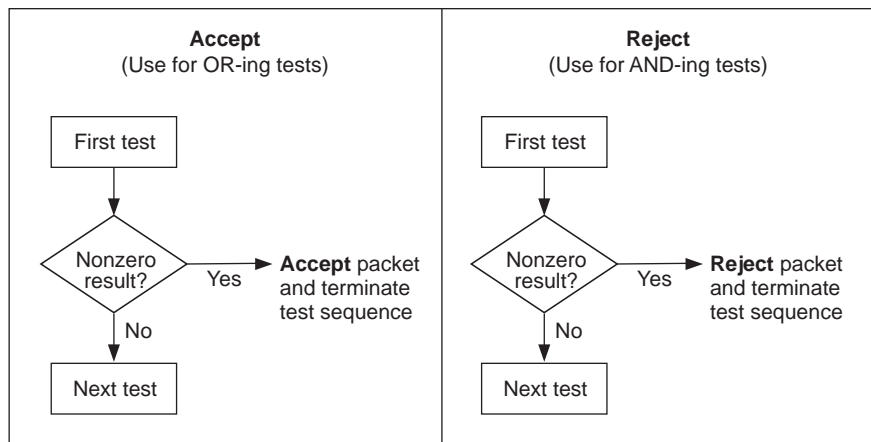


Figure 13-2 Accept and Reject Instructions

The following example shows the use of both accept and reject in a packet filter. This packet filter was created for a network running both Phase I and Phase II AppleTalk™ software. The goal of the filter is to eliminate the AppleTalk traffic.

```

Name      "Filter AppleTalk datagrams"
pushField.w      12      # Get the type field.
pushTop          # Make a copy.
pushLiteral      0x809b  # EtherTalk Phase I type.
eq               # Test if the packet type is
                 # equal to the AppleTalk type,
reject           # reject the packet and end.
                 # Otherwise,
pushLiteral.w    0x5dc   # Largest 802.3 packet size
lt               # If this value is less than the
                 # value in the packet's
                 # type/length field, then this
                 # is an Ethernet frame, so
accept           # accept the packet if it is not
                 # 802.3, otherwise...
pushField.a      16      # get the SNAP OUI and Ethertype
pushLiteral.a    0x03080007809b # value to compare.
ne               # If not equal, then forward the
                 # packet, otherwise drop it.

```

Preprocessed and Run-time Storage

A packet filter program is stored in a preprocessed format to minimize the space required by the packet filter definition. When assigned to a port, the packet filter is converted from the stored format to a run-time format to optimize the performance of the filter. Each LANplex system is limited to a maximum of 16 packet filter programs.

Storage for preprocessed packet filters Each system provides a maximum of 2048 bytes of nonvolatile storage for *preprocessed* packet filter programs. In the preprocessed stored format:

- A single packet filter program is limited to 254 bytes.
- Each instruction in the packet filter program requires 1 byte for the opcode and size, plus additional bytes for any explicit operands.
- System overhead is 22 bytes, plus a per-packet-filter overhead of 13 bytes. For example, assume a packet filter program requires 200 bytes for storing the instructions in the program. If this packet filter is the only one loaded, the nonvolatile memory required is 22 bytes (for system overhead) plus 13 bytes (for packet filter overhead) plus 200 bytes (for the program itself) — a total of 235 bytes.

*Run-time storage of
packet filters*

For *run-time* storage of packet filter programs, each LANplex system provides a maximum of 8192 bytes. There is no explicit system or per-packet-filter overhead; however, performance considerations can result in unused areas of the run-time storage.

The run-time format is approximately eight times the size of the stored format. Thus a 200-byte packet filter program in stored format expands to approximately 1600 bytes in the run-time format. A single packet filter program cannot exceed 2048 bytes in the run-time format.

Procedure for Writing a Filter

This section shows the process of writing a packet filter. Detailed examples are provided in the section “Examples of Creating Filters” on page 13-11.

You write the instructions for the packet filter using the following syntax:

```
<opcode>[.<size>]  [<operand>...] [# <comment>]
```

The opcode descriptions are in Appendix A: *Packet Filter Opcodes, Examples, and Syntax Errors*. The description of the supported operand sizes can be found in Table 13-2. The operand value is determined by what you are testing (for example, an address or a length).



Implicit operands for an instruction must be of the size expected by the instruction. Any mismatch in implicit operand size results in an error “operand size mismatch” when you load the program into the system.

When writing a packet filter, be sure that you use comments (preceded by #) to describe each step in the filter. This habit helps you to revise filters and enables others to understand and use the filters you create.

To write a packet filter:

- 1 Assign a unique, descriptive name to the filter using the `NAME` opcode.
- 2 Specify what to test. For example, use the `PUSHFIELD` opcode to select a field in the packet.
- 3 Specify what to compare to the value in step 2. For example, use the `PUSHLITERAL` opcode to select a constant value.

- 4 Apply a logic operation to the values in steps 2 and 3. The operator you use depends on what comparison you want to make.

Variations on these four basic steps of writing packet filters include:

- Use `pushTop` for each additional comparison you intend to make with the `pushField` value. This opcode makes a duplicate of the `pushField` value and places it on top of the original `pushField` on the stack. The `pushTop` instruction makes a copy of the field more efficiently than if you use a second `pushField` instruction.
- Use `accept` or `reject` with `and` and `or` operators when you have sequential tests and you would like the filter to accept or reject a packet before the entire expression has been evaluated. Using `accept` and `reject` can significantly improve the performance of certain types of filters. See “Implementing Sequential Tests in a Packet Filter” on page 13-8 for more information.
- Use `pushSAGM`, `pushDAGM`, `pushSPGM`, or `pushDPGM` for filtering by address or port groups. See Chapter 14 for more information.

Examples of Creating Filters

The following solution shows a complex packet filter built from three simple packet filters. Each of the shorter, simpler packet filters can be used on its own to accomplish its own task. Combined, these filters create a solution for a larger filtering problem.

Filtering Problem

Your network contains market data feed servers that receive time-critical financial data needed for trading floor applications. At the center of the trading floor networks is a LANplex system that is being used to switch Ethernet traffic and to concentrate the market data feed servers onto the FDDI departmental backbone.

The difficulty is that the market data feed servers transmit data to users with broadcast packets that are forwarded to all stations on all segments attached to the LANplex system. Not all of the segments attached to the LANplex system have stations that require these broadcast updates. In order to optimize the performance of these Ethernet segments, you need to filter the broadcasts.

Packet Filter Solution

The solution described here is to create a highly sophisticated packet filter that prevents only the broadcast packets from the market data servers from being forwarded onto the segments that are not part of an active trading floor.

Before writing the packet filter, it is important to understand the functions that the filter must provide. The broadcast packets that are transmitted by the servers are based on either TCP/IP or XNS protocol. In both cases, the broadcast packets have socket values that are greater than 0x076c and less than 0x0898. The socket value is located 24 bytes into the packet in IP datagrams and 30 bytes into the packet in XNS datagrams.

You can use this information to create pseudocode that simplifies the process of writing the actual filter. It helps to first write the pseudocode in outline form, as shown here:

- 1** Determine if the packet has a broadcast address. (Use the packet filter path assignment.)
- 2** Determine if the packet is an XNS datagram.
- 3** Check socket values and discard the packet if:
 - a** The socket value is greater than or equal to 0x76c
AND
 - b** The socket value is less than 0x898
- 4** Determine if the packet is an IP datagram.
- 5** Check socket values and discard the packet if:
 - a** The socket value is greater than or equal to 0x76c
AND
 - b** The socket value is less than 0x898
- 6** End the filter.

The pseudocode translates into this complex packet filter:

```

Name      "IP XNS ticker bcast filter"
          # Assign this filter in the multicast path
          # of a port only--this is very important.
          #
          # XNS FILTERING SECTION
          #
pushField.w      12      # Get the type field of the packet and
pushLiteral.w    0x0600  # place it on top of the stack.
eq               # Put the type value for XNS on top of
                 # the stack.
pushLiteral.w    0x76c   # If the two values on the top of the
ge               # stack are equal, then return a non-zero
                 # value.
pushField.w      30      # Put the lowest socket value on top of
                 # the stack.
pushField.w      30      # Put the value of the socket from the
ge               # packet on top of the stack.
                 # Compare if the value of the socket is
pushLiteral.w    0x0898  # greater than or equal to lower bound.
                 # Put the highest socket value on top of
pushField.w      30      # the stack.
lt               # Put the value of the socket from the
                 # packet on top of the stack.
and              # Compare if the value of the socket is
                 # less than the upper bound
                 # "and" together with "ge" and "lt" test
                 # to determine if the socket value is
                 # "within" the range. If it is, place a
                 # "one" on the stack
and              #
                 # Compare if XNS & in range
                 #
                 # IP FILTERING SECTION
                 #
pushField.w      12      # Get the type field of the packet and
pushLiteral.w    0x0800  # place it on top of the stack.
eq               # Put the type value for IP on top of
                 # the stack.
pushLiteral.w    0x76c   # If the two values on the top of the
ge               # stack are equal, then return a non-zero
                 # value.
pushField.w      24      # Put the lowest socket value on top of
                 # the stack (1900).
pushField.w      24      # Put the value of the socket from the
ge               # packet on top of the stack.
                 # Compare if the value of the socket is
pushLiteral.w    0x0898  # greater than or equal to lower bound.
                 # Put the highest socket value on top of
pushField.w      24      # the stack (2200).
lt               # Put the value of the socket from the
                 # packet on top of the stack.
and              # Compare if the value of the socket is
                 # less than the upper bound
                 # "and" together with "ge" and "lt".
                 # Test to determine if the socket value is
                 # "within" the range. If it is in range,
                 # place a "one" will on the stack.
or               # Compare if IP and in range.
                 # Determine if the type field is either
not              # XNS or IP.
                 # Discard if (IP & in range) and (XNS & in
                 # range).

```

The rest of this section concentrates on the parts of the complex filter, showing you how to translate the pseudocode's requirements into filter language. The large filter on page 13-13 is broken down into subsets to show how you can create small filters that perform one or two tasks, and then combine them for more sophisticated filtering. Table 13-5 shows how the purpose of each pseudocode step is accomplished in the small series of packet filters.

Table 13-5 Pseudocode Requirements Mapped to the Packet Filter

Step	Accomplished Through...
1	The path to which you assign the packet filter. For administrative purposes, this path is specified in the first two comment lines in the filter definition. The filter must be assigned to a multicast path to filter packets that have broadcast addresses.
2	Packet Filter One — Forwarding XNS packets
3	Packet Filter Two — Checking for specified socket range
4 & 5	Combining a Subset of Filters — Forwarding IP packets within specified socket range

Packet Filter One. This filter is designed to forward XNS packets. These steps show how to create this filter.

- 1 Name the filter:

```
"Forward only XNS packets"
```

It is important to distinguish the function of each filter when it is loaded onto a LANplex system that has more than one filter stored in memory. Naming is also useful for archiving filters on an ftp server so that the filters can be saved and loaded on one or more LANplex systems.

- 2 Enter executable instruction #1:

```
pushField.w 12 # Get the type field of the packet and  
# place it on top of the stack.
```

- 3 Enter executable instruction #2:

```
pushLiteral.w 0x0600 # Put the type value for XNS on top  
# of the stack.
```

4 Enter executable instruction #3:

```
eq # If the two values on the top of the stack are equal,  
  # then return a non-zero value.
```

Packet Filter Two. This filter is designed to accept packets within the socket range of 0x76c and 0x898. These steps show how to create this filter.

1 Name the filter:

```
"Socket range filter"
```

2 Enter executable instruction #1:

```
pushLiteral.w 0x76c # Put the lowest socket value on top  
  # of the stack.
```

3 Enter executable instruction #2:

```
pushField.w 30 # Put the value of the socket from the  
  # packet on top of the stack.
```

4 Enter executable instruction #3:

```
ge # Compare if the value of the socket is greater than  
  # or equal to the lower bound.
```

5 Enter executable instruction #4:

```
pushLiteral.w 0x0898 # Put the highest socket value on  
  # top of the stack.
```

6 Enter executable instruction #5:

```
pushField.w 30 # Put the value of the socket from the  
  # packet on top of the stack.
```

7 Enter executable instruction #6:

```
lt # Compare if the value of the socket is less than the  
  # upper bound.
```

8 Enter executable instruction #7:

```
and # "and" together with "ge" and "lt" test to determine  
  # if the socket value is "within" the range. If it is,  
  # place a "one" on the stack.
```

Combining a Subset of the Filters. The next filter accepts IP packets with a socket range of 0x76c (1900) and 0x898 (2200). The filter combines packet filters one and two, modifying them for IP. These steps show how to create this filter.

- 1 Name the filter:
"Only IP pkts w/in socket range"
- 2 Perform steps 2 through 4 as described in "Packet Filter One" on page 13-14, except give the pushLiteral instruction (in step 3) a value of 0x0800 for IP.
- 3 Perform steps 2 through 8 as described in "Packet Filter Two" on page 13-15, except the socket value for IP (in step 3) is located 24 bytes into the packet (instead of 30 as for XNS).
- 4 Add an *and* statement to compare the results of step 2 with the results of step 3:
and # Compare if IP and in range.

This combination looks like this:

```
Name      "Only IP pkts w/in socket range"
pushField.w    12      # Get the type field of the packet and
                  # place it on top of the stack.
pushLiteral.w  0x0800  # Put the type value for IP on top of
                  # the stack.
eq           # If the two values on the top of the
                  # stack are equal, then return a non-zero
                  # value.
pushLiteral.w  0x76c   # Put the lowest socket value on top of
                  # the stack (1900).
pushField.w    24      # Put the value of the socket from the
                  # packet on top of the stack.
ge           # Compare if the value of the socket is
                  # greater than or equal to the lower bound
pushLiteral.w  0x0898  # Put the highest socket value on top of
                  # the stack (2200).
pushField.w    24      # Put the value of the socket from the
                  # packet on top of the stack.
lt           # Compare if the value of the socket is
                  # less than the upper bound.
and          # "and" together with "ge" and "lt" test
                  # to determine if the socket value is
                  # "within" the range. If it is in range,
                  # place a "one" will on the stack.
and          # Compare if IP and in range.
```

Combining All the Filters. Together, the four packet filters work to perform the solution to the problem: filtering the broadcast packets from the market data servers. These steps show how to create this filter:

1 Name the filter:

```
"Discard XNS & IP pkts w/in socket range"
```

2 Perform steps 2 through 4 as described in "Packet Filter One" on page 13-14.

3 Perform steps 2 through 8 as described in "Packet Filter Two" on page 13-15.

4 Add an *and* statement to compare the results of step 2 and the results of step 3:

```
and # compare if XNS & in range
```

5 Perform steps 2 through 4 as described in "Combining a Subset of the Filters" on page 13-16.

6 Add an *or* statement:

```
or # determine if the type field is either XNS or IP
```

7 Add a *not* statement to discard any matching packets:

```
not # discard if (IP & in range) & (XNS & in range)
```

The complete packet filter that discards IP and XNS packets that are within the specified range is shown on page 13-13.

Tools for Writing a Filter

You can create a new packet filter using either an ASCII-based text editor (such as *EMACS* or *vi*) or the line editor built into the Administration Console. Using an ASCII-based text editor allows you to create multiple copies of the packet filter definition, which you can then store and copy onto one or more LANplex systems from a networked workstation. This method also allows you to archive copies of filter definitions.

Using the Built-in Line Editor

The Administration Console's built-in text editor provides a minimal set of editing functions that you can use to edit a packet filter definition one line at a time. A single line is limited to no more than 79 characters. The number of lines is limited only by available memory.



The maximum length of a packet filter definition is 4096 bytes.

The built-in editor assumes a terminal capability no higher than a glass tty (that is, it does not assume an addressable screen). You can place any ASCII printable character into the editing buffer at the cursor position. If a character exceeds the maximum line length, the character is discarded and a bell sounds. The built-in editor initially operates in *insert* mode. The commands supported by the editor are summarized in Table 13-6.

To use the built-in line editor to create a packet filter definition:

Top-Level Menu

system		
ethernet	display	list
fdi	mode	display
atm	ipFragme	create
bridge	ipxSnapT	delete
ip	addressT	edit
snmp	agingTim	load
analyze	stpState	assign
script	stpPriority	unassign
logout	stpMaxAc	addressGroup
	stpHelloT	portGroup
	stpForwardDelay	
	stpGroupAddress	
	port	
	packetFilter	

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

```
bridge packetFilter create
```

The packet filter line editor appears.

- 2 Enter the definition for the packet filter. See the commands in Table 13-6.
- 3 Save the packet filter by pressing Ctrl+W.

The software checks the syntax of the filter. If the software detects any errors, it displays them and re-enters the editor is at the line containing the first error. After correcting the errors, attempt to save the packet filter again.

After you have corrected all errors and successfully saved the packet filter, the software converts the file to internal form and stores it on the system.

Table 13-6 Commands for the Administration Console's Built-in Packet Filter Editor

Command	Keys	Description
List buffer	Ctrl+l	Displays each of the lines in the editing buffer and then redisplay the line currently being edited
Next Line	Ctrl+n	Moves cursor to next line; positions cursor at start of line
Previous Line	Ctrl+p	Moves cursor to previous line; positions cursor at start of line
Start of Line	Ctrl+a	Moves cursor within a line to the start of the present line
End of Line	Ctrl+e	Moves cursor within a line to the end of the present line
Left 1 Character	Ctrl+b	Moves cursor <i>left</i> one character within a line
Right 1 Character	Ctrl+f	Moves cursor <i>right</i> one character within a line
Insert Line	Enter	Inserts a new line. The new line becomes the current line, with the cursor positioned at the start. If the cursor is positioned over the first character on a line when you press [Enter], a blank new line is inserted before the current line. Otherwise, the current line is split at the cursor position, with the current line retaining the characters before the cursor, followed by the new line containing the remainder of the characters.
Delete Previous Character	Ctrl+h	Deletes a single character preceding the cursor and shifts the remainder of the line <i>left</i> one position
Delete Current Character	Ctrl+d	Deletes a single character under the cursor and shifts the remainder of the line <i>left</i> one position
Delete Line	Ctrl+k	Deletes the remainder of the line from the current cursor position. If the cursor is positioned over the first character, all of the characters on the line are deleted, but the line is retained. A second Delete Line command removes the line from the edit buffer.
Insert/Overstrike Toggle	Ctrl+o	Toggles between the insert mode and overstrike mode
Write Changes	Ctrl+w	Writes (saves) the current contents of the edit buffer into the packet filter definition. No syntax checking of the definition is performed at this point other than to verify that the length of the source is within the maximum limits. If the source is too long, the message <code>Error: Edit buffer exceeds maximum length</code> is displayed. The contents of the edit buffer are unaffected; however, the packet filter definition contains only those lines that fit entirely within the length limitation.
Exit Editor	ESC	Allows you to leave the editor. You receive a warning if the edit buffer has not been successfully written since the last modification. You can either discard the changes or return to the editor. Note that only those changes made since the last Write Changes command are discarded.

Using an External Text Editor

To use an ASCII-based editor to create a packet filter:

- 1 Create the definition in a text file.
- 2 From a networked workstation, ftp the file to the LANplex system on which you want to load the filter.
- 3 Load the filter as described in “Loading Packet Filters” on page 13-22.

Deleting Packet Filters

Deleting a packet filter removes the filter from the LANplex system.

To delete a packet filter:

- 1 From the top level of the Administration Console, enter:
bridge packetFilter delete
- 2 Enter the id of the filter to delete. To find the id of the filter, list the filters as described in “Listing Packet Filters” on page 13-2.
You are prompted to confirm the deletion.
- 3 Enter **y** (yes) to delete or **n** (no) to return to the previous menu.

Top-Level Menu

```

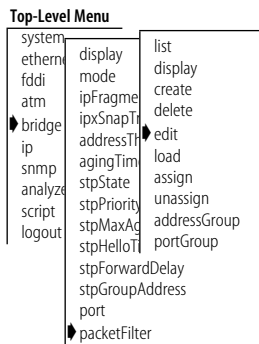
system
ethernet
fdi
atm
bridge
ip
snmp
analyze
script
logout
display
mode
ipFragme
ipxSnapT
addressTr
agingTim
stpState
stpPriority
stpMaxAd
stpHelloT
stpForwardDelay
stpGroupAddress
port
packetFilter
list
display
create
delete
edit
load
assign
unassign
addressGroup
portGroup

```

Editing, Checking, and Saving Packet Filters

You can use the built-in line editor to edit packet filters. Once you save the packet filter, the software checks it for syntax errors. The LANplex system software will not allow you to assign the packet filter to a port until the filter is error-free.

You can also edit a packet filter using an ASCII-based text editor such as *EMACS* or *vi*. You can then use ftp to send the filter text to the LANplex system from a networked workstation.



To edit a packet filter using the built-in line editor:

- 1 From the top level of the Administration Console, enter:
bridge packetFilter edit
- 2 Enter the packet filter id number.
Specifying a filter id loads that filter into the edit buffer.
- 3 Edit the filter. For more information, see the section “Using the Built-in Line Editor” on page 13-18.
- 4 Press [Esc] to exit the line editor.
- 5 At the `Edit` buffer has been changed. Quit anyway? prompt, enter **y** (yes) to end the editing session or **n** (no) to return to editing.
- 6 You have three choices of what to do next:
 - To overwrite the existing filter with the contents of the edit buffer, enter **y** at the `Replace existing filter?` prompt.
 - To store the definition as a new filter, enter **n** at the `Replace existing filter?` prompt and **y** at the `Store as new filter?` prompt. The packet filter is assigned a number.
 - To exit from the editor without saving changes, enter **n** at both prompts.

Correcting errors in a packet filter

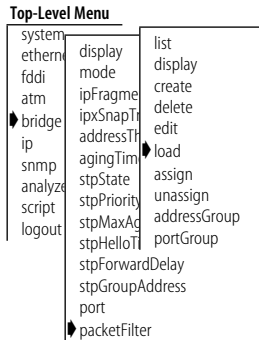
When you save a packet filter edited with the built-in text editor, the system checks the syntax of the filter definition. If any errors are detected, the errors are displayed and the editor is re-entered at the line containing the first error. After correcting the errors, you must exit the editor and attempt to save the packet filter again.

After you have corrected all errors and saved the packet filter, it is converted to internal form and updated on the system.

Loading Packet Filters

When you create packet filters using an external text editor, you must load the filters onto the system from the network host on which you created them. Once loaded, the packet filter definition is converted into the internal format that is used by the packet filter code in the system.

To load a packet filter:



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

```
bridge packetFilter load
```

You are prompted for a host IP address, file path name, user name, and password. To use the value in brackets, press [Return] at any prompt.

- 2 Enter the host IP address.
- 3 Enter the path name.
- 4 Enter your user name.
- 5 Enter your password.

The packet filter is loaded onto the LANplex system.

Any syntax errors in the packet filter definition are reported to you at this time. See Appendix A: *Packet Filter Opcodes, Examples, and Syntax Errors* for a description of these errors. If errors are detected, you are offered the option of editing the filter definition or terminating the load.

The load might fail if the system has insufficient nonvolatile RAM to store the filter. In this case, an error message tells you that the system did not accept the load.

Assigning Packet Filters to Ports

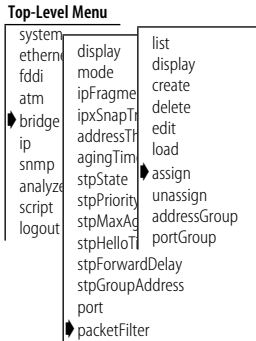
To assign a packet filter to one or more ports, the packet filter must reside on the system. Each path (transmit all, transmit multicast, receive all, and receive multicast) of a port can have only one packet filter assigned to it; however, you can assign a single packet filter to multiple paths and ports.

Packet filter path assignments

Placing a filter on the transmit path confines the packet to the segment it originated from if it does not meet the forwarding criteria. Placing a filter on the receive path prohibits a packet from accessing certain segments unless

it meets the forwarding criteria. A packet that does not meet the forwarding criteria defined in the filter is discarded.

To assign a packet filter:



- 1 From the top level of the Administration Console, enter:
bridge packetFilter assign
- 2 Enter the id number of the packet filter to be assigned. To find the id of the packet filter, follow the instructions in “Listing Packet Filters” on page 13-2.
- 3 Enter the port type (**Ethernet, FDDI, ATM-Lec, all**).
- 4 Enter the port(s) to assign the filter.
- 5 Enter the path(s) you want to place the filter (**txA, txM, rxA, rxM, all**).

In this example, the assignment is to the transmit all (txA) path and the receive all (rxA) path on port 1.

```

Select filter [1-n]: 1
Select port type(s) (Ethernet,FDDI,ATM-Lec|all)
[Ethernet,FDDI,ATM-Lec]: FDDI

Select port(s) (1-16|all) [1-16]: 1
Select path(s) (txA,txM,rxA,rxM|all): txA,rxA

```

The ports are limited to those that have at least one path unassigned, while the paths are limited to those that are unassigned. Because you can specify multiple selections at each level, you can assign a wildcard that attempts to assign the filter to the set indicated by the ports and paths taken in combination.



One or more assignments might fail because of a previous assignment.

Unassigning Packet Filters from Ports

Top-Level Menu

```

system
ethernet
fddi
atm
bridge
ip
snmp
analyze
script
logout
  display
  mode
  ipFragme
  ipxSnapT
  addressT
  agingTim
  stpState
  stpPriority
  stpMaxAg
  stpHelloT
  stpForwardDelay
  stpGroupAddress
  port
  packetFilter
  list
  display
  create
  delete
  edit
  load
  assign
  unassign
  addressGroup
  portGroup

```

To unassign a packet filter from one or more ports, the packet filter must have been previously assigned to at least one port.

To unassign a packet filter:

- 1 From the top level of the Administration Console, enter:
bridge packetFilter unassign
- 2 Enter the id number of the packet filter you want to unassign.
- 3 Enter the port type (**Ethernet, FDDI, ATM-Lec, all**).
- 4 Enter the port number(s) of the packet filter you want to unassign.
- 5 Enter the path(s) of the packet filter you want to unassign.

In this example, the unassignment is from the transmit all (txA) paths on port 1.

```

Select filter [1-n]: 1
Select port type(s) (Ethernet,FDDI,ATM-Lec|all)
[Ethernet,FDDI,ATM-Lec]: FDDI
Select port(s) (1-16|all) [1-16]: 1
Select path(s) (txA,txM,rxA,rxM|all) [txA,rxA]: txA

```

Because you can specify multiple selections at each level, you can assign a wildcard that attempts to unassign the filter from the set indicated by the ports and paths taken in combination.



One or more of the unassignments might fail if the filter is not assigned.

14

CONFIGURING ADDRESS AND PORT GROUPS TO USE IN PACKET FILTERS

This chapter describes how to use address and port groups as filtering criteria in a packet filter, and how to administer address and port groups.

Using Groups in Packet Filters

You can use address groups (a list of MAC addresses) and port groups (a list of Ethernet and FDDI ports) as filtering criteria in a packet filter.



For more information about address and port group concepts, see Chapter 7: User-defined Packet Filtering in the LANplex® 2500 Operation Guide.

A packet filter uses a group to make filtering decisions by accessing the group's source group mask and destination group mask. You reference these group masks using the opcodes SAGM (source address group mask), DAGM (destination address group mask), SPGM (source port group mask), and DPGM (destination port group mask). What follows are some examples of using address and port groups in packet filters.

Address group packet filter example

In this example, the filter only forwards packets among stations that are within the same address group.

```
Name      "Accept Same Source and Destination"
pushSAGM   # Get source address group mask.
pushDAGM   # Get destination address
           # group mask.
and        # Compare if source address and
           # destination address are common
           # members of an address group (result
           # is either zero or non-zero).
pushLiteral.1  0 # Put a zero on the stack.
ne         # If not equal, return a "one" to
           # stack, resulting in packet
           # forwarded.
```

Port group packet filter example

In this example, packets are not forwarded to ports in groups 3 and 8.

```
Name      "Discard Groups 3 and 8"
pushSPGM          # Get source port group mask.
pushLiteral.1    0x0084 # Select bits 3 and 8.
and              # If port group bits 3 & 8 are common
                # with SPGM, then non-zero value is
                # pushed onto stack.
pushLiteral.1    0      # Push zero.
eq              # Only if SPGM is not in port groups
                # corresponding to bits 3 & 8, then
                # packet is forwarded.
```

In the Administration Console you can:

- List the groups
- Display specific information about a group
- Create a new group
- Delete a group
- Copy a group from one module to another (address groups only)
- Add and remove addresses and ports to or from a group

Listing Groups

You can list the address and port groups currently defined for your LANplex® system. The group id, group name (if any), and group mask are displayed.

Top-Level Menu

```
system
ethernet  display
fddi      mode
atm       ipFragm
          create
bridge    ipxSnap delete
          address edit
          agingTim load
          snmp      stpState assign
          analyze  stpPrior unassign
          script   stpMax addressGroup
          logout  stpHello portGroup
                stpForw
                stpGroupAddress
                port
                packetFilter
```

- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter:

```
bridge packetFilter portGroup
```

- 2 To list the currently defined groups, enter:

```
list
```

The system displays the listing of address or port groups, as shown in the next example.

Address group example

In this example, three address groups are defined in the system. The first address group has an id of 1 and the name *Accounting*. This group uses an address group mask of 1 (the bit set in the mask).

```
Address Groups
Address Group 1 - Accounting
    Address group mask - bit 1
Address Group 2 - Development
    Address group mask - bit 6
Address Group 3 - Sales
    Address group mask - bit 3
```

Port group example

In this example of listing port groups, two port groups are defined in the system. The first port group has an id of 1 and the name *Sales*. This group uses a port group mask of 7 (the bit set in the mask).

```
Port Groups
Port Group 1 - Sales
    Port group mask - bit 7
Port Group 2 - Manufacturing
    Port group mask - bit 23
```

Displaying Groups

The display of an address or port group shows the group id, the name of the group, and all the addresses or ports included in that group.

To display address or port groups:

- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter:

```
bridge packetFilter portGroup
```

- 2 To display a port or group, enter:

```
display
```


- 3 Enter the id number of the address or port group you want to display.

The system displays the address or port group you selected.

Address group example

In this example, address group 2 is displayed. The address group id and the name (if any) are displayed, followed by Ethernet addresses that are

Top-Level Menu



members of the group. The name of the address group in this example is *Development*, and the group has five members.

```
Select address group to be displayed [1-n]: 2
```

```
Address Group 2 - Development
05-39-24-56-ab-ee      08-29-34-fd-32-14      08-29-34-dd-ee-01
09-34-56-32-12-e3     00-14-32-54-fd-4e
```

Port group example

In this example, port group 2 is displayed. The port group id and the name (if any) are displayed, followed by the ports that are members of the group. The name of the port group in this example is *Manufacturing*, and the group has three members.


```
Select port group to be displayed [1-n]: 2
```

```
Port Group 2 - Manufacturing
Ethernet port 1  Ethernet port 5  FDDI port 1
```

Creating New Groups

When you create a new address or port group, an unused address or port group must be available. A port group is limited to the number of ports on the system.

Top-Level Menu



- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter the following command:

```
bridge packetFilter portGroup
```

- 2 Enter this command:
create
- 3 For address groups, enter the address group mask.
For port groups, enter the port group mask.
- 4 Enter the address or port group name.
- 5 Enter the addresses or ports to add to the new group.
Enter the addresses in MAC format as:

MAC address format xx-xx-xx-xx-xx-xx

Enter the ports in this syntax:

Port syntax < Ethernet | E | FDDI | F > [port] < port number >

As you enter each address or port, the system attempts to add it to the group. If the address or port you enter is already a member of the group, the system displays a message, as shown next, and the address or port is ignored.

Error: Address already in address group.

OR

Error: Port already in port group.

For an address group, if the system fails to accept the additional address, the address is not added to the group and the system displays an error message as follows:

Error: No room in group for an additional address.

When this message occurs, the specified address is ignored and creation of the address group stops. All addresses entered up to the last address are added to the group, and the new group is loaded on the system.

If you enter an invalid port name, the port is not added to the group, and you receive one of the following error messages:

Error: No port type specified for the port.

Error: No port number specified for the port.

The correct format is < Ethernet | E | FDDI | F > [port]
< port number >

Specified port number is invalid.

Valid FDDI port for this group is 1 or 2.

6 Type **q** after entering the last addresses or ports.

*Address group
example*

In this example, a new address group is created and loaded on the system. The address group mask for the group is 5 and the name of the group is *Marketing*. Two Ethernet addresses are entered and assigned to the group.

Select a bit in the address group mask [3-8, 14-32]: 5

Enter the address group name: **Marketing**

Enter the addresses for the group - type **q** to return to the menu:

Address: **08-32-45-fe-76-d3**

Address: **08-32-45-e3-32-21**

Address: **q**

Address Group 4 - Marketing - has been loaded

Port group example

In this example, a new port group is created and loaded on the system. The bit in the port group mask for the group is 12 and the name of the group is *Education*. One port is entered and assigned to the group.

```
Select a bit in the port group mask [3-8, 14-32]: 12
Enter the port group name: Education
Enter the ports for the group - type q to return to the menu:
Port: Ethernet 2
Port: q
Port Group 6 - Education - has been loaded
```

Deleting Groups

When you delete address or port groups from the system, those groups are no longer available for use in packet filters.



If you want to use a group later but want to delete it now, first save it to an ASCII file.

To delete an address or port group:

Top-Level Menu

```
system
ethernet
fddi
atm
bridge
ip
snmp
analyze
script
logout
  display
  mode
  ipFragm
  ipxSnap
  address
  agingTim
  stpState
  stpPrior
  stpMaxA
  stpHello
  stpForw
  stpGroupAddress
  port
  packetFilter
  list
  display
  create
  delete
  edit
  load
  assign
  unassign
  addressGroup
  portGroup
```

- 1 For *address* groups, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup
```

OR, for *port* groups, enter:

```
bridge packetFilter portGroup
```

- 2 Enter:

```
delete
```

You are prompted for the ID of the address group or port group that you want to delete.

- 3 Enter the ID number of the group you want to delete.

Adding Addresses and Ports to Groups

When adding addresses or ports to an existing group, you can either enter the addresses or ports at the prompts or import them from a file. At least one address group or port group must exist before you can add addresses or ports. (See “Creating New Groups” on page 14-4.) The same address may be in multiple address groups.

Address group size

An address group for the LANplex 2500 system supports a maximum of 8192 addresses in both 802.1d Bridging mode and Express Switching mode. When you load an address group, the addresses that are not currently in the table are added. Therefore, the actual number of entries that you can add to an address group is limited by the address table size.

Port group size

The maximum number of ports a port group can contain is 18, which is the maximum number of ports on a LANplex 2500 system.

To add addresses or ports to an existing group:

Top-Level Menu

system		
ethernet	display	list
fdi	mode	
atm	ipFragm	display
bridge	ipxSnap	create
ip	address	delete
snmp	agingTim	edit
analyze	stpState	load
script	stpPrior	assign
logout	stpMax	unassign
	stpHello	addressGroup
	stpForw	portGroup
	stpGroupAddress	
	port	
	packetFilter	

- 1 To add an *address* group, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup addAddress
```

OR, to add a *port* group, enter:

```
bridge packetFilter portGroup addPort
```

- 2 Enter the number of the group you want to modify.
- 3 Enter the addresses or ports you want to add to the group.

MAC address format

Enter the addresses in MAC format as:

```
xx-xx-xx-xx-xx-xx
```

Enter the ports in this syntax:

Port syntax < Ethernet | E | FDDI | F > [port] < port number >

As you enter each address or port, the system attempts to add it to the group.

If the address or port you enter is already a member of the group, a message is displayed, as shown next, and the address or port is ignored.

```
Error: Address already in address group.
```

OR

```
Error: Port already in port group.
```

For address groups, if the system fails to accept the additional address, the address is not added to the group and an error message is displayed as follows:

```
Error: No room in group for additional address.
```

The point at which the system runs out of room for additional addresses depends on:

- The number of addresses currently in the address table.
- The number of unique addresses configured across all address groups on the system. (Each statically configured address and each unique address assigned to one or more address groups consumes one address storage location.)

For port groups, entering an invalid port specification results in error message, similar to those described on page 14-5.

4 Enter **q** after entering all the addresses or ports.

*Address group
example*

In this example, two additional addresses are added to the *Development* address group.

```
Select address group to be modified [1-4]: 2
```

```
Adding addresses to group 2 - Development
```

```
Enter the addresses to be added - type q to return to the menu:
```

```
Address: 08-21-42-62-98-ab
```

```
Address: 08-37-21-65-78-c4
```

```
Address: q
```

Port group example

This example shows a port successfully added to the *Manufacturing* port group.

```
Select port group to be modified [1-4]: 2
```

```
Adding ports to group 2 - Manufacturing
```

```
Enter the ports to be added - type q to return to the menu:
```

```
Port: Ethernet 3
```

```
Port: q
```

Removing Addresses or Ports from a Group

At least one group must exist before you can remove an address or port.

To remove addresses or ports from an address or port group:

Top-Level Menu

system		
ethernet	display	list
fddi	mode	
atm	ipFragm	display
bridge	ipxSnap	create
ip	address	delete
snmp	agingTim	edit
analyze	stpState	load
script	stpPrior	assign
logout	stpMaxA	unassign
	stpHello	addressGroup
	stpForw	portGroup
	stpGroupAddress	
	port	
	packetFilter	

- 1 To remove an *address* group, enter the following command from the top level of the Administration Console:

```
bridge packetFilter addressGroup removeAddress
```

OR, to remove a *port* group, enter:

```
bridge packetFilter portGroup removePort
```

- 2 Enter the number of the group you want to modify.
- 3 Enter the addresses or ports to remove from the group.

Enter the addresses in MAC format as:

```
xx-xx-xx-xx-xx-xx
```

Enter the ports in this syntax:

```
< Ethernet | E | FDDI | F > [port] < port number >
```

As you enter addresses and ports, the system attempts to remove them from the designated group.

If the address or port is not found in the group, the system displays a warning message, as shown here:

```
Warning: Specified address was not a member of the
address group.
```

OR

```
Warning: Specified port was not a member of the port
group.
```

The specified address or port is ignored, and the system prompts you for the next one you want to remove.

4 Type **q** after entering all the addresses or ports.

*Address group
example*

In this example, the user is removing two Ethernet addresses from the *Marketing* address group.

```
Select address group to be modified [1-4]: 4
Removing addresses from group 4 - Marketing
Enter the addresses to be removed - type q to return to the menu:
Address: 08-37-21-65-78-c4
Address: 08-42-21-84-78-f1
Address: q
```

Port group example

In this example, the user is removing an Ethernet and an FDDI port from the *Education* port group.

```
Select port group to be modified [1-4]: 4
Removing ports from group 4 - Education
Enter the ports to be removed - type q to return to the menu:
Port: FDDI 1
Port: Ethernet 4
Port: q
```


Loading Groups

The Administration console has no explicit menu item for loading address and port groups that are defined in a file on a remote host. However, you can “load” groups by creating a script on a remote host (which includes your address or port group) and then running that script on your Administration Console host.

The following example shows a script that builds an address group:

```
bridge packetFilter addressGroup create
08-37-21-65-78-c4
08-32-18-55-40-a0
08-22-12-65-78-05
08-18-23-00-82-00
08-52-12-65-5f-22
08-25-43-41-6e-09
08-00-65-23-00-ee
08-5a-42-77-8a-01
08-22-13-66-00-2a
08-8e-54-11-78-3b
08-77-12-65-78-8c
q
```


When you run the script, your system address group is automatically created and stored on the system. For more information on running scripts, see “Running Scripts of Administration Console Tasks,” on page 2-13.




APPENDIXES

Appendix A Packet Filter Opcodes, Examples, and Syntax Errors

Appendix B Technical Support





PACKET FILTER OPCODES, EXAMPLES, AND SYNTAX ERRORS

This appendix:

- Describes the specific opcodes you can use when creating a packet filter
- Provides numerous examples of commonly used packet filters
- Describes the possible syntax errors you might receive when loading a packet filter



For information on creating and using packet filters, see Chapter 13.

Opcodes

Opcodes are instructions used in packet filter definitions. The available opcodes are described in this section:

name "<name>"

Description:

Assigns a user-defined <name> to the packet filter. The name may be any sequence of ASCII characters other than quotation marks. The name is limited to 32 characters. You can include only a single name statement in each packet filter program.

Storage Requirements:

2 + n bytes of packet filter storage, where n is the length of the <name>

pushField.size <offset>

Description:

Pushes a field from the target packet onto the stack. Packet data starting at <offset> is copied onto the stack. The most significant byte of the field is the byte at the specified offset. The size field of the instruction determines the number of bytes pushed. The pushField instruction provides direct access to any 1, 2, 4, or 6 byte field contained within the first 65535 bytes of the target packet.

Certain implementations of the packet filter language further limit the maximum offset, based on the packet lengths supported by the underlying network. Ethernet-based packet filters are limited to accessing fields in the first 1518 bytes of the target packet.

Specify the offset as an octal, decimal, or hexadecimal number.

- Precede an octal number by a "0".
- Precede a hexadecimal number by either "0x" or "0X".
- Use either upper or lower case letters for the hexadecimal digits "a" through "f".

Storage Requirements:

3 bytes

pushLiteral.size <value>

Description:

Pushes a literal constant <value> onto the stack. The most significant byte of the <value> is the first byte of the literal. Bytes are copied directly from the instruction stream onto the stack. The size field of the instruction determines number of bytes pushed.

Specify the value as either an octal, decimal, or hexadecimal number.

- Precede an octal number by a "0".
- Precede a hexadecimal number by either "0x" or "0X".
- Use either upper or lower case letters for the hexadecimal digits "a" through "f".

Storage Requirements:

1 (.b), 2 (.w), 4 (.l), or 6 (.a) bytes—depending on the size of the operand

pushTop

Description:

Pushes the current top of the stack onto the stack (that is, it reads the top of the stack and pushes the value onto the stack). The size of the contents of the stack determines the size of the push.

Storage Requirements:

1 byte

pushSAGM

Description:

Pushes the source address group mask (SAGM) onto the top of the stack. The SAGM is a bitmap representing the groups to which the source address of a packet belongs. This instruction pushes 4 bytes onto the stack.

Each address group is represented by a single bit in the SAGM.

Multicast addresses (including broadcast addresses) are in all groups.

Storage Requirements:

1 byte

pushDAGM

Description:

Pushes the destination address group mask (DAGM) onto the top of the stack. The DAGM is a bitmap representing the groups to which the destination address of a packet belongs. This instruction pushes 4 bytes onto the stack.

Each address group is represented by a single bit in the DAGM.

Multicast addresses (including broadcast addresses) are in all groups.

Storage Requirements:

1 byte

pushSPGM

Description:

Pushes the source port group mask (SPGM) onto the top of the stack. The SPGM is a bitmap representing the groups to which the source port of a packet belongs. This instruction pushes 4 bytes on to the stack.

Each port group mask is represented by a single bit in the SPGM bitmap. Port group masks are assigned to the bitmap in sequence, starting with port group mask 1 as the least significant bit through port group mask 32 as the most significant bit.

Storage Requirements:

1 byte

pushDPGM

Description:

Pushes the destination port group mask (DPGM) onto the top of the stack. The DPGM is a bitmap representing the groups to which the destination port of a packet belongs. This instruction pushes 4 bytes on to the stack.

Each port group mask is represented by a single bit in the DPGM bitmap. Port group masks are assigned to the bitmap in sequence, starting with port group mask 1 as the least significant bit through port group mask 32 as the most significant bit.

Storage Requirements:

1 byte

eq (equal)

Description:

Pops two values from the stack and compares them. If they are equal, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The contents of the stack determines the size of the operands.

Storage Requirements:

1 byte

ne (not equal)

Description:

Pops two values from the stack and compares them. If they are not equal, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Requirements:

1 byte

lt (less than)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is less than the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The contents of the stack determines the size of the operands.

Storage Requirements:

1 byte

le (less than or equal to)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is less than or equal to the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The contents of the stack determines the size of the operands.

Storage Requirements:

1 byte

gt (greater than)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is greater than the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The contents of the stack determines size of the operands.

Storage Requirements:

1 byte

ge (greater than or equal to)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is greater than or equal to the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The contents of the stack determines the size of the operands.

Storage Requirements:

1 byte

and (bit-wise AND)

Description:

Pops two values from the stack and pushes the bit-wise *AND* of these values back onto the stack. The contents of the stack determines the size of the operands and the result.

Storage Requirements:

1 byte

or (bit-wise OR)

Description:

Pops two values from the stack and pushes the bit-wise *OR* of these values back onto the stack. The size of the operands and the result are determined by the contents of the stack.

Storage Requirements:

1 byte

xor (bit-wise exclusive-OR)

Description:

Pops two values from the stack and pushes the bit-wise *exclusive-OR* of these values back onto the stack. The size of the operands and the result are determined by the contents of the stack.

Storage Requirements:

1 byte

not

Description:

Pops a byte from the stack; if its value is non-zero, a byte containing 0 is pushed back onto the stack. Otherwise, a byte containing the value is pushed back onto the stack.

Storage Requirements:

1 byte

accept

Description:

Conditionally accepts the packet that is being examined. Pops a byte from the stack. If its value is non-zero, the packet is accepted and evaluation of the filter ends immediately; otherwise, filter evaluation continues with the next instruction.

Storage Requirements:

1 byte

reject

Description:

Conditionally rejects the packet being examined. Pops a byte from the stack. If its value is non-zero, the packet is rejected and evaluation of the filter ends immediately; otherwise, filter evaluation continues with the next instruction.

Storage Requirements:

1 byte

shifl (shift left)

Description:

Pops two values from the stack and shifts the first operand left by the number of bits specified by the second operand. Bits shifted out of the left side of the operand are discarded, and zeros are shifted in from the right. The resulting value is pushed back onto the stack. The contents of the top of the stack determines the size of the first operand and the size of the result. The second operand is always 1 byte and only the low 5 bits of the byte are used as the shift count.

Storage Requirements:

1 byte

shifr (shift right)

Description:

Pops two values from the stack and shifts the first operand right by the number of bits specified by the second operand. Bits shifted out of the right side of the operand are discarded, and zeros are shifted in from the left. The resulting value is pushed back onto the stack. The contents of the stack determines the size of the first operand and the size of the result. The second operand is always 1 byte and only the low 5 bits of the byte are used as the shift count.

Storage Requirements:

1 byte

Packet Filter Examples

The following examples of packet filters, built using the packet filter language start with basic concepts.

Destination Address Filter

This filter operates on the destination address field of a frame. It allows packets to be forwarded that are destined for stations with an Organizationally Unique Identifier (OUI) of 08-00-02. To customize this filter to another OUI value, change the literal value loaded in the last **pushLiteral.l** instruction. Note that the OUI must be padded with an additional 00 to fill out the literal to 4 bytes.

```
name          "Forward to 08-00-02"
pushField.l   0          # Get first 4 bytes of
                  # destination address.
pushLiteral.l 0xffffffff # Set up mask to isolate first
                  # 3 bytes.
and
pushLiteral.l 0x08000200 # Top of stack now has OUI
                  # Load OUI value.
eq            # Check for match.
```

Source Address Filter

This filter operates on the source address field of a frame. It allows packets to be forwarded that are from stations with an OUI of 08-00-02. To customize this filter to another OUI value, change the literal value loaded in the last **pushLiteral.l** instruction. Note that the OUI must be padded with an additional 00 to fill out the literal to 4 bytes.

```
name          "Forward from 08-00-02"
pushField.l   6          # Get first 4 bytes of source
                  # address.
pushLiteral.l 0xffffffff # Set up mask to isolate first
                  # 3 bytes.
and
pushLiteral.l 0x08000200 # Top of stack now has OUI
                  # Load OUI value.
eq            # Check for match.
```

Length Filter

This filter operates on the length field of a frame. It allows packets to be forwarded that are less than 400 bytes in length. To customize this filter to another length value, change the literal value loaded in the **pushLiteral.w** instruction.

```
name          "Forward < 400"
pushField.w   12         # Get length field.
pushLiteral.w 400        # Load length limit.
lt            # Check for frame length <
                  # limit.
```

Type Filter This filter operates on the type field of a frame. It allows packets to be forwarded that are IP frames. To customize this filter to another type value, change the literal value loaded in the **pushLiteral.w** instruction.

```

name                "Forward IP frames"
pushField.w         12                # Get type field.
pushLiteral.w       0x0800           # Load IP type value.
eq                  # Check for match.

```

Ethernet Type IPX and Multicast Filter This filter *rejects* frames that have either a Novell IPX Ethernet type field (8134 hex) or a multicast destination address.

```

name                "Type > 900 or Multicast"
pushField.w         12                # Get type field.
pushLiteral.w       0x900            # Push type value to test
                                     # against.
gt                  # Is type field > 900 (hex)?
reject              # If yes: reject frame (done).
pushLiteral.b       0x01            # Multicast bit is low-order
pushField.b         0                # bit
and                 # Get 1st byte of destination
not                 # Isolate multicast bit
                                     # Top of stack 1 to accept,
                                     # 0 to reject

```

Multiple Destination Address Filter This filter operates on the destination address field of a frame. It allows packets to be forwarded that are destined for one of four different stations. To customize this filter to other destination stations, change the literal values.

```

name                "Forward to four stations"
pushField.a         0                # Get destination address.
pushTop             # Make 3 copies of address.
pushTop             #
pushTop             #
pushLiteral.a       0x367002010203 # Load allowed destination
                                     # address.
eq                  # Check for match.
accept              # Forward if valid address.
pushLiteral.a       0x468462236526 # Load allowed destination
                                     # address.
eq                  # Check for match.
accept              # Forward if valid address.
pushLiteral.a       0x347872927352 # Load allowed destination
                                     # address.
eq                  # Check for match.
accept              # Forward if valid address.
pushLiteral.a       0x080239572897 # Load allowed destination
                                     # address.
eq                  # Check for match.

```

Source Address and Type Filter

This filter operates on the source address and type fields of a frame. It allows XNS packets to be forwarded that are from stations with an OUI of 08-00-02. To customize this filter to another OUI value, change the literal value loaded in the last **pushLiteral.l** instruction. Note that the OUI must be padded with an additional 00 to fill out the literal to 4 bytes. To customize this filter to another type value, change the literal value loaded into the **pushLiteral.w** instruction.

```

name                "XNS from 08-00-02"
pushField.w         12                # Get type field.
pushLiteral.w       0x0600           # Load type value.
ne                  # Check for mis-match.
reject              # Toss any non-XNS frames.
pushLiteral.l       0xffffffff00     # Set up mask to isolate first 3
                                     # bytes.
pushField.l         6                # Get first 4 bytes of source
                                     # address.
and                 # Top of stack now has OUI.
pushLiteral.l       0x09000200       # Load OUI value.
eq                  # Check for match.

```

Accept XNS or IP Filter

This filter operates on the type field of a frame. It allows packets to be forwarded that are XNS or IP frame. Note the use of the **pushTop** instruction to make a copy of the type field.

```

name                "Forward IP or XNS"
pushField.w         12                # Get type field.
pushTop             # Push copy of type.
pushLiteral.w       0x0800           # Load IP type value.
eq                  # Check for match.
pushLiteral.w       0x0600           # Load XNS type value.
eq                  # Check for match.

```

XNS Routing Filter

This filter operates on the type and data fields of a frame. It discards all XNS routing packets.

```

name                "Drop XNS Routing"
pushField.w         12                # Get type field.
pushLiteral.w       0x0600           # Load XNS type value.
ne                  # Check for non-XNS packet.
accept             # Forward if non-XNS packet.
pushLiteral.b       0x01             # Load XNS routing type.
pushField.b         19                # Get XNS type.
ne                  # Check for non-XNS routing
                                     # packet.

```

Address Group Filter This filter accepts only frames whose source and destination address are in the same group.

```

name                "Forward Same Source and Destination"
pushSAGM            # Get source address group mask.
pushDAGM            # Get destination address group
                    # mask.
and                 # Compare if source and
destination         # groups are common members of
                    # an address group (result is
                    # either zero or non-zero)
                    # address group masks.
pushLiteral.1      0 # Put a zero on the stack.
ne                  # If not equal, returns a "one"
                    # to stack, resulting in packet
                    # forwarded.

```

Port Group Filter This filter discards all frames sourced from a port in group three or group eight.

```

name                "Discard Port Groups 3 and 8"
pushSPGM            # Get source port group mask.
pushLiteral.1      0x0084 # Select bits 3 and 8.
and                 # If port group bits 3 or 8 are
                    # common with SPGM, then
                    # non-zero value is
                    # pushed onto stack.
pushLiteral.1      0 # Push zero
eq                  # only if SPGM is not in port
                    # groups corresponding to bits
                    # 3 or 8, then packet is
                    # forwarded.

```

Common Syntax Errors

When you load a packet filter definition, the software checks the definition for syntax errors. The syntax errors and their causes are listed in Table A-1.


Table A-1 Syntax Errors When Loading Packet Filters

Syntax Error	Description
Opcode not found Unknown opcode	An opcode was expected on the line and was not found. The opcode must be one of those described in "Opcodes" on page A-1 and must include the size, if any. The opcode and size must be separated by a single "." with no intervening spaces. Any mix of uppercase and lowercase characters is permitted.
Operands are not the same size	The opcode requires two operands of the same size. The top two operands currently on the stack are of different sizes.
Stack underflow	The opcode requires one or more operands. An insufficient number of operands are currently on the stack.
Stack overflow	The opcode pushes an operand on the stack. The stack does not have sufficient room for the operand.
No result found on top of stack	The program must end with a byte operand on the top of the stack. After the last instruction in the program is executed, the stack is either empty or contains an operand other than a byte.
Extra characters on line	The source line contains extraneous characters that are not part of the instruction and are not preceded by a comment character (#).
Expected a byte operand	The opcode requires a byte operand as one of its parameters. The operand is of a size other than a byte.
Offset not found	The opcode requires an offset to be specified. None was found on the line.
Literal not found	The opcode requires a literal value to be specified. None was found on the line.
String not found	The opcode requires a quoted string to be specified. None was found on the line.

(continued)

Table A-1 Syntax Errors When Loading Packet Filters (continued)

Syntax Error	Description
Invalid characters in number	<p>The number specified as an offset or literal is improperly formatted. Possible causes are 1) lack of white space setting off the number, and 2) invalid characters in the number.</p> <p>Note: The radix of the number is determined by the first 1 or 2 characters of the number.</p> <ul style="list-style-type: none"> ■ A number with a leading "0x" or "0X" is treated as hexadecimal. ■ All other numbers with a leading 0 are treated as octal. ■ All other numbers are treated as decimal.
Number is too large	<p>The number specified as an offset or literal is too large. An offset is limited to 1518 minus the size of the operand. For example, the offset for pushField.b can be no more than 1517, and the offset for pushField.w no more than 1516. A literal value is limited to the number of bytes in the operand size (1, 2, 4, or 6).</p>
Missing open quote on string	<p>The string specified does not have a starting quotation mark (").</p>
String is too long	<p>The string specified is too long. Strings are limited to 32 characters exclusive of the opening and closing quotation marks.</p>
Missing close quote on string	<p>The string specified does not have an ending quotation mark (").</p>
Multiple name statements in program	<p>More than one name statement was found in the program. Only a single name statement is allowed.</p>
Program too large	<p>The program exceeds the maximum size allowed. The causes of this error include a source definition exceeding 4096 bytes, a stored format exceeding 254 bytes, or a run-time format exceeding 2048 bytes. All of these boundary conditions are checked when the filter is loaded. See Table 13-2 for more information on packet filter sizes.</p>
Too many errors - compilation aborted	<p>The program contains an excessive number of errors. No further syntax errors will be reported. The program stops compiling when this condition occurs.</p>



TECHNICAL SUPPORT

3Com provides easy access to technical support information through a variety of services. This appendix describes these services.

Online Technical Services

3Com offers worldwide product support seven days a week, 24 hours a day, through the following online systems:

- 3Com Bulletin Board Service (3ComBBS)
- World Wide Web site
- 3ComForum on CompuServe® online service
- 3ComFactsSM automated fax service

3Com Bulletin Board Service

3ComBBS contains patches, software, and drivers for all 3Com products, as well as technical articles. This service is available via modem or ISDN, 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. Call the telephone number nearest you:

Country	Data Rate	Telephone Number
Australia	up to 14400 bps	(61) (2) 9955 2073
France	up to 14400 bps	(33) (1) 69 86 69 54
Germany	up to 9600 bps	(49) (89) 627 32 188 or (49) (89) 627 32 189
Hong Kong	up to 14400 bps	(852) 2537 5608
Italy (fee required)	up to 14400 bps	(39) (2) 273 00680
Japan	up to 14400 bps	(81) (3) 3345 7266
Singapore	up to 14400 bps	(65) 534 5693
Taiwan	up to 14400 bps	(886) (2) 377 5840
United Kingdom	up to 28800 bps	(44) (1442) 278278
United States	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 is a CompuServe service containing patches, software, drivers, and technical articles about all 3Com products, as well as a messaging section for peer support. To use 3ComForum, you need a CompuServe account.

To use 3ComForum:

- 1 Log on to CompuServe.
- 2 Enter **go threecom**
- 3 Press [Return] to see the 3ComForum Main menu.

3ComFactsSM Automated Fax Service

3Com Corporation's interactive fax service, 3ComFacts, provides data sheets, technical articles, diagrams, and troubleshooting instructions on 3Com products 24 hours a day, seven days a week.

Call 3ComFacts service using your Touch-Tone® telephone. International access numbers:

Country	Telephone Number
Hong Kong	(852) 2537 5610
United Kingdom	(44) (1442) 278279
United States	(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	United Kingdom	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 United States and Canada, call **(800) 876-3266** for customer service.

If you are outside the United States and Canada, contact your local 3Com sales office to find your authorized service provider:

Country	Telephone Number	Country	Telephone Number
Australia*	(1800) 678 515	Japan	(81) (3) 3345 7251
Belgium*	0800 71429	Mexico	(525) 531 0591
Brazil	(55) (11) 546 0869	Netherlands*	06 0227788
Canada	(416) 498-3266	Norway*	800 13376
Denmark*	800 17309	Singapore	(65) 538 9368
Finland*	0800 113153	South Africa	(27) (11) 803 7404
France*	05 917959	Spain*	900 983125
Germany*	0130 821502	Sweden*	120 795482
Hong Kong	(852) 2501 1111	Taiwan	(886) (2) 577 4352
Ireland*	1 800 553117	United Arab Emirates	(971) (4) 349049
Italy*	1678 79489	United Kingdom*	0800 966197
		United States	(1) (408) 492 1790

* These numbers are toll-free.

Returning Products for Repair

Before you send a product directly to 3Com for repair, you must first be assigned a Return Materials Authorization (RMA) number. A product sent to 3Com without an RMA number will be returned to the sender unopened, at the sender's expense.

To obtain an RMA number, call or fax:

Country	Telephone Number	Fax Number
United States and Canada	(800) 876 3266, option 2	(408) 764 7120
Europe	31 30 60 29900, option 5	(44) (1442) 275822
Outside Europe, the United States, and Canada	(1) (408) 492 1790	(1) (408) 764 7290

8/20/96

INDEX

Numerics

3Com Bulletin Board Service (3ComBBS) B-1
3Com sales offices B-4
3ComFacts B-3
3ComForum B-2
802.1d bridging, enabling mode 11-4

A

abort
 at prompts 2-9
 enabling CTL+C 2-12
accept opcode 13-8, A-7
access levels 2-1
address
 adding static 12-12
 aging time 11-7
 filters A-9
 flushing 12-13
 for SNMP trap reporting 3-18
 freezing 12-14
 in routing table 3-7
 IP 3-4
 IP to MAC, translating 3-11
 maximum number in group 14-7
 removing static 12-12
address group
 adding addresses 14-7 to 14-9
 as filtering criteria 14-1
 copying 14-7
 creating 14-4
 deleting 14-6
 displaying contents 14-3
 listing 14-2
 loading on system 14-11
 removing addresses 14-9
 used in packet filter 14-1
Address Resolution Protocol. *See* ARP
address threshold
 values 11-7
addresses
 listing for ATM 9-14
addressThresholdEvent 11-7

administer access example 2-2
Administration Console
 command strings 2-9
 Control keys 2-12
 entering values 2-9
 exiting 2-17
 initial user access 2-1
 interface parameters 2-10, 2-11
 locking 2-12
 menu descriptions 2-3 to 2-7
 menu hierarchy, moving up 2-9
 menu options, selecting 2-8
 password access 2-1, 4-2
 preventing disconnections 2-12
 restarting 2-11
 screen height, setting 2-10
 scripts 2-13
 top-level menu 2-3
advertisement address 3-5
aging time
 setting for bridge 11-7
 values 11-7
analyzer
 connecting 10-4
 MAC address display 10-4
 removing port 10-4
 setting up monitored port 10-5
and (bit-wise AND) opcode A-6
AppleTalk
 packet filter 13-9
ARP
 See also ARP cache
ARP cache
 defined 3-11
 flushing 3-12
 removing entry 3-11
ASCII-based editor
 and scripts 2-13
 for packet filters 13-17
ATM
 Broadcast and Unknown Server (BUS) 9-2
 Classical IP (CLIP) 9-1
 commands, quick 1-6, 1-7
 LAN Emulation (LANE) 9-1, 9-2

- LAN Emulation Client (LEC) 9-2
- LAN Emulation Configuration Server (LECS) 9-2
- LAN Emulation Server (LES) 9-2
 - menu 2-5
- ATM port
 - administering 9-16 to 9-22
 - displaying statistics 9-16
 - labeling 9-19
 - listing network prefixes and addresses 9-14
 - virtual channel connection information 9-19 to 9-22

B

- backup
 - saving NV data 6-2
- baseline
 - displaying current 5-2
 - enabling and disabling 5-2
 - reasons for 5-1
 - setting 5-2
- baud rate
 - serial port 3-2
- bell, warning 4-2
- blocking state 12-5
- bridge
 - address threshold, setting 11-7
 - aging time, setting 11-7
 - designated 12-3
 - IP fragmentation, enabling 11-6
 - IPX Snap Translation, enabling 11-6
 - menus 2-6
 - mode, setting 11-4
 - Spanning Tree
 - bridge priority, setting 11-8
 - enabling 11-8
 - forward delay, setting 11-10
 - hello time, setting 11-10
 - maximum age, setting 11-9
 - statistics, displaying 11-1
- bridge port
 - MAC addresses
 - adding 12-12
 - flushing 12-13
 - freezing 12-14
 - listing 12-11
 - removing 12-12
 - multicast limit, setting 12-7
 - Spanning Tree
 - enabling 12-8
 - path cost, setting 12-9
 - port priority, setting 12-10
 - states defined 12-5

- statistics, displaying 12-1
 - bridge. *See also* packet filter
- bridging
 - commands, full 11-1 to 11-11
 - commands, quick 1-4
 - mode defined 11-4
- Broadcast and Unknown Server (BUS) 9-2
- bulletin board service B-1

C

- Classical IP (CLIP) over ATM 9-1
- code. *See* scripting *and* packet filter
- commands
 - and entering values 2-9
 - quick 1-1
 - using 2-9
- community strings
 - setting 3-16
 - values 3-15
- CompuServe® B-2
- connect state
 - UME 9-15
- connectPolicy
 - configuring 8-4
- Control keys
 - enabling 2-12
- conventions
 - notice icons 3
- cost
 - of IP interface 3-5
 - Spanning Tree settings 11-4, 12-3, 12-9
 - See also* metric
- CTL+C (abort) 2-12
- CTL+X (reboot) 2-12

D

- datagrams, statistics 3-14
- date
 - formats 4-4
 - setting system 4-4
- default route
 - defined 3-8
 - displayed 3-9
 - removing 3-11
 - setting 3-10
- destination address
 - for SNMP trap reporting 3-18
- destination address group mask (DAGM) 14-1
- destination IP address
 - in routing table 3-7

destination port group mask (DPGM) 14-1
direct, route status 3-8
documentation
 for the LANplex system 4
duplex mode
 setting 7-7

E

editor
 for packet filters 13-18
 for scripts 2-13
EMACS editor 2-13, 13-17
Emulated LAN (ELAN)
 creating 9-4
eq opcode A-4
Ethernet
 analyzing segments 10-1 to 10-6
 commands, quick 1-6
 fragmenting packets 11-6
 menus 2-4
 packet fields 13-6
 portState 7-9
 station MAC addresses 12-11
Ethernet address
 and restoring NV data 6-3
 for the monitored port 10-5
Ethernet port
 analyzer attached 10-3
 displaying information 7-1
 label 7-4
 labeling 7-8
 See also Roving Analysis
 setting state (on-line or off-line) 7-9
 static MAC addresses 12-12
 statistics 7-3
Express switching
 enabling mode 11-4

F

fan, warning 4-2
Fast Ethernet
 Roving Analysis, and 10-1
Fast Ethernet ports
 full-duplex mode 7-7
fax service B-3
FDDI
 commands, quick 1-6, 1-7
 fragmenting packets 11-6
 management 8-1
 menus 2-5

 packet fields 13-6
 port label 8-20
 rings 8-6
 station MAC addresses 12-11
 wrapped ring 8-6
FDDI MAC
 condition report 8-16
 defined 8-9
 FrameErrorThreshold, setting 8-16
 LLC Service, enabling 8-18
 NotCopiedThreshold, setting 8-17
 statistics, displaying 8-10
FDDI path
 defined 8-6
 maxT-Req, setting 8-9
 statistics, displaying 8-6
 tmaxLowerBound, setting 8-8
 tvxLowerBound, setting 8-7
FDDI port
 and roving analysis 10-6
 defined 8-19
 labeling 8-22
 lerAlarm, setting 8-20
 lerCutoff, setting 8-21
 statistics, displaying 8-19
FDDI station
 and SMT 8-1
 and SRFs 8-2, 8-5
 connection policies, setting 8-4
 defined 8-1
 statistics, displaying 8-2
 status reporting, enabling 8-5
 T-notify, setting 8-5
filter id 13-2
flushing
 ARP cache 3-12
 learned routes 3-10
 MAC addresses 12-13
 SNMP trap addresses 3-19
forward delay 11-10
forwarding state 12-5
FrameErrorThreshold
 defined 8-16
 setting 8-16
freezing addresses 12-14
ftp
 IP address 3-1, 3-4
full-duplex mode 7-7

G

- gateway
 - IP address 3-8
 - See also* route
- ge opcode A-6
- group address
 - Spanning Tree, setting 11-11
- group. *See* address group *or* port group
- gt opcode A-6

H

- hello time 11-10
- Help
 - Administration Console 2-16
 - topical 2-16

I

- ICMP
 - and ping 3-13
 - echo (request and reply) 3-13
- in-band management 3-2
- instructions
 - opcodes 13-5, A-1
 - operands 13-5, 13-7
 - operators 13-7
- interface
 - Administration Console parameters 2-10, 2-11
 - defining 3-6
 - displaying 3-5
 - parts of 3-4
 - removing definition 3-7
- Internet Control Message Protocol. *See* ICMP
- IP
 - address translation 3-11
 - ARP cache 3-11
 - interface 3-4
 - management access 3-1
 - menus 2-7
 - pinging 3-13
 - RIP mode 3-12
 - route table 3-8
 - routes 3-7
 - statistics, displaying 3-14
- IP address
 - and restoring NV data 6-3
 - configuring 3-5
 - for IP interface 3-4
 - in routing table 3-7
- IP fragmentation

- enabling 11-6
- IP interface
 - address 3-4
 - advertisement address 3-5
 - cost 3-5
 - defining 3-4, 3-6
 - displaying 3-5
 - removing definition 3-7
 - state 3-5
 - subnet mask 3-4
- IP packets filter 13-12, 13-16
- IP route
 - default 3-8, 3-10
 - defining static 3-9
 - destination address 3-7
 - gateway IP address 3-8
 - metric 3-8
 - removing from table 3-9, 3-10
 - status 3-8
 - subnet mask 3-7
- IPX Snap Translation
 - enabling 11-6

L

- LAN Emulation (LANE) 9-1, 9-2
- LAN Emulation Client (LEC) 9-2
 - administering 9-5 to 9-12
 - defining 9-10
 - displaying statistics 9-5
 - joining existing ELAN 9-5
 - modifying parameters 9-9
 - removing 9-12
- LAN Emulation Configuration Server (LECS) 9-2
- LAN Emulation Server (LES) 9-2
- LANplex
 - administration overview 1-1
 - and network monitoring 10-1
 - bell warning 4-2
 - documentation 4
 - fan warning 4-2
 - naming 4-3
 - NV data restoration 6-3
 - power supply warning 4-2
 - rebooting 4-5
 - resetting to system defaults 6-6
 - system backup 6-2
 - system configuration, displaying 4-1
 - system date and time 4-4
 - temperature warning 4-2
 - user access levels 2-1
 - warning messages 4-2

le opcode A-5
 learned, route status 3-8
 learning state 11-10, 12-5
 LER (Link Error Rate)
 alarm value 8-20
 cutoff value 8-21
 lerAlarm
 and lerCutoff value 8-21
 defined 8-20
 setting 8-21
 lerCutoff
 and lerAlarm value 8-21
 defined 8-21
 Link Error Rate. *See* LER
 listening state 11-10, 12-5
 LLC
 enabling 8-18
 service description 8-18
 Logical Link Control. *See* LLC
 lt opcode A-5

M

MAC (Media Access Control) address
 adding 12-12
 and ARP 3-11
 configuring 12-11
 displaying 12-11
 dynamic to static 12-14
 flushing 12-13
 removing static 12-12
 roving analysis configuration 10-3
 management
 and naming the system 4-3
 configuring system access 3-1 to 3-13
 FDDI 8-1
 in-band 3-2
 IP interface 3-1, 3-4
 out-of-band 3-2
 port labels 7-8, 8-22
 setup, quick commands 1-3
 SNMP community strings 3-15
 system name 4-3
 Transcend® Enterprise Manager 1-1
 maximum age 11-9
 maxT-Req
 defined 8-9
 setting 8-9
 menu
 analyzer (roving analysis) 2-8
 and command strings 2-9
 ATM 2-5

 bridge 2-6
 ethernet 2-4
 fddi 2-5
 IP 2-7
 moving up hierarchy 2-9
 selecting options 2-8
 SNMP 2-7
 system 2-4
 metric
 in routing table 3-8
 modem 3-1
 connecting to 3-3
 escape sequence 3-3
 serial port speed 3-3
 modules
 revision numbers 4-2
 multicast frames
 and packet filters 13-1
 multicast limit
 configuring 12-7
 defined 12-7

N

name opcode A-1
 naming the LANplex 4-3
 ne opcode A-5
 neighbor notification
 and LLC Service 8-18
 network monitoring. *See* roving analysis *and* analyzer
 network supplier support B-3
 network troubleshooting 10-1
 not opcode A-7
 NotCopiedThreshold
 defined 8-17
 setting 8-17
 Novell
 in packet filter A-10
 NV data
 and packet filters 13-3
 backup 6-1
 contents saved 6-1
 examining a saved file 6-5
 file information 6-1
 resetting 6-6
 restoring 6-3
 saving 6-2
 transferring 6-1

O

off-line port state 7-9
 on-line Help 2-16
 on-line port state 7-9
 on-line technical services B-1
 opcode
 and packet filter language 13-4
 and writing packet filters 13-10
 descriptions A-1 to A-8
 operand 13-5
 and opcodes 13-7
 sizes supported 13-5
 operator
 and opcodes 13-7
 purpose 13-7
 or opcode A-7
 OUI
 in packet filter A-11
 out-of-band management 3-2

P

packet
 Ethernet type 13-6
 FDDI type 13-6
 fields for operands 13-7
 packet filter
 address group example 14-1
 assigning to ports 13-22
 basic elements 13-6
 concepts 13-4 to 13-11
 correcting errors 13-21
 creating 13-3 to 13-17
 definitions 13-3
 deleting 13-20
 displaying contents 13-3
 editing 13-20
 editor
 commands 13-19
 description 13-18
 using 13-18
 examples 13-11 to 13-17, A-9 to A-12
 external editor 13-20
 filter id 13-2
 filtering criteria, groups 14-1
 instructions 13-5
 language description 13-3, 13-4
 listing 13-2
 loading 13-22
 opcodes A-1
 operands 13-5
 port group example 14-2
 procedure for writing 13-10
 processing paths 13-1, 13-22
 pseudocode 13-12
 run-time storage 13-10
 sequential tests 13-8
 stack 13-5
 storage space 13-9
 syntax errors A-13, A-14
 unassigning from ports 13-24
 See also address group *and* port group
 password
 configuring 4-2
 initial system access 2-1
 levels of user access 2-1
 path cost
 defined 12-9
 setting 12-9
 path. *See* FDDI path *and* backplane paths
 PHY
 and FDDI ports 8-19
 ping
 IP station 3-13
 PMD
 and FDDI ports 8-19
 port
 bridging priority 12-10
 for analyzer 10-3
 label 8-20
 maximum number in group 14-7
 path cost 12-9
 speed, setting 3-2, 3-3
 state, setting 7-9
 types 8-19
 See also FDDI port, Ethernet port, and ATM port
 port group
 adding ports 14-7 to 14-9
 as filtering criteria 14-1
 copying 14-7
 creating 14-4
 deleting 14-6
 displaying contents 14-3
 listing 14-2
 loading on system 14-11
 removing ports 14-9
 used in packet filter 14-2
 power supply warning 4-2
 probe. *See* roving analysis *and* analyzer
 pushDAGM opcode 14-1, A-3
 pushDPGM opcode 14-1, A-4
 pushField.opcode A-2
 pushLiteral.opcode A-2
 pushSAGM opcode 14-1, A-3

pushSPGM opcode 14-1, A-4
 pushTop opcode A-3

R

read access example 2-3
 reboot
 enabling CTL+X 2-12
 resetting the system 4-5
 reboot system 2-11
 receive all
 packet processing path 13-1
 receive multicast
 packet processing path 13-1
 reject opcode 13-8, A-8
 remote sessions
 enabling timeout 2-13
 setting timeout interval 2-13
 restart, Administration Console 2-11
 returning products for repair B-4
 RIP
 and broadcast address 3-5
 default mode 3-12
 displaying state 3-5
 Off mode 3-12
 Passive mode 3-12
 setting mode 3-12
 rlogin
 and exiting the Console 2-17
 and rebooting the system 4-5
 route
 default 3-8
 defining static 3-9
 destination IP address 3-7
 flushing learned routes 3-10
 gateway IP address 3-8
 metric 3-8
 removing default 3-11
 removing from table 3-9, 3-10
 See also routing table
 status 3-8
 subnet mask 3-7
 Routing Information Protocol. *See* RIP
 routing table
 contents 3-7
 default route, setting 3-10
 display routes 3-8
 flushing learned routes 3-10
 removing default route 3-11
 removing routes 3-9
 roving analysis
 adding analyzer port 10-3
 and Spanning Tree 10-4

configuration rules 10-2
 configuration, displaying 10-3
 configuring 10-3
 defined 10-1
 menu 2-8
 process overview 10-2
 remote Fast Ethernet connections, and 10-1
 removing analyzer port 10-4
 starting port monitoring 10-5
 stopping port monitoring 10-6

S

SAGM (source address group mask) 14-1
 screen height
 adjusting 2-10
 scripts for the Administration Console
 examples 2-15
 running 2-13
 serial port
 reasons for disconnecting 2-12
 serial port (modem)
 setting baud rate 3-3
 serial port (terminal)
 setting baud rate 3-2
 Service Access Points (SAPs)
 and packet filters 13-4
 shiftl opcode A-8
 shiftr opcode A-8
 SMT (Station Management)
 and FDDI stations 8-1
 lerAlarm value 8-21
 lerCutoff value 8-21
 SMT event
 enabling proxying 3-20
 proxying defined 3-19
 Sniffing. *See* roving analysis *and* analyzer
 SNMP
 community strings
 setting 3-16
 values 3-15
 displaying configurations 3-15
 menus 2-7
 proxying remote SMT events 3-20
 trap reporting
 and SMT event proxying 3-19
 configuring destinations 3-18
 displaying configuration 3-16
 flushing addresses 3-19
 See also trap *and* community strings

- SNMP agent
 - accessing through IP 3-1
 - defined 3-15
 - SNMP trap
 - Address Threshold 3-17
 - addressThresholdEvent 11-7
 - Authentication Failure 3-17
 - Coldstart 3-17
 - Link Down 3-17
 - Link Up 3-17
 - MAC Duplicate Address Condition 3-17
 - MAC Frame Error Condition 3-17
 - MAC Neighbor Change 3-17
 - MAC Not Copied Condition 3-17
 - MAC Path Change 3-17
 - New Root 3-17
 - Port EB Error Condition 3-17
 - Port LER Condition 3-17
 - Port Path Change 3-17
 - Port Undesired Connection 3-17
 - SMT Hold Condition 3-17
 - SMT Peer Wrap Condition 3-17
 - System Overtemperature 3-17
 - Topology Change 3-17
 - socket values filter 13-12, 13-15
 - software
 - backup NV data 6-1, 6-2
 - build date and time 4-2
 - from factory 1-1
 - version number 4-2
 - source address group mask (SAGM) 14-1
 - source port group mask (SPGM) 14-1
 - Spanning Tree Protocol. *See* STP
 - SPGM (source port group mask) 14-1
 - SRF (Status Report Frames)
 - and FDDI stations 8-2, 8-5
 - and IerAlarm 8-20
 - stack 13-5
 - state
 - of IP interface 3-5
 - static route status 3-8
 - station. *See* FDDI station
 - Station Management. *See* SMT
 - statistics
 - baselining 5-1
 - Ethernet ports 7-3
 - FDDI MAC 8-10, 8-11
 - FDDI path 8-6
 - FDDI station 8-2
 - IP 3-14
 - Status Report Frames. *See* SRF
 - status reporting
 - configuring 8-5
 - defined 8-5
 - STP (Spanning Tree Protocol)
 - bridge priority, setting 11-8
 - designated bridge 12-3
 - designated cost 12-3
 - designated port 12-3
 - designated root 12-3
 - enabling on bridge 11-8
 - enabling on bridge port 12-8
 - forward delay, setting 11-10
 - group address, setting 11-11
 - hello time, setting 11-10
 - maximum age, setting 11-9
 - port priority 12-10
 - states 12-5
 - subnet mask
 - for IP address 3-4
 - in routing table 3-7
 - system configuration
 - displaying 4-1
 - system menus 2-4
-
- T**
- T_Opr 8-9
 - technical support B-1
 - telnet
 - enabling timeout 2-13
 - rebooting the system 4-5
 - setting timeout interval 2-13
 - temperature, warning 4-2
 - terminal emulation
 - and the serial port 3-1
 - terminal serial port
 - setting baud rate 3-2
 - text editor, built-in 13-18
 - time
 - displaying system up time 4-5
 - formats 4-4
 - setting system 4-4
 - timing out route status 3-8
 - tmaxLowerBound
 - defined 8-8
 - setting 8-8
 - T-notify
 - configuring 8-5
 - defined 8-5
 - token
 - and FDDI MAC 8-9
 - transmit all
 - packet processing path 13-1
 - transmit multicast

- packet processing path 13-1
- Transparent
 - enabling mode 11-4
- trap reporting
 - configuring destinations 3-18
 - flushing addresses 3-19
 - removing destinations 3-19
- T-Req 8-9
- tvxLowerBound
 - defined 8-7
 - setting 8-8

U

- UNI Management Entity (UME)
 - administering 9-12 to 9-16
 - connect state, setting 9-15
 - displaying information 9-13
 - virtual channel identifier, setting 9-16
 - virtual path identifier, setting 9-15
- UNIX
 - and terminal emulation with LANplex 3-1
- up time
 - displaying 4-5

V

- vi editor 2-13, 13-17
- virtual channel connection
 - displaying 9-19
- virtual channel identifier
 - setting 9-16
- virtual path identifier
 - setting 9-15

W

- warning messages for system 4-2
- wrapped ring 8-6
- write access example 2-2

X

- XNS
 - in packet filter 13-12, 13-14, A-11
- xor opcode A-7

