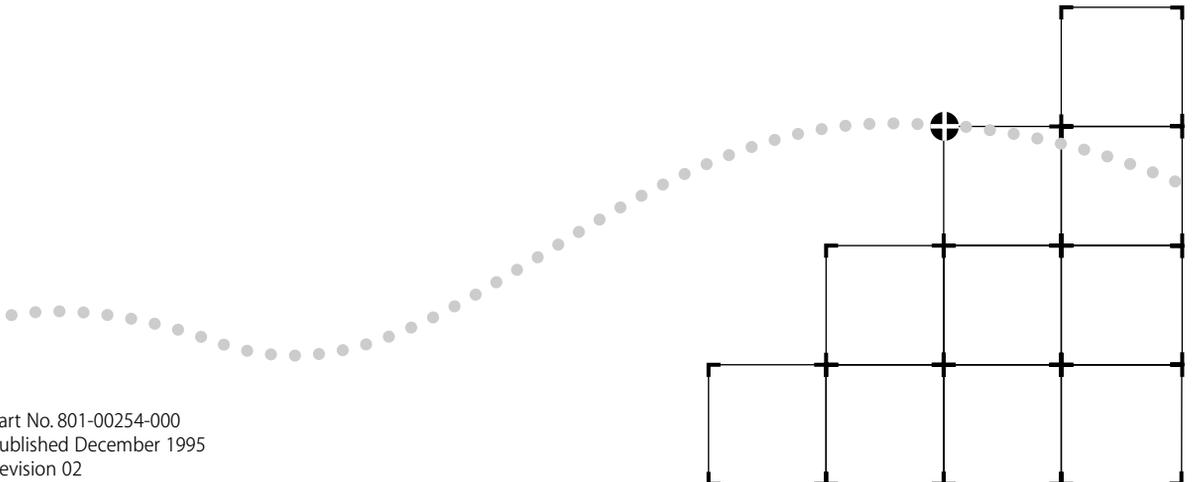




LANPLEX 6000 ADMINISTRATION CONSOLE USER GUIDE



Part No. 801-00254-000
Published December 1995
Revision 02

3Com Corporation ■ 5400 Bayfront Plaza ■ Santa Clara, California ■ 95052-8154

© 3Com Corporation, 1995. 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 Unpublished rights reserved under the copyright laws of the United States.

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

3Com, LANplex, LanScanner, LinkBuilder, NETBuilder, NETBuilder II, ViewBuilder, EtherDisk, EtherLink, EtherLink Plus, EtherLink II, TokenLink, TokenLink Plus, and TokenDisk are registered trademarks of 3Com Corporation. 3Com Laser Library, 3TECH, Boundary Routing, CacheCard, FDDILink, LinkSwitch, NetProbe, Parallel Tasking, SmartAgent, Star-Tek, and Transcend are trademarks of 3Com Corporation. 3ComFacts, Ask3Com, CardFacts, NetFacts, and CardBoard are service marks of 3Com Corporation.

Sniffer is a registered trademark of Network General Corp. CompuServe is a registered trademark of CompuServe, Inc.

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

Guide written, edited, and illustrated by Patricia Crawford, Lynne Gelfand, Michael Jenness, Patricia L. Johnson, Michael Taillon, and Iain Young.

CONTENTS

ABOUT THIS GUIDE

- Introduction 1
- How to Use This Guide 2
- Conventions 3
- LANplex 6000 Documentation 4
- Documentation Comments 6

PART I GETTING STARTED

1 LANPLEX 6000 ADMINISTRATION OVERVIEW

- About LANplex Administration 1-1
- Configuration Options 1-2

2 HOW TO USE THE ADMINISTRATION CONSOLE

- 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
 - Console Menu Structure 2-4
 - System Menu 2-4
 - Ethernet Menu 2-4
 - FDDI Menu 2-5
 - Token Ring Menu 2-6
 - Bridge Menu 2-6
 - IP Menu 2-7
 - SNMP Menu 2-8
 - Analyzer Menu 2-8
 - Selecting Menu Options 2-9
 - Entering a Command String 2-9
 - Entering Values 2-10
 - Getting Out 2-10

- Console Interface Parameters 2-11
 - Adjusting the Screen Height 2-11
 - Preventing Console Disconnections 2-12
- Configuring Control Panel Write Access 2-13
 - Disabling Reboot and Abort Keys 2-13
- Running Scripts of Console Tasks 2-14
- Getting Help in the Console 2-16
 - Online Help 2-16
 - Outlining 2-16
- Exiting the Administration Console 2-17

PART II SYSTEM-LEVEL FUNCTIONS

3 CONFIGURING MANAGEMENT ACCESS TO THE SYSTEM

- About Management Access 3-1
 - Using a Serial Connection 3-1
 - Using an IP Interface 3-2
 - In-band or Out-of-band? 3-2
- Setting up the Terminal Serial Port 3-3
- Setting up the Modem Serial Port 3-3
 - Setting the Port Speed 3-4
 - Configuring the External Modem 3-4
- Setting up an IP Interface for Management 3-5
 - General Setup Process 3-5
 - Administering Interfaces 3-5
 - Displaying Interfaces 3-7
 - Defining an Interface 3-7
 - Modifying an Interface. 3-9
 - Removing an Interface 3-9
 - Administering Routes 3-9
 - Displaying the Routing Table 3-11
 - Defining a Static Route 3-11
 - Removing a Route 3-12
 - Flushing a Route 3-12
 - Setting the Default Route 3-13
 - Removing the Default Route 3-13
 - Administering the ARP Cache 3-14
 - Displaying the ARP Cache 3-14
 - Removing an ARP Cache Entry 3-14
 - Flushing an ARP Cache Entry 3-15
- Setting the RIP Mode 3-15

- Pinging an IP Station 3-16
- Displaying IP Statistics 3-17
- Setting Up SNMP on Your System 3-18
 - Configuring the SNMP Mode 3-18
 - Configuration Guidelines 3-19
- Displaying SNMP Settings 3-20
 - Setting the Mode 3-21
- Configuring Community Strings 3-21
- Administering SNMP Trap Reporting 3-23
 - Displaying Trap Information 3-23
 - Configuring Trap Reporting 3-25
 - Removing Trap Destinations 3-26
 - Flushing Trap Destinations 3-27
 - Setting up SMT Event Proxying 3-27

4 ADMINISTERING YOUR SYSTEM ENVIRONMENT

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

5 UPDATING SYSTEM SOFTWARE

- About Updating Software 5-1
- Copying Software to a Hard Disk 5-1
 - Copying to UNIX 5-2
 - Copying to DOS 5-3
- Loading Software 5-4

6 BASELINING STATISTICS

- About Setting Baselines 6-1
- Displaying the Current Baseline 6-1
- Setting Baselines 6-2
- Enabling or Disabling Baselines 6-2

7 SAVING, RESTORING, AND RESETTNG NONVOLATILE DATA

- About Working with Nonvolatile Data 7-1
- Saving NV Data 7-2
- Restoring NV Data 7-3
- Examining a Saved NV Data File 7-6
- Resetting NV Data to Defaults 7-7

PART III ETHERNET, FDDI, & TOKEN RING

8 ADMINISTERING ETHERNET PORTS

- Displaying Ethernet Port Information 8-1
 - Frame Processing and Ethernet Statistics 8-5
 - Labeling a Port 8-7
 - Setting the Port State 8-8
-

9 ADMINISTERING FDDI RESOURCES

- About Configuring FDDI Resources 9-1
 - The Backplane
 - Path Mode 9-2
 - Modules and FDDI Resources 9-3
- Configuring the Backplane Path Mode 9-3
- Administering FDDI Stations 9-5
 - Displaying Station Information 9-5
 - Setting the Connection Policies 9-7
 - Setting Neighbor Notification Timer 9-9
 - Enabling/Disabling Status Reporting 9-10
- Administering FDDI Paths 9-11
 - Displaying Path Information 9-11
 - Setting tvxLowerBound 9-13
 - Setting tmaxLowerBound 9-14
 - Setting maxT-req 9-15
- Administering FDDI MACs 9-16
 - Displaying MAC Information 9-16
 - Frame Processing and FDDI MAC Statistics 9-22
 - Setting the Frame Error Threshold 9-23
 - Setting the Not Copied Threshold 9-24
 - Enabling/Disabling LLC Service 9-25
 - Assigning MACs to Stations in Multi-station Mode 9-25
 - Setting the MAC Paths 9-26
- Administering FDDI Ports 9-28
 - Displaying Port Information 9-28
 - Setting lerAlarm 9-30
 - Setting lerCutoff 9-31
 - Setting Port Labels 9-32
 - Assigning Ports to Stations in Multi-Station Mode 9-32
 - Setting the Port Paths 9-33

10 ADMINISTERING TOKEN RING PORTS

- Displaying Token Ring Port Information 10-1
 - Frame Processing and Token Ring Statistics 10-6
- Labeling a Port 10-8
- Setting the Port State 10-8
- Setting the Port Speed 10-9
- Setting the Port Mode 10-9

11 SETTING UP THE SYSTEM FOR ROVING ANALYSIS

- About Roving Analysis 11-1
- Displaying the Roving Analysis Configuration 11-3
- Adding an Analyzer Port 11-4
- Removing an Analyzer Port 11-5
- Starting Port Monitoring 11-6
- Stopping Port Monitoring 11-7

PART IV BRIDGING

12 ADMINISTERING THE BRIDGE

- Displaying Bridge Information 12-1
- Setting the Bridging Mode 12-4
- Enabling/ Disabling IP Fragmentation 12-6
- Enabling/ Disabling IPX Snap Translation 12-6
- Setting Protocol Address Translation from Token Ring to FDDI 12-7
- Setting the Address Threshold 12-8
- Setting the Aging Time 12-8
- Administering STP Bridge Parameters 12-9
 - Enabling/Disabling STP on a Bridge 12-9
 - Setting the Bridge Priority 12-10
 - Setting the Bridge Maximum Age 12-10
 - Setting the Bridge Hello Time 12-11
 - Setting the Bridge Forward Delay 12-12

13 ADMINISTERING BRIDGE PORTS

- Displaying Bridge Port Information 13-1
 - Frame Processing and Bridge Port Statistics 13-6
- Setting the Multicast Limit 13-7
- Administering STP Bridge Port Parameters 13-8
 - Enabling/Disabling STP on a Port 13-8

- Setting the Port Path Cost 13-9
- Setting the Port Priority 13-10
- Setting the Source Route Ring Number 13-11
- Administering Port Addresses 13-12
 - Listing Addresses 13-12
 - Adding New Addresses 13-13
 - Removing Addresses 13-13
 - Flushing All Addresses 13-14
 - Flushing Dynamic Addresses 13-15
 - Freezing Dynamic Addresses 13-15

14 CREATING AND USING PACKET FILTERS

- About Packet Filtering 14-1
- Listing Packet Filters 14-2
- Displaying Packet Filters 14-3
- Creating Packet Filters 14-3
 - Concepts for Writing a Filter 14-4
 - How the Packet Filter Language Works 14-4
 - Basic Elements of a Packet Filter 14-6
 - Implementing Sequential Tests in a Packet Filter 14-8
 - Pre-processed and Run-time Storage 14-9
 - Procedure for Writing a Filter 14-10
 - Examples of Creating Filters 14-11
 - Filtering Problem 14-11
 - Packet Filter Solution 14-12
 - Tools for Writing a Filter 14-18
 - Using the Built-in Line Editor 14-18
 - Using an External Text Editor 14-20
- Deleting Packet Filters 14-20
- Editing Packet Filters 14-21
- Loading Packet Filters 14-22
- Copying Packet Filters 14-23
- Assigning Packet Filters to Ports 14-23
- Unassigning Packet Filters from Ports 14-24

15 CONFIGURING ADDRESS AND PORT GROUPS TO USE IN PACKET FILTERS

- Using Groups in Packet Filters 15-1
- Listing Groups 15-2
- Displaying Groups 15-3
- Creating New Groups 15-4
- Deleting Groups 15-6

Copying Groups 15-7
Adding Addresses and Ports to Groups 15-7
Removing Addresses or Ports from a Group 15-9
Loading Groups 15-11

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

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

INDEX

ABOUT THIS GUIDE

Introduction

The *LANplex 6000 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 6000 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 inter-connected LANs.



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



Throughout this guide, information that applies to the LANplex Management Module Plus (LMM+) also applies to the original LANplex Management Module (LMM).

How to Use This Guide

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

Table 1 Description of Guide Parts

Part	Contents
I: Getting Started	<ul style="list-style-type: none"> Introducing LANplex system administration Learning about the various system configurations and the quick commands to perform them Learning about password access to the Console Learning about the Administration Console menu structure and maneuvering within the Console (using commands and moving between menus) Getting help Setting interface parameters (screen height and control keys) Running scripts of Console tasks
II: System-level Functions	<ul style="list-style-type: none"> Setting up the system for management access (through serial ports or using IP and setting up SNMP) Configuring SNMP community strings Setting up trap reporting Configuring system parameters, such as name, date/time, and passwords Installing new system software Baselining statistics Saving, restoring, and resetting nonvolatile data
III: Ethernet, FDDI, and Token Ring	<ul style="list-style-type: none"> Displaying statistics for and labeling Ethernet ports Configuring the system's FDDI backplane paths for single station or multi-station Displaying statistics for and configuring various parameters for FDDI stations, ports, MACs, and paths Setting up the system to monitor Ethernet port activity Displaying statistics for and configuring various parameters for Token Ring ports

(continued)

Table 1 Description of Guide Parts (continued)

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

Conventions

The following tables list icon and text conventions that are used throughout this guide.

Table 2 Notice Icons

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

Table 3 Text Conventions

Convention	Description
"Enter"	"Enter" means type something, then press the [Return] or [Enter] key.
"Syntax" vs. "Command"	<p>"Syntax" indicates that the general command syntax form is provided. You must evaluate the syntax and supply the appropriate value; for example:</p> <p>Set the date by using the following syntax:</p> <pre>mm/DD/yy hh:mm:ss xm</pre> <p>"Command" indicates that all variables in the command syntax form have been supplied and you can enter the command as shown in text; for example:</p> <p>To update the system software, enter the following command:</p> <pre>system software Update</pre>
Text represented as screen display	This typeface represents text that appears on your terminal screen; for example: NetLogin:
Text represented as commands	This typeface represents commands that you enter; for example: bridge port stpState
<i>Italic</i>	<i>Italic</i> is used to denote emphasis and buttons.
Bold	Bold is used to denote key features, menus, and menu options.
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 6000 Documentation

The following documents comprise the LANplex 6000 documentation set. If you want to order a document that you do not have or order additional documents, contact your sales representative for assistance.

- *LANplex 6000 Unpacking Instructions*
Describe how to unpack your LANplex system. It also provides you with an inventory list of all the items shipped with your system. (Shipped with system)
- *LANplex 6000 Software Release Notes*
Provide 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 6000 Planning Your Site*

Provides information on the planning requirements you should consider when preparing your site for a LANplex 6000 system. (Shipped with system/Part No. 801-00251-000)
- *LANplex 6000 Getting Started*

Describes all the procedures necessary for installing, cabling, powering up, and troubleshooting your LANplex system. (Shipped with system/Part No. 801-00252-000)
- *LANplex 6000 Operation Guide*

Provides information to help you understand system management and administration, FDDI technology, and bridging. It also describes how these concepts are implemented in the LANplex system. (Shipped with system/Part No. 801-00253-000)
- *LANplex 6000 Administration Console User Guide*

Provides information about using the Administration Console to configure and manage your LANplex system. (Shipped with system/Part No. 801-00254-000)
- *LANplex 6000 Extended Switching User Guide*

Describes how the routing protocols are implemented in the LANplex system as well as providing information about using the Administration Console to configure and manage your routing protocols. (Shipped with the option package/Part No. 801-00257-000)
- *Command Quick Reference for the LANplex 6000 Administration Console*

Contains all of the Administration Console configurations for the LANplex system. (Shipped with system/Part No. 801-00258-000)
- *LANplex 6000 Control Panel User Guide*

Provides information about using the LANplex 6000 control panel to configure and manage your LANplex system. (Shipped with system/Part No. 801-00131-000)
- *Module Installation Guides*

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

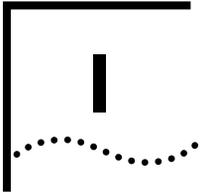
Documentation Comments

Your suggestions are very important to us and will help make LANplex documentation more useful to you. Please email comments about this guide to 3Com at: **sdtechpubs_comments@3Mail.3Com.com**

Please include the following information when commenting:

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

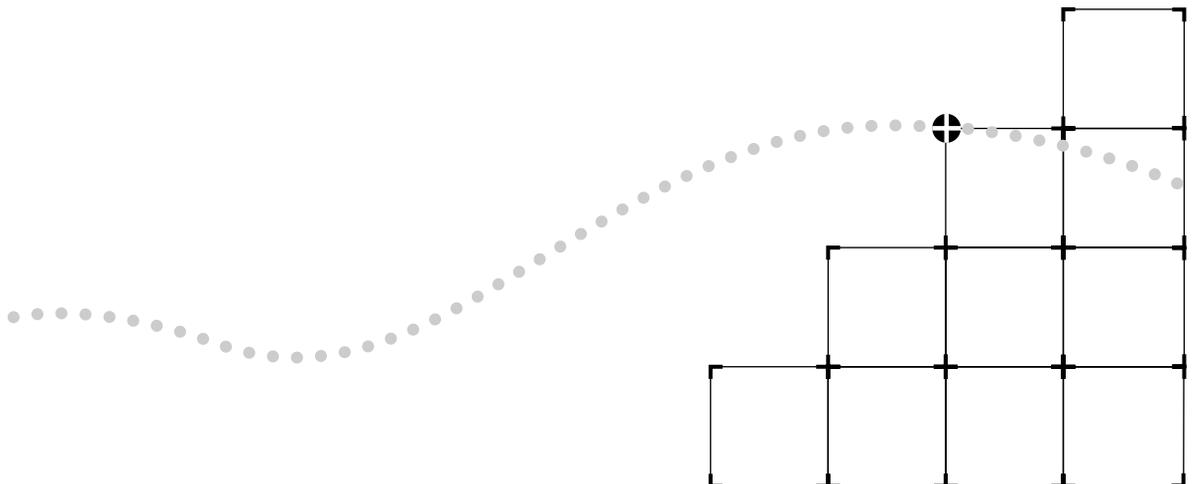
*Example: LANplex 6000 Planning Your Site
Part No. 801-00128-000
Page 2-5 (chapter 2, page 5)*



GETTING STARTED

Chapter 1 LANplex 6000 Administration Overview

Chapter 2 How to Use the Administration Console



1

LANPLEX 6000 ADMINISTRATION OVERVIEW

This chapter introduces you to LANplex 6000 administration and briefly describes the switching system parameters that you can configure.



For information on LANplex Extended Switching functionality refer to the LANplex 6000 Extended Switching User Guide.

About LANplex Administration

The LANplex system switching software is installed at the factory in flash memory on the LANplex Management Module Plus (LMM+). 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 may need to configure certain parameters for the system to operate effectively in your particular networking environment. Your LANplex system may contain several different module types, each of which provide certain types of functionality. A listing of the modules and their configurable parameters is provided below:

- LANplex Management Module Plus (LMM+) — System-level configuration parameters and management configuration parameters.
- FDDI Concentrator Module (FCM) — FDDI configuration parameters.
- Ethernet Switching Module (ESM) — Ethernet, bridging, and routing configuration parameters.
- Ethernet/FDDI Switching Module (EFSM) — Ethernet, FDDI, bridging, and routing configuration parameters.
- Token Ring Switching Module (TRSM) — Token Ring bridging configuration parameters.



Unless specifically referenced by name, the ESM, EFSM, and TRSM are referred to as "switching modules" in this guide.

Additionally, when managing your LANplex system, you may want to view important MAC, port, bridge, and IP statistics. The LANplex 6000 Administration Console allows you to configure your system and display these important statistics. For more complete network management, you can use an external application, such as 3Com's Transcend Enterprise Manager.

Configuration Options

You can configure the following parameters for your LANplex system from the Administration Console:

- System-level (Table 1-1)
- Management Access (Table 1-2)
- Ethernet, FDDI, and Token Ring (Table 1-3)
- Bridging (Table 1-4)

These tables provide a brief description of each parameter, along with the Administration Console command to access the configuration task quickly. They also tell you where to look in the documentation for additional information about configuring the specific parameter.

Table 1-1 System-level Configuration Options

Option	Description & Quick Access	For More Information, Go to...
Install software into LANplex flash memory	Update your LANplex system software. Software is initially installed at the factory. Quick Command: <code>system softwareUpdate</code>	Chapter 5: <i>Updating System Software</i> page 5-4
Set baseline for statistics	Baseline Ethernet, FDDI, and bridging statistics so you can more easily evaluate recent activity in your system and on your network. Quick Command: <code>system baseline set</code>	Chapter 6: <i>Baselining Statistics</i> page 6-2
Set Console password(s)	Prevent unauthorized users from accessing the Administration Console. You can set passwords for different levels of access (read, write, administer). Quick Command: <code>system password</code>	Chapter 4: <i>Administering Your System Environment</i> page 4-3

Table 1-1 System-level Configuration Options (continued)

Option	Description & Quick Access	For More Information, Go to...
(continued)		
Name the system	Assign a unique name to the system for management. For example, you may choose to name the system based on its location: <i>LANplex-Floor2</i> . Quick Command: <code>system name</code>	Chapter 4: <i>Administering Your System Environment</i> page 4-5
Set system date and time	Ensure that messages are accurately logged. The system's internal clock is preset at the factory. Quick Command: <code>system time</code>	Chapter 4: <i>Administering the System Environment</i> page 4-5
Lock the Console for your session	Prohibit remote access to the Console. Quick Command: <code>system consoleLock</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-33
Lock the control panel to prevent unauthorized use	Prevent write access of LANplex system parameters from the control panel. Quick Command: <code>system panelLock</code>	Chapter 2: <i>How to Use the Administration Console</i> page 2-19
Set screen height	Adjust the console screen height for your particular terminal. Quick Command: <code>system screenHeight</code>	Chapter 2: <i>How to Use the Administration Console</i> page 2-17
Enable the [Control] keys when working in the Console	Enable quick keys for the reboot (CTL-X) and abort (CTL-C) functions. Quick Command: <code>system ctlKeys</code>	Chapter 2: <i>How to Use the Administration Console</i> page 2-18
Save, restore, or reset nonvolatile data in the system	Provide backup for nonvolatile data, restore nonvolatile data to the system, or reset nonvolatile data to the factory defaults. Quick Command: <code>system nvData</code>	Chapter 7: <i>Saving, Restoring, & Resetting Nonvolatile Data</i> page 7-2, page 7-3, and page 7-7

Table 1-2 Management Access Configuration Options

Option	Description & Quick Access	For More Information, Go to...
Configure terminal serial port baud rate	View messages during power up diagnostics and access the Administration Console using the terminal serial port. The factory defaults for the baud rate allow you to connect a VT or tty type terminal or terminal emulator to the serial port with a null modem cable. Quick Command: <code>system serialPort terminalSpeed</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-3 <i>LANplex 6000 Control Panel User Guide</i>
Set up for external modem	Manage your system remotely with an external modem. You can set the modem serial port baud rate and configure an external modem. Quick Commands: <code>system serialPort modemSpeed</code> <code>system serialPort connectModem</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-3 <i>LANplex 6000 Control Panel User Guide</i>
Configure IP interfaces	Establish the relationship between module ports and the subnets in your IP network. Allows you to communicate with the system using SNMP, rlogin, or telnet. Quick Command: <code>ip interface define</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-9 <i>LANplex 6000 Control Panel User Guide</i>
Configure routes to other IP networks, subnets, and hosts	Add static routes to the routing table, remove routes, or define the default route (the route used when no other routes match). Quick Command: <code>ip route</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-13
Maintain the ARP cache	Remove and flush entries from the ARP cache when necessary. Quick Command: <code>ip arp</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-17
Set RIP's operational mode	Process RIP messages in one of two ways: Off (ignores RIP messages and does not generate any) or Passive (responds to RIP but does not broadcast its own), and Active (responds and broadcasts RIP messages) Quick Command: <code>ip rip</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-18
Ping an IP station or the LANplex system	Find out if the system can reach an IP station or ensure that the system is on the network. (You must first assign the system an IP address by configuring an IP interface.) Quick Command: <code>ip ping</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-19

(continued)

Table 1-2 Management Access Configuration Options (continued)

Option	Description & Quick Access	For More Information, Go to...
Set the SNMP mode	Specify whether the system is managed with a single agent or with multiple agents — one for each logical device (chassis, bridge, router) in the system. Quick Command: <code>snmp mode</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-22
Configure SNMP community strings	Specify the type of authorization for SNMP management (read-only or read-write). Quick Command: <code>snmp community</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-25
Configure SNMP trap reporting	Specify which SNMP traps from which IP addresses are reported to an SNMP-based external network management application. Quick Command: <code>snmp trap</code>	Chapter 3: <i>Configuring Management Access to the System</i> page 3-26

Table 1-3 Ethernet, FDDI, & Token Ring Configuration Options

Option	Description & Quick Access	For More Information, Go to...
Label an Ethernet port	Give an Ethernet port a unique name. Useful for port identification when managing the system. Quick Command: <code>ethernet label</code>	Chapter 8: <i>Administering Ethernet Ports</i> page 8-7
Set the Ethernet port state	Enable or disable an Ethernet port to control whether the port sends and receives frames. Quick Command: <code>ethernet portState</code>	Chapter 8: <i>Administering Ethernet Ports</i> page 8-8
Set the FDDI backplane ring mode	Set the FDDI backplane rings for single station mode or multi-station mode. Perform prior to other FDDI setup. Quick Command: <code>fddi backplane</code>	Chapter 9: <i>Administering FDDI Resources</i> page 9-3
Set FDDI station parameters	Set parameters for connection policies, the neighbor notification timer, and status reporting. Quick Commands: <code>fddi station connectPolicy</code> <code>fddi station tNotify</code> <code>fddi station statusReporting</code>	Chapter 9: <i>Administering FDDI Resources</i> page 9-4, page 9-5

(continued)

Table 1-3 Ethernet, FDDI, & Token Ring Configuration Options (continued)

Option	Description & Quick Access	For More Information, Go to...
Set FDDI path parameters	<p>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.</p> <p>Quick Commands: <code>fddi path tvxLowerBound</code> <code>fddi path tmaxLowerBound</code> <code>fddi path maxTreq</code></p>	<p>Chapter 9: <i>Administering FDDI Resources</i></p> <p>page 9-8, page 9-9</p>
Set FDDI MAC parameters	<p>Set the parameters for the frame error threshold and the not copied threshold, and enable/disable LLC service.</p> <p>Quick Commands: <code>fddi mac frameErrorThreshold</code> <code>fddi mac notCopiedThreshold</code> <code>fddi mac llcService</code> <code>fddi mac path</code></p>	<p>Chapter 9: <i>Administering FDDI Resources</i></p> <p>page 9-16, page 9-17, page 9-18</p>
Assign FDDI MACs to stations and paths	<p>After setting the backplane paths mode, select a station for each MAC in your system and assign the MACs to paths. Station and path selections vary depending on the backplane path mode.</p> <p>Quick Command: <code>fddi mac station</code> <code>fddi mac path</code></p>	<p>Chapter 9: <i>Administering FDDI Resources</i></p> <p>page 9-25, page 9-26</p>
Set FDDI port parameters	<p>Set the parameters for link error rate alarm threshold and link error rate cut-off threshold. Assign the A and B ports to either the primary or secondary rings.</p> <p>Quick Commands: <code>fddi port lerAlarm</code> <code>fddi port lerCutoff</code> <code>fddi port path</code></p>	<p>Chapter 9: <i>Administering FDDI Resources</i></p> <p>page 9-21, page 9-22, page 9-23</p>
Label an FDDI port	<p>Give an FDDI port a unique name. Useful for port identification when managing the system.</p> <p>Quick Command: <code>fddi port label</code></p>	<p>Chapter 9: <i>Administering FDDI Resources</i></p> <p>page 9-22</p>
Assign FDDI ports to stations and paths	<p>After setting the backplane paths mode, select a station for each FDDI port in your system and assign the ports to paths. Station and path selections vary depending on the backplane path mode.</p> <p>Quick Commands: <code>fddi port station</code> <code>fddi port path</code></p>	<p>Chapter 9: <i>Administering FDDI Resources</i></p> <p>page 9-32, page 9-33</p>

(continued)

Table 1-3 Ethernet, FDDI, & Token Ring Configuration Options (continued)

Option	Description & Quick Access	For More Information, Go to...
Label a Token Ring port	Give a Token Ring port a unique name. Useful for port identification when managing the system. Quick Command: <code>tokenring label</code>	Chapter 10: <i>Administering Token Ring Ports</i> page 10-8
Configure Token Ring ports	Specify the Token Ring port state (enabled or disabled), the port speed (4Mbps, 16Mbps, or 16MbpsEarly Token Release), and the port mode (station or lobe) Quick Command: <code>tokenring portState</code> <code>tokenring portSpeed</code> <code>tokenring portMode</code>	Chapter 10: <i>Administering Token Ring Ports</i> page 10-8, page 10-9, page 10-9
Configure Ethernet ports to be monitored by a network analyzer	Configure the system for Roving Analysis to be able to analyze data forwarded through a LANplex Ethernet port. 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). Quick Commands: <code>analyzer add</code> <code>analyzer start</code>	Chapter 11: <i>Setting up the System for Roving Analysis</i> page 11-5, page 11-7

Table 1-4 Bridging Configuration Options

Option	Description & Quick Access	For More Information, Go to...
Set the bridging mode	Configure the system to operate in either IEEE 802.1d bridging mode. Quick Command: <code>bridge mode</code>	Chapter 12: <i>Administering the Bridge</i> page 12-6
Set the bridge address reporting threshold	Generates the SNMP trap <code>addressThresholdEvent</code> when the threshold is reached. Quick Command: <code>bridge addressThreshold</code>	Chapter 12: <i>Administering the Bridge</i> page 12-9
Set the bridge address aging timer	Allows the bridging ports to age dynamically-learned addresses in a timely manner and thereby prevent needless packet flooding. Quick Command: <code>bridge agingTime</code>	Chapter 12: <i>Administering the Bridge</i> page 12-13

Table 1-4 Bridging Configuration Options (continued)

Option	Description & Quick Access	For More Information, Go to...
Configure Spanning Tree Protocol (STP) parameters for a bridge or bridge port	Block redundant routes, creating a loopless network that operates as if only one link connects each LAN. Quick Commands: <code>bridge</code> <code>bridge port</code>	Chapter 12: <i>Administering the Bridge</i> page 12-12 Chapter 13: <i>Administering Bridge Ports</i> page 13-9
Statically configure bridge port addresses	Assign new MAC addresses to selected bridge ports. Can be used as a means of network security because a statically-configured address is not aged. Quick Command: <code>bridge port address add</code>	Chapter 13: <i>Administering Bridge Ports</i> page 13-13
Create packet filters	Restrict which packets are forwarded through the system. Packet filtering is provided for various packet processing paths on each bridge port of the system. When packet filtering is configured for a port, the packets are filtered according to a user-defined packet filter definition. Quick Command: <code>bridge packetFilter</code>	Chapter 14: <i>Creating and Using Packet Filters</i> page 14-2
Create address and port groups to use as filtering criteria	Specify groups of addresses or ports that you want to use in a packet filter definition. The opcodes <i>pushSAGM</i> , <i>pushDAGM</i> , <i>pushSPGM</i> , and <i>pushDPGM</i> identify both source and destination address and port group masks within the filter definition. Quick Commands: <code>bridge packetFilter addressGroup</code> <code>bridge packetFilter portGroup</code>	Chapter 15: <i>Configuring Address and Port Groups to Use in Packet Filters</i> page 15-4

2

HOW TO USE THE ADMINISTRATION CONSOLE

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

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

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. These levels are described in Table 2-1.

Table 2-1 Password Access Levels

Access Level	Kind of User	Allows Users to...
Administer	Those who need to perform system set-up and management (usually a single network administrator)	Perform system-level administration (such as setting passwords, loading new software, etc.)
Write	Those who need to perform active network management	Configure network parameters (such as setting the aging time for a bridge)
Read	Those who only need to view system parameters	Access “display” menu items (display, summary, detail) only

Access level prompt

Each time you access the Administration Console, you are prompted for an access level and password, as shown below:

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

The passwords are stored in nonvolatile memory. You must enter the password correctly before you are allowed to continue. The first time you access the Console, the password is null.

Initial user access As the initial user, access the system at the *administer* level and press return at the password prompt.

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

Administer Access Example If you have administer access, each menu contains all options, as shown in the **system** menu example below:

```
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
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
panelLock       - Allow/Disallow control panel write access
ctlKeys         - Enable/Disable Ctl-X (reboot) and Ctl-C (abort)
nvData          - Save, restore, or reset nonvolatile data
reboot          - Reboot the LANplex system
```

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

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

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

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

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

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

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

```

Menu options: -----
      display                - Display the system configuration
      baseline                - Administer 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 Console by selecting options from this menu and from others below it. Each menu option is accompanied by a brief description. See the top-level menu below:

```

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

Options
(These vary per
level of access.)

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

```

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.

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, you would enter **system** at the top-level menu, **serialPort** at the system menu, then **connectModem** at the serialPort menu.

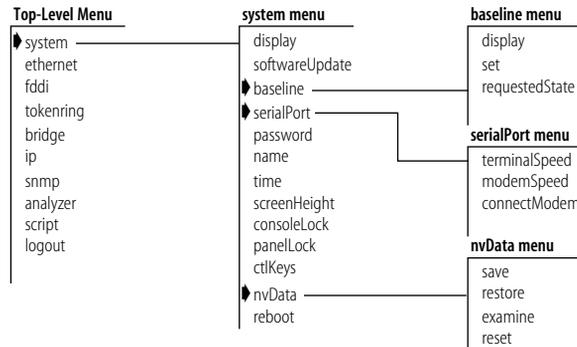


Figure 2-1 System-level Functions Menu Hierarchy

Ethernet Menu

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

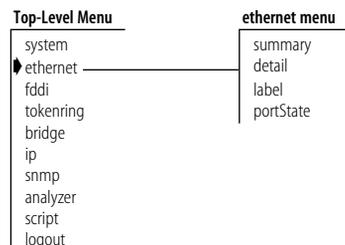


Figure 2-2 Ethernet Port Menu Hierarchy

FDDI Menu

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

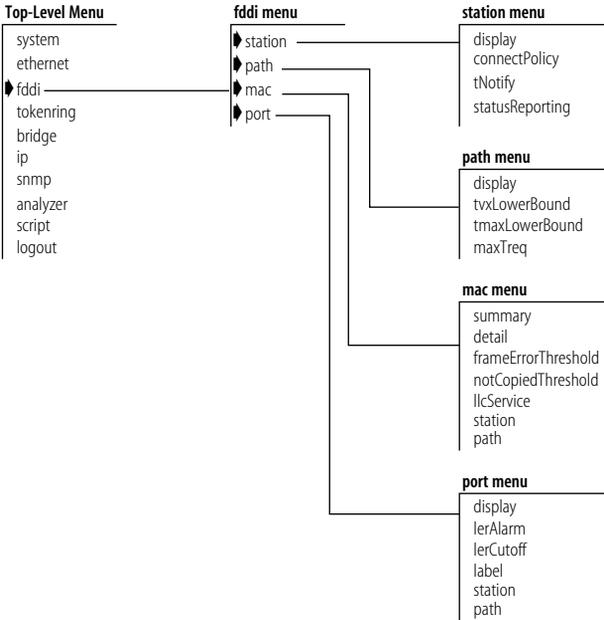


Figure 2-3 FDDI Resources Menu Hierarchy

Token Ring Menu

From the **tokenring** menu, you can view information for Token Ring ports and name Token Ring ports (see Figure 2-4). You can also enable or disable Token Ring ports, configure the Token Ring port speed, and set the Token Ring port mode. For example, to set the Token Ring port state, you would enter **tokenring** at the top-level menu, then **portState** at the tokenring menu.

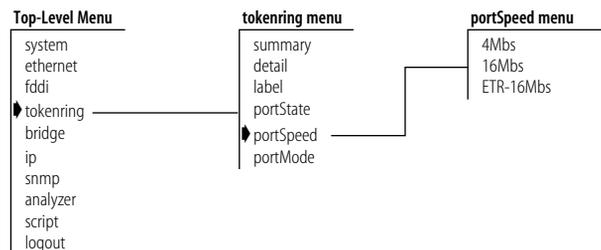


Figure 2-4 Token Ring Port Hierarchy

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, if you wanted to set the Spanning Tree state for a bridge port, you would enter **bridge** at the top-level menu, **port** at the bridge menu, and **stpState** at the bridge menu.

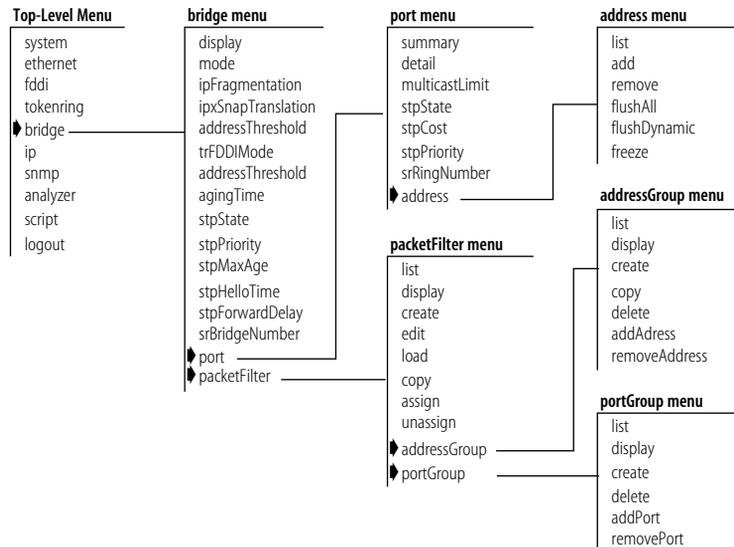


Figure 2-5 Bridging Menu Hierarchy

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), the Routing Information Protocol (RIP), and ping IP stations (see Figure 2-6). For example, to define a new IP interface, you would enter **ip** at the top-level menu, **interface** at the ip menu, then **define** at the interface menu.

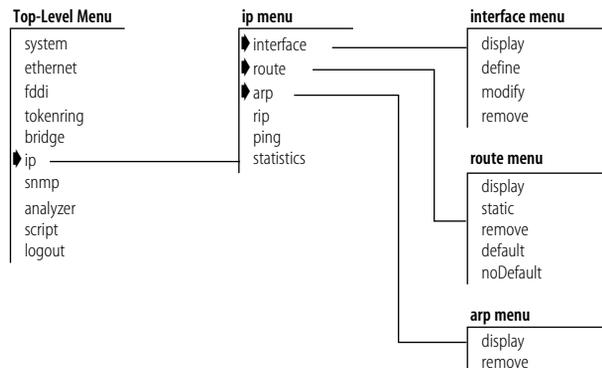


Figure 2-6 IP Menu Hierarchy

SNMP Menu

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

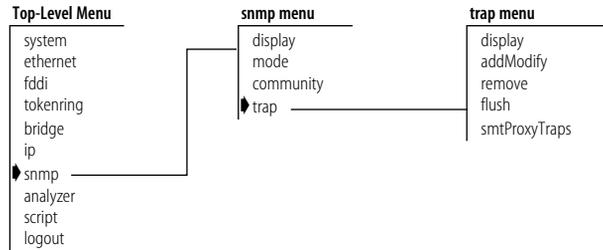


Figure 2-7 SNMP Menu Hierarchy

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, you would enter **analyzer** at the top-level menu, then **add** at the analyzer menu.

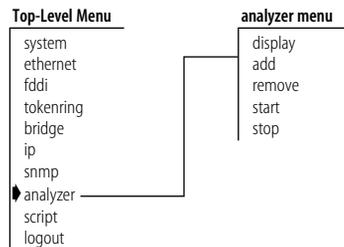


Figure 2-8 Analyzer Menu Hierarchy

Selecting Menu Options

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

When a new menu appears, the selection prompt 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 to the following:

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

Entering a Command String

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

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

The most abbreviated version of the same command string would be:

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

When you enter a command string, you move to the last item entered in the command string, and information relevant to that command is displayed, whether it is a menu, prompt, or screen display.

If a command is entered 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 it became incorrect.

Entering Values

When you reach the level where 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 may be the system default or the current user-defined value of that parameter.

The valid values are displayed in parenthesis and the default value is in brackets. In the example below, disabled and enabled are the valid values and 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, you enter:

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

Getting Out

To go to the previous menu or cancel an operation that you are currently performing, enter **q**, followed by [Return].

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

To completely leave the Administration Console, see the section "Exiting the Administration Console" on page 2-23 for instructions.

Console Interface Parameters

You can change three Console interface parameters: screen height, console disconnections, and the reboot and abort control keys.

Adjusting the Screen Height

Most terminal screens have a height of 24 lines. You can change the Administration Console's screen height to increase or decrease the space available for displaying information. You can configure the screen height to be between 20 to 200 lines or zero (0) for infinite; the default is 24.

Each time the screen output reaches the 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 may scroll beyond the screen, depending on your screen size.



The screen height setting does not affect the way the system displays menus. The screen height setting controls the way the system displays information you request using the menus, such as statistical summaries.

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 (20 to 200). To receive no prompts, set the screen height to infinite (0).

See the example below:

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

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

See the example below:

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

Top-Level Menu

```

system
ethernet
fddi
tokenring
bridge
ip
snmp
analyzer
script
logout
display
softwareUpdate
baseline
serialPort
password
name
time
screenHeight
consoleLock
panelLock
ctlKeys
nvData
reboot

```

Preventing Console Disconnections

Because only a single shell is supported by the Administration Console, you may be disconnected from your session if someone else remotely accesses the 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 Console. Remote access is prohibited only for that particular session.



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

To lock the Console:

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

```
system consoleLock
```

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

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

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

Top-Level Menu

```

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

```

Configuring Control Panel Write Access

You can permit or prohibit the modification of system configurations and MIB parameters from the LANplex system's control panel. See the *LANplex 6000 Control Panel User Guide*, for more information about the parameters you can set from the control panel. As shipped, modifications from the control panel are permitted.

To change the control panel setting:

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

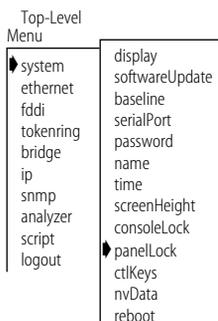
```
system panelLock
```

The current setting is displayed (locked or unlocked).

- 2 Enter **y** to configure the opposite of the current setting. Enter **n** to maintain the current access setting.

See the following example:

```
The Control Panel is currently unlocked.
Do you want to lock the Control Panel? (y/n): y
```



Disabling Reboot and Abort Keys

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



CAUTION: *If you disable the control keys, only use Ctl-C if instructed to by a Technical Support representative. Using Ctl-C may irregularly terminate a Console session.*

To enable or disable the reboot and abort control keys:

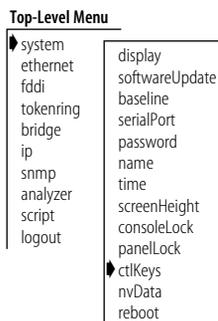
- 1 From the top level of the Administration Console, enter:

```
system ctlKeys
```

You are prompted for whether you want to enable or disable the functionality as shown below:

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

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



Running Scripts of Console Tasks

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

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

To run a script:

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

script

You are prompted for information about where the script you want to run is stored: 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 Console.

The example below shows how you can script the following 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

Top-Level Menu

```
system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout
```

```
# 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
all                 # ports
#
ip interface display
#
#
# Validate access to management workstation
#
ip ping
158.101.112.26      # management workstation address
#
# Enable the Spanning Tree Protocol
#
bridge stpState enabled
#
# Configure my node as an SNMP trap destination
#
snmp trap add
158.101.112.26      # management workstation address
all                 # turn on all traps
q                   # no more trap destinations
#
snmp trap display
#
```

Getting Help in the Console

If you need assistance when using the Administration Console, the Console has online help and an outlining feature, both of which can be accessed from anywhere in the Console. These features are described below.

Online Help

The Administration Console online help provides an overview of the Console and lets you access information about any menu option in the Console from anywhere in the Console.

General online help

To get help using the Administration Console, enter `?`, followed by [Return]. General instructions for using the Administration Console are displayed.

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.

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

```
? ethernet
```

Instructions for using that option are displayed, if available.

Outlining

The outlining feature allows you to list the menu options that fall in the hierarchy below the current menu. The default displays three levels of options (if available).

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

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

```
outline 2
```

Exiting the Administration Console

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

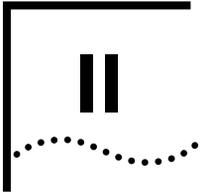
To exit from the Administration Console:

- 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

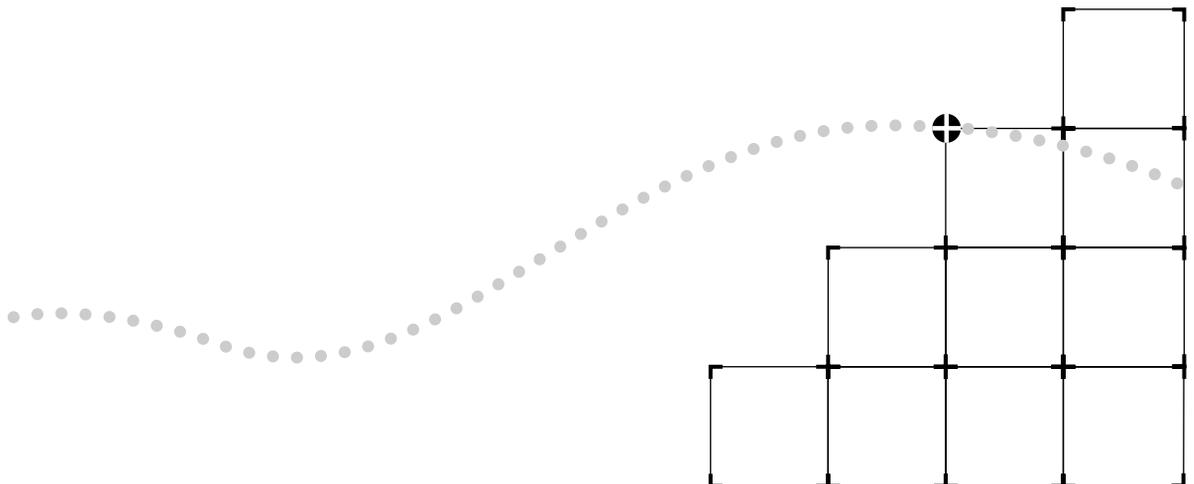
Top-Level Menu

```
system
ethernet
fddi
tokenring
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** Updating System Software
- Chapter 6** Baselineing Statistics
- Chapter 7** Saving, Restoring, & Resetting Nonvolatile Data



3

CONFIGURING MANAGEMENT ACCESS TO THE SYSTEM

This chapter describes how to configure management access to the LANplex system through a serial connection or an IP interface and describes how to configure the LANplex system to be managed using the Simple Network Management Protocol (SNMP). It also explains how to prevent console disconnections while you are using the Console and how to prohibit control panel write access.

About Management Access

You can access the Administration Console directly through the LMM+ *terminal* serial port or through the LMM+ *modem* serial port. Alternatively, from a PC or workstation, you can access the Administration Console through the LMM+ Ethernet port or the first LMM+ FDDI MAC. Once you establish an IP interface, you can also set up the system to be managed by an SNMP-based network management application, such as 3Com's Transcend Enterprise Manager.

Using a Serial Connection

Direct access through the terminal serial port (serial port 1) is often preferred because it allows you to stay attached during system reboots. You can also access the Administration Console through an external modem attached to the modem serial port (serial port 2).



See the LANplex Management Module Installation Guide for serial port pin-outs.

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

Using an IP Interface

Through the LMM+'s Ethernet port, you can rlogin or telnet to the Administration Console. A unique IP address on a separate IP network from the in-band network is used to identify this interface. This connection is faster than the serial port, and it allows you to connect to the Administration Console using TCP/IP from a host.

You can also rlogin or telnet to the Administration Console using the first LMM+ FDDI MAC. The SNMP agent(s) can be accessed through this interface as well. A unique IP address is used to identify this interface. By default, the first LMM+ FDDI MAC (FDDI MAC 1) is placed on the primary FDDI path. Therefore, it is possible for stations connected to either the ESM Ethernet ports, EFSM Ethernet/FDDI ports, or FCM FDDI ports to access the Administration Console through the FDDI MAC 1.

In-band or Out-of-band?

If you manage your LANplex system and its attached LANs over the same network that carries your regular data traffic, then you are managing your network *in-band*. This is often the most convenient and inexpensive way to access your LANplex system. The disadvantage, however, of using in-band management is that if your data network is faulty, you may not be able to diagnose the problem because the management requests are sent over the same network. If you are using a dedicated network for management data, then you are managing your network *out-of-band*.



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

Setting up the Terminal Serial Port

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



CAUTION: *Baud rate changes take effect immediately. You will be unable to communicate using the serial port until you adjust the baud rate of your terminal or terminal emulator appropriately.*



The cable attached to the terminal serial port (serial port 1) must be a 6-conductor RJ-12 cable (which has the RJ-12 to DB-25 adapter crossed internally). This cable is shipped with the LANplex system.

To set the baud rate for the terminal 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 terminal serial port when you set the baud rate for that serial port, the following message is displayed:

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

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

Top-Level Menu

```

system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout
  display
  softwareUpdate
  baseline
  serialPort
  password
  name
  time
  screenHeight
  consoleLock
  panelLock
  ctrlKeys
  nvData
  reboot
    terminalSpeed
    modemSpeed
    connectModem

```

Setting up the Modem Serial Port

For the modem serial port, you can set the port speed to match your external modem baud rate, 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 may need to change the baud rate to match your external modem.

To set the baud rate for the modem serial port:

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

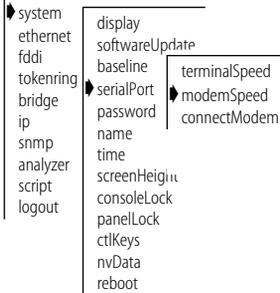
```
system serialPort modemSpeed
```

- 2 Enter the baud rate for the 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 setting 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.



The cable attached to the modem serial port (serial port 2) must be a 6-conductor RJ-12 cable (which has the RJ-12 to DB-25 adapter crossed internally). This cable is shipped with the LANplex system.

To configure the modem port:

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

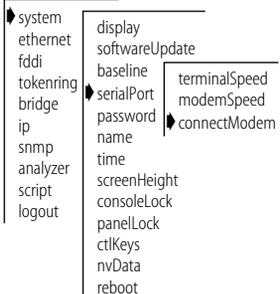
```
system serialPort connectModem
```

You can now issue the commands to the attached modem 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 Once the modem is configured, enter the escape sequence `~1` with no intervening characters.

When you enter the escape sequence, the connection to the modem serial port is broken and you are returned to the previous menu.

Top-Level Menu



Setting up an IP Interface for Management

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

Each switching module operates as a separate IP router. This means that each module has its own interfaces, routing table, ARP cache, etcetera.

General Setup Process

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



You should assign an IP host address to all ports for system management.

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

- Routes
- Address Resolution Protocol (ARP) cache
- Routing Information Protocol (RIP)

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 switching modules and the subnets in your IP network. You must define one interface for each group of ports that are connected to the same subnet. This means that every switching module has one interface defined for each subnet to which it is directly connected.

An IP interface has the following information associated with it:

- **IP Address**

This is the address specific to your network. It should be chosen from the range of addresses assigned to your organization. An interface's IP address serves two functions. First, it is the address that is used when sending IP packets to or from the switching module itself. Second, the IP address defines the network and subnet numbers of the segments connected to that interface.



Packets to be forwarded by the switching module contain the IP addresses of the original source and the ultimate destination.

- **Subnet Mask**

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

- **Broadcast Address**

This is the IP address to be used by the switching module when it broadcasts packets to other stations on the same subnet. In particular, this address is used for sending RIP updates. By default, the switching module uses a directed broadcast (all ones in the host field).

- **Cost**

This is a number between one and fifteen that is used when calculating route metrics. Unless your network has special requirements, you should assign a cost of **1** to all interfaces.

- **Ports**

A single interface may contain several bridge ports. All of the ports corresponding to one interface share the same IP address, subnet mask, broadcast address, and cost. An ESM contains nine ports: one FDDI and nine Ethernet. The port indices are always the following: 1 = FDDI and 2 – 9 = Ethernet. An EFSM contains a maximum of eighteen ports: two FDDI and sixteen Ethernet. The port indices for the maximum configuration are the following: 1, 2 = FDDI and 3 – 18 = Ethernet.

Ensure that the port to which your management station is attached is included in an interface. IP packets will not be forwarded to ports that are not assigned to an IP interface.

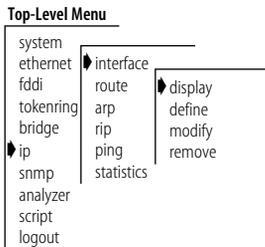


You should assign an IP host address to all ports for system management.

Displaying Interfaces

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

To display IP interface information:



- 1 From the Administration Console top-level menu, enter:
ip interface display
- 2 Enter the slot(s) of the switching module(s) for which you want to display the interface information. Separate non-consecutive slots with commas (.). Enter a consecutive series of slots using a dash (-).

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

Slot 3 - IP forwarding is enabled, RIP is passive.

Index	IP address	Subnet mask	Cost	Ports (1-2=FDDI, 3-18=Ethernet)
1	158.101.112.225	255.255.255.0	1	3

Defining an Interface

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

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

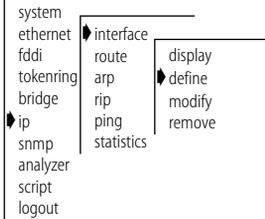
Table 3-1 Recommended Settings for IP Management Access

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

If you do not assign all ports to this interface, ensure that you include the port to which your network management station is attached.

To define an IP interface:

Top-Level Menu



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

```
ip interface define
```

- 2 Enter the slot of the switching module for which you want to define an interface.

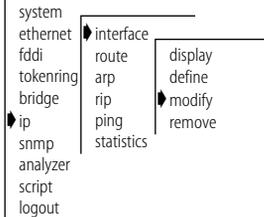
You are prompted for the interface's parameters. To use the value in brackets, press [Return] at the prompt.

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

See the example below:

```

Select slot {3-4} [3-4]: 3
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter broadcast address [158.101.1.255]:
Enter cost [1]:
Enter ports (1=FDDI, 2-9=Ethernet) (1-9|all): 2-4,8
  
```

Top-Level Menu**Modifying an Interface.**

To modify an IP interface that you have already defined:

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

```
ip interface modify
```

- 2 Enter the slot of the switching module for which you want to modify an interface.

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

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

Removing an Interface

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

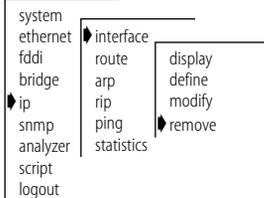
To remove an IP interface definition:

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

```
ip interface remove
```

- 2 Enter the slot of the switching module from which you want to remove an interface.

- 3 Enter the index number(s) of the interface(s) you want to remove.

Top-Level Menu**Administering Routes**

Each switching module 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 switching modules to use RIP to exchange routing information automatically.

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 corresponding to the bits set in the route subnet mask match the route destination address. When forwarding a packet, if the switching module 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. The switching module includes the metric in its RIP updates to allow other 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 switching module 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

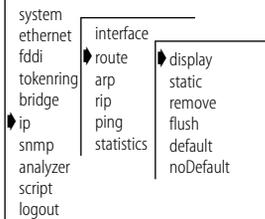
In addition to the routes to specific destinations, the routing table may contain an additional entry called the *default route*. The switching module uses the default route to forward packets that do not match any other routing table entry. You may want to use a default route in place of routes to numerous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

Displaying the Routing Table

You can display the routing tables for the switching modules in a system to determine which routes are configured and if they are operational.

To display the contents of the routing table:

Top-Level Menu



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

```
ip route display
```

- 2 Enter the slot(s) of the switching module(s) for which you want to display the routing table. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

In the following example, routes for an ESM in slot 3 are displayed. The configuration of IP forwarding and RIP is indicated in the display. The default route is displayed as "Default Route".

Slot 3 - IP forwarding is enabled, RIP is passive.

Destination	Subnet mask	Metric	Gateway	Status
Default Route	--	2	158.101.112.250	Learned (RIP)
10.0.0.0	255.0.0.0	8	158.101.112.254	Learned (RIP)
129.213.0.0	255.255.0.0	7	158.101.112.254	Learned (RIP)
137.39.0.0	255.255.0.0	2	158.101.112.250	Learned (RIP)
139.87.0.0	255.255.0.0	4	158.101.112.254	Learned (RIP)

Defining a Static Route

You may want to define a static route to transmit system traffic, such as system pings or SNMP response, through a consistent route. Prior to defining static routes on a given switching module, you must define at least one IP interface (see the section "Defining an Interface" on page 3-9). 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.

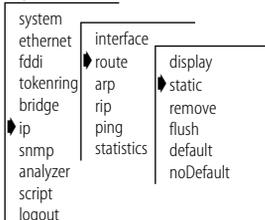
To define a static route:

- 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.

Top-Level Menu



- 2 Enter the slot of the switching module for which you want to define a static route.
- 3 Enter the destination IP address of the route.
- 4 Enter the subnet mask of the route.
- 5 Enter the gateway IP address of the route.

A static route is defined in the following example:

```
Select slot {3-4} [3-4]: 3
Enter destination IP address: 158.101.4.0
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter gateway IP address: 158.101.2.8
```

Removing a Route

To remove a route:

- 1 From the top level of the Administration Console, enter:
ip route remove
- 2 Enter the slot of the switching module for which you want to remove a static route.
- 3 Enter the destination IP address of the route.
- 4 Enter the subnet mask of the route.

The route is immediately deleted from the routing table.

Flushing a Route

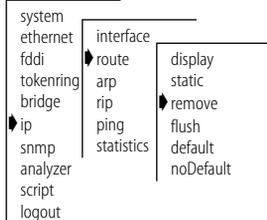
Flushing deletes all learned routes from the routing table.

To flush all learned routes

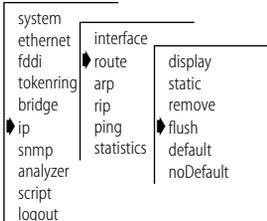
- 1 From the top level of the Administration Console, enter:
ip route flush
- 2 Enter the slot of the switching module for which you want to delete the learned routes.

All learned routes are immediately deleted from the routing table.

Top-Level Menu



Top-Level Menu



Setting the Default Route

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

If a switching module'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 this occurs, then the module 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 slot of the switching module for which you want to set a default route.
- 3 Enter the gateway IP address of the route.

The default route is immediately added to the routing table.

Removing the Default Route

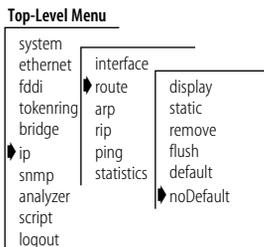
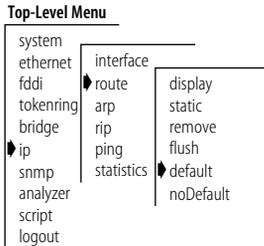
To remove a default route:

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

```
ip route noDefault
```

- 2 Enter the slot of the switching module for which you want to remove the default route.

The default route is immediately removed from the routing table.



Administering the ARP Cache

The switching modules use the Address Resolution Protocol (ARP) to find the MAC addresses corresponding to the IP addresses of hosts and other routers on the same subnets. Each device participating in routing maintains an ARP cache — a table of known IP addresses and their corresponding MAC addresses.

Displaying the ARP Cache

To display the contents of the ARP cache:

- 1 From the Administration Console top-level menu, enter:
ip arp display
- 2 Enter the slot(s) of the switching module(s) for which you want to display the ARP cache. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

The contents of the ARP cache are displayed as shown in the example below.

Slot 3- IP forwarding is enabled, 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

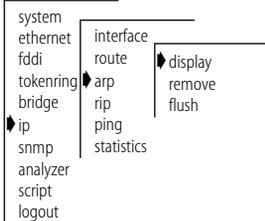
You may want to remove an entry from the ARP cache if the MAC address has changed.

To remove an entry from the ARP cache:

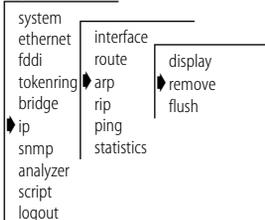
- 1 From the top level of the Administration Console, enter:
ip arp remove
- 2 Enter the slot of the switching module for which you want to remove an ARP cache entry.
- 3 Enter the IP address you want to remove.

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

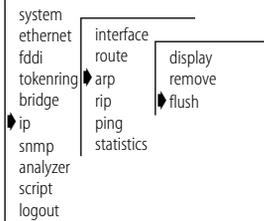
Top-Level Menu



Top-Level Menu



Top-Level Menu



Flushing an ARP Cache Entry

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

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

```
ip arp flush
```

- 2 Enter the slot of the switching module for which you want to remove an ARP cache entry.

The ARP cache entries are immediately removed from the table.

Setting the RIP Mode

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

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

RIP default mode

By default, RIP operates in passive mode.

To set the RIP operating mode:

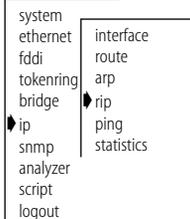
- 1 From the top level of the Administration Console, enter:

```
ip rip
```

- 2 Enter the slot(s) of the switching module(s) for which you want to set the RIP mode. Separate non-consecutive slots with commas (.). Enter a consecutive series of slots using a dash (-).

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

Top-Level Menu



See the following example:

```
Select slot(s) (3-4|all) [3-4]: all
Slot 3 - Select RIP mode (off, passive) [passive]: passive
Slot 4 - Select RIP mode (off, passive) [passive]: passive
```

Pinging an IP Station

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

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

- Alive
- No answer
- Network is unreachable

A network is unreachable when there is no route to that network.

To ping an IP station:

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

```
ip ping
```

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

```
IP Address: 192.9.200.40
```

You may receive one of the following responses:

```
192.9.200.40 is alive
```

OR

```
no answer from 192.9.200.40
```

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

```
Network is unreachable
```

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

Top-Level Menu

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

To display IP statistics:

- 1 From the Administration Console top-level menu, enter:
ip statistics
- 2 Enter the slot(s) of the switching module(s) for which you want to view IP routing statistics. Separate non-consecutive slots with commas (.). Enter a consecutive series of slots using a dash (-).

Statistics are displayed, as shown in the example below:

Slot 3 - IP forwarding is enabled,

```

inReceives    forwDatagramsn    inDelivers    outrequest
      51213                49743                3227                2285
                                outNoRoutes    inHdrErrors    inAddrErrors
                                     273                7                0

```

Slot 4 - IP forwarding is enabled, RIP is active.

```

inReceives    forwDatagrams    inDelivers    outRequests
      11                11                11                20
                                outNoRoutes    inHdrErrors    inAddrErrors
                                0                0                0

```

Top-Level Menu

```

system
ethernet
fdci
tokenring
bridge
ip
snmp
analyzer
script
logout

```

```

interface
route
arp
rip
ping
statistics

```

Setting Up SNMP on Your System

To manage the LANplex 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 (an SNMP manager) sends requests to the system where they are processed by the LANplex SNMP agent (or multiple agents, if configured).

The SNMP agent provides access to the collection of information about the LANplex system. Different views of MIB information are available depending on the LANplex SNMP management method you choose. Additionally, a LANplex 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* of the *LANplex 6000 Operation Guide*.

Configuring the SNMP Mode

The LANplex system supports two SNMP modes: **single agent** and **multiple agent**. These modes provide three methods of management:

- Method I – system management with a single SNMP agent
- Method II – direct management of individual devices in a system with multiple agents
- Method III – indirect management of individual devices in a system with multiple agents

For more information about each management method, along with a list of the advantages and disadvantages of using that particular method, see Chapter 3: *Management Access: Protocols* of the *LANplex 6000 Operation Guide*.

Configuration Guidelines

The configuration guidelines for each method of management are described below:

- Method I* To configure SNMP for system management with a single SNMP agent you must:
- Assign an IP address to at least one of the LMM+ FDDI MACs or the LMM+ Ethernet port.
 - Set the SNMP mode to singleAgent.
 - Set the destination IP address(es) of the management station(s) where traps should be forwarded by the system agent.
- Method II* To configure SNMP for direct management with multiple agents you must:
- Assign at least one IP address for the LMM+ and for each switching module (ESM, EFSM, and TRSM).
 - Set the SNMP mode to multipleAgent.
 - Set the destination IP address(es) of the management station(s) where traps should be forwarded by each agent.
- Method III* To configure SNMP for indirect management with multiple agents you must:
- Assign a single IP address for the LMM+ or for one of the ESMs, EFSMs, and TRSM's. This will be the proxy agent.
 - Set the SNMP mode to multipleAgent.
 - Set the proxy agent as the destination agent through which traps should be proxied by the other agents in the system. For the proxy agent, configure the destination IP address(es) of the management station(s) where traps should be forwarded.
 - Find the community strings needed for proxy in *IpsSnmpInternalProxyTable* located in the LANplex Systems MIB.
- Combining Methods II and III* To configure SNMP for both direct and indirect management with multiple agents, you must:
- Assign at least one IP address for the LMM+ and for each switching module that you want to manage directly.

- Set the SNMP mode to multipleAgent.
- For directly managed agents, set the destination IP address(es) of the management station(s) where traps should be forwarded by each agent.
- For indirectly managed agents, set the destination agent through which traps should be proxied.
- For indirectly managed agents, find the community strings needed for proxy in *IpsSnmpInternalProxyTable* located in the LANplex Systems MIB.

Displaying SNMP Settings

You can display the current LANplex SNMP configurations for the SNMP mode and community strings.

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

snmp display

If the system is set for the single agent SNMP mode, the mode setting and community string values for the agent are displayed as shown in the example below:

```
Current SNMP mode is singleAgent
Read-only community is public
Read-write community is private
```

If the system is set for multiple agent SNMP mode, the mode setting and community string values for each agent (by slot) are displayed as shown in the following example:

```
Current SNMP mode is multipleAgent
Slot 1 - Chassis
  Read-only community is public
  Read-write community is private
Slot 2 - ESM
  Read-only community is public
  Read-write community is private
Slot 3 - EFSM
  Read-only community is public
  Read-write community is private
```

The SNMP trap configuration display is available with the trap menu.

Top-Level Menu

```
system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout
```

```
display
mode
community
trap
```

Setting the Mode

The system is shipped with the SNMP mode default set for single agent. You should determine which SNMP mode you want to operate the system in and set the SNMP mode prior to performing any other SNMP configurations in the system.



CAUTION: *When you change the SNMP mode and reboot the system, you reset all the SNMP configurations you have set in the Administration Console to the default values.*

To configure the SNMP agent mode:

Top-Level Menu

```

system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout
  
```

```

display
mode
community
trap
  
```

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

```
snmp mode
```

- 2 Enter in a new mode (`singleAgent` or `multipleAgent`) or enter **q** to return to the previous menu.

If you entered in a new mode, the following prompts is displayed:

```
The new mode is different from the current mode. Changing
modes resets the SNMP configuration to the default values.
You must reboot the system for the new mode to take
effect.
```

```
Do you want to reboot the system now (n,y) [y]:
```

- 3 Enter **y** (yes) or **n** (no) at the prompt. If you entered **y** (yes) the system reboots. If you entered **n** (no), you are returned to the previous menu. The next time the system reboots, the new mode will take effect.

Configuring Community Strings

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

- *Read-only* with the default "public"
- *Read-write* with the default "private"

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

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

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

To set a community string:

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

```
snmp community
```

If the system is set for single agent SNMP mode, proceed to step 3. If the system is set for the multiple agent SNMP mode, you are prompted for the agent for which you want to configure community strings. You can configure only one agent at a time.

- 2 If multiple agent SNMP mode is enabled, enter the agent for which you want to set community string values. You can configure only one agent at a time.

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

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

The following example shows setting community strings when the system is operating in multiple agent SNMP mode:

```
Select agent by slot {1-4}: 2
Enter new read-only community [public]:
Enter new read-write community [private]: secret
```

Top-Level Menu

```
system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout

display
mode
community
trap
```

Administering SNMP Trap Reporting

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

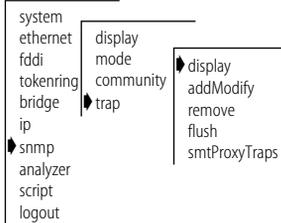
If the system is set for multiple agent SNMP mode, trap reporting is based on destination address (IP address of the SNMP manager) and/or a destination agent (number of an agent used for proxying traps to an SNMP manager). In both the trap information display and configuration, destination address and destination agent are used.

Displaying Trap Information

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

To show the trap reporting information configured:

Top-Level Menu



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

```
snmp trap display
```

If the system is set for single agent SNMP mode, the trap settings automatically appear for the single agent. If the system is set for the multiple agent SNMP mode, you are prompted for the agents for which you want to display trap information.

- 2 If multiple agent SNMP mode is enabled, enter the agent(s) for which you want to view the information. You can select individual agents or **all** to display trap information for all the agents in the system.

The display of the SNMP trap reporting information for a single agent is shown in the following example:

Trap Descriptions:

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

Trap Destinations Configured:

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

Proxying of remote SMT events is disabled



Trap 6: Power Supply Failure only appears in the trap list if the LANplex 6000 has a dual power supply. Trap 3: New Root, Trap 4: Topology Change, and Trap 10: Address Threshold only appear in the trap list if the LANplex system contains at least one ESM, EFSM, and TRSM.

If the system is set for multiple agent mode, the traps configured for each agent and the setting for proxying of remote SMT events is displayed.

```
Slot 1 - Chassis
  Trap Destinations Configured
  Address                Trap Numbers Enabled
  158.101.112.3         1-10, 12-21
  Proxying of remote SMT events is disabled

Slot 2 - ESM
  Trap Destination Agents Configured
  Agent Trap Numbers Enabled
  11-10, 12-21
  Proxying of remote SMT events is disabled

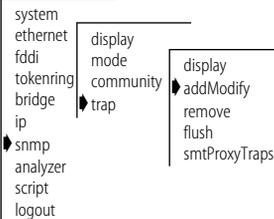
Slot 3 - EFSM
  No trap destination info configured
  Proxying of remote SMT events is disabled
```

Configuring Trap Reporting

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

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
```

If the system is set for single agent SNMP mode, proceed to step 3. If the system is set for multiple agent SNMP mode, you are prompted for either a destination agent or a destination address.

- 2 Enter **dstAgents** or **dstAddress** at the prompt.
- 3 Enter an IP address of the SNMP manager (destination address) or the number of the destination agent(s) at the prompt.
- 4 Enter the trap number(s) you want to enable for that destination at the prompt.

Separate a series of more than two trap numbers with a dash (-), and non-sequential trap numbers by commas. Enter **a11** 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 or agent(s) when the corresponding event occurs. Any traps that are not listed are not transmitted to the destination when the corresponding event occurs.

The following example shows a trap configuration for a system in single agent SNMP mode:

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

The following example shows trap configuration when the system is in multiple agent SNMP mode:

```
Select dst proxy agents or address (dstAgents,dstAddress)
[dstAddress]: dstAgents
Select agent(s) by slot (2-3|all) [2-3]: 2
Enter the trap numbers to enable (1-10,12-22|all)
[1-10,12-22]: all
```

Address Error

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

```
Trap address invalid or unreachable
```

Removing Trap Destinations

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

To remove a destination:

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

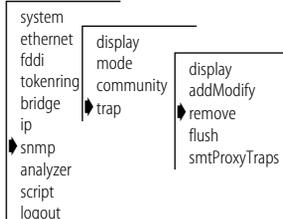
```
snmp trap remove
```

If the system is set for single agent SNMP mode, you are prompted for a trap reporting destination address. If the system is set for multiple agent SNMP mode, you are prompted for either a destination agent or a destination address.

- 2 Enter the SNMP trap reporting destination agent or address you want to remove at the prompt.

The destination agent(s)/address is removed and you are returned to the previous menu.

Top-Level Menu

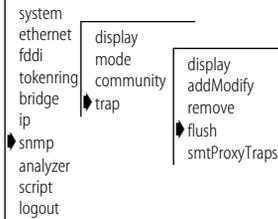


Flushing Trap Destinations

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

To flush all SNMP trap reporting destinations:

Top-Level Menu



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

```
snmp trap flush
```

If the system is set for single agent SNMP mode, you are automatically prompted to flush the trap destinations. Proceed to step 3. If the system is set for multiple agent SNMP mode, you are prompted for the agent(s) by slot.

- 2 Enter the slot number(s).

You receive the following prompt:

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

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

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

Setting up SMT Event Proxying

FDDI SMT events, which occur on the FDDI ring, may 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 remote SMT events to be reported by a LANplex 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 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 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 earlier section “Configuring Trap Reporting” 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 and 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 other non-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
```

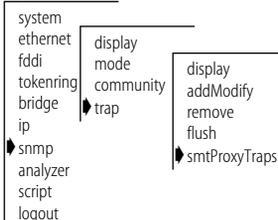
If the system is in single agent SNMP mode, proceed to step 3. If the system is in multiple agent SNMP mode, you are prompted for agents.

- 2 Enter the agent(s) for which you want to enable or disable proxying of remote SMT events.

- 3 Enter **enable** or **disable** at the prompt.

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

Top-Level Menu



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 date and time
- Rebooting the system

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
```

See the following example of a LANplex 6004 display:

```
LANplex 6004 (rev 0.2) - System ID 003000  
Software version 4.1.0 - Built 04/05/95 03:10:27 PM
```

Slot	Status	Contents
1	On-line	LANplex Management Module - Single Mac
2	On-line	Ethernet Switching Module (rev. 1.3)
3	On-line	Ethernet/FDDI Switching Module (rev. 0
4	On-line	Ethernet Switching Module (rev. 2.8)

Top-Level Menu

```
system  
ethernet  
fddi  
tokenring  
bridge  
ip  
snmp  
analyzer  
script  
logout  
display  
softwareUpdate  
baseline  
serialPort  
password  
name  
time  
screenHeight  
consoleLock  
panellLock  
ctlKeys  
nvData  
reboot
```

The display contains the following general system information:

- The system type (LANplex 6004)
- System ID
- Software version
- Software build date and time

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

- Slot number
- Module status (On-line or Off-line)
- Type of module (LMM+, FCM, ESM, EFSM, TRSM)
- Hardware revision number of each module

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

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



NOTE: *You can turn off the system bell by pressing any key on the LANplex system's control panel.*

Setting Passwords

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

The initial passwords stored in the nonvolatile memory of the LMM+ are null. You can only change the password if you enter the Console using the *administer* access level.

Top-Level Menu

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

To set a password:

- 1 From the top level of the Administration Console, enter:
system password
- 2 At the prompt requesting you to enter a password access level, enter one of the following:
read
write
administer
- 3 At the prompt for your old password, enter the old password. If you have never set a password or the password is null, press [Return].
- 4 Enter the new password.
The password can have up to 32 characters and is case sensitive. To enter a null password, press [Return].
- 5 Retype the new password for verification.
See the example below:
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 aid in system management. For example, you may want to name the system according to its physical location, LP2500 ENGLAB.

To name the system:

Top-Level Menu

```

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

```

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

system name

You are prompted for the name of the system, as shown below:

```
Enter new value [LANplex]:
```

- 2 Enter a name that is both unique on the network and meaningful to you. The new system name will appear the next time you display the system configuration.

Changing the Date and Time

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 time:

Top-Level Menu

```

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

```

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

system time

The current date and time are displayed, along with a prompt asking you if the date and time are correct. See the example below:

```
The current system time is 08/24/94 04:37:57 PM.
Is this correct? (y/n):
```

- 2 Enter **y** (yes) or **n** (no) at the prompt. If you respond **y**, you are returned to the main menu. If you respond **n**, you are prompted for the correct date and time.
- 3 Enter the correct date and time in the format, `mm/dd/yy hh:mm:ss xM`. The formats are described in Table 4-1.

Table 4-1 Date and Time Formats

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

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

See the example below:

```
Enter the new system time (mm/dd/yy hh:mm:ss xM):
09/30/94 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 Console through reboots so that you can view diagnostic information, your system must be connected through the terminal serial port.

To reboot the system:

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

```
system reboot
```

The following message appears:

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

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

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

Top-Level Menu

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



5

UPDATING SYSTEM SOFTWARE

This chapter explains how to install software updates to your system.



Refer to the LANplex 6000 Release Notes for the latest system software update information.

About Updating Software

LANplex system software is installed at the factory in flash memory on the LANplex Management Module Plus (LMM+). The software boots from flash memory. To update your system software, you can install a new version from any host running ftp.



CAUTION: *In order to run software version 6.0 on the LANplex 6000 you must have the new LMM Plus (LMM+) installed in your system. To verify if you have an LMM Plus installed check the module's ejector tab to ensure it says LMM+.*

To install or upgrade the system software, you must perform two tasks:

- Copy the software from the diskette to your UNIX-based or DOS-based computer's hard disk.
- Load the system software from your computer's hard disk to flash memory.

Copying Software to a Hard Disk

The software is distributed for both UNIX and DOS platforms. The following media types are used to distribute software releases:

- UNIX tar format 3 1/2-inch double-sided, high-density 1.44 MB diskette
- DOS format 3 1/2-inch double-sided, high-density 1.44 MB diskette

The software files are compressed on the media.

Copying to UNIX

The LANplex software for a UNIX-based hard disk is distributed on four floppy diskettes. Diskettes #1, #2, and #3 contain the LANplex software. Diskette #4 contains the SNMP MIBs.

The SNMP MIBs, on diskette #4, are provided so that you can compile on 3rd party applications.

To copy software to a UNIX hard disk, follow the instructions below:



If the directory "/usr/lp6000" does not exist on your computer, create the directory before proceeding. If your "/usr" directory is full, you can use a different directory. In this case, substitute the actual directory used for "/usr" in this and subsequent examples.

- 1 Insert diskette #1 containing the LANplex software file into a disk drive (these instructions assume drive rfd0).
- 2 Extract the first part of the LANplex software file using the following commands:

```
cd /usr/lp6000
tar xvf /dev/rfd0
```
- 3 Remove diskette #1 using the following command:

```
# eject
```
- 4 Insert diskette #2 containing the LANplex software file into a disk drive and extract the second part of the file using the following command:

```
tar xvf /dev/rfd0
```
- 5 Remove diskette #2 using the following command:

```
# eject
```
- 6 Insert diskette #3 containing the LANplex software file into a disk drive and extract the second part of the file using the following command:

```
tar xvf /dev/rfd0
```
- 7 Remove diskette #3 using the following command:

```
# eject
```

The following files should be in your current default directory:

- README1
- lp600000
- lp600001
- lp600002
- restore_lpx

- 8 Use the supplied script to decompress and restore the split file (lp600000, lp600001, and lp600002).

```
# ./restore_lpx
```

Restoring the split file creates the uncompressed file **lp6000**. See the readme file for size and checksum information.

Copying to DOS

The LANplex software for a DOS-based hard disk is distributed on two floppy diskettes. Diskette #1 contains the LANplex software. Diskette #2 contains the SNMP MIBs.

To copy software to a DOS hard disk, follow the instructions below:



If the directory "lp6000" does not exist on your computer, create the directory before proceeding.

- 1 Insert the diskette containing the software file into a disk drive (these instructions assume drive B:).
- 2 Copy the software file to the **lp6000** directory of your computer using the following commands:

```
cd lp6000  
copy b:lp6000.exe
```



The file lp6000.exe is a self-extracting archive. It decompresses and creates the loadable lanplex file.

- 3 Decompress the file using the following command:

```
lp6000
```

This creates a file called **lp6000**, which you can then load into flash memory.

Loading Software

Before loading the system software on the LMM Plus, verify that the host machine, which has a copy of the updated system software, is connected to the system by one of the methods described in Chapter 3: *Configuring Management Access to the System*.



You can load the system software into flash memory while the system is operating. You do not need to bring the system down. After the flash install is completed, a reboot will put the newly-loaded software to use.



If you are loading software from a PC, the ftp server must be running on the PC before beginning this procedure.

How long will a software load take?

Loading software into flash memory takes approximately 10 to 15 minutes to complete, depending on your network load.

To load the system software:

Top-Level Menu

```

system
ethernet
fddi
tokenring
bridge
ip
snmp
analyzer
script
logout
display
softwareUpdate
baseline
serialPort
password
name
time
screenHeight
consoleLock
panelLock
ctlKeys
nvData
reboot

```

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

system softwareUpdate

You are prompted for the Host IP address, Install file path name, User name, and Password. The current values are displayed in brackets []. To use the value in brackets, press [Return]. The password field does not display what you enter.

- 2 Next to `Host IP address`, enter the IP address of the host machine from which you are installing the software (such as a Sun workstation or PC).

In the following example, the IP address of the host is **192.9.200.96**.

- 3 Next to `Install file name`, enter the complete path and file name.



For DOS system syntax, you must precede the full pathname with a forward slash (/). For example, if you are loading software from a DOS host, enter the following at the Install Filename prompt:

- 4 Next to `User name`, enter your user name.
- 5 Next to `Password`, enter your password. You *must* enter a value for this field.

See the following screen for an example of the software installation prompts.

```
Host IP address [192.9.200.14]:192.9.200.96  
Install file path name [/usr/lp6000/lp6000]:  
User name: ronnyk  
Password:
```

After the software is loaded, you are notified that installation has been completed:

```
Installation complete.
```



If the LANplex executable software image stored in EPROM is corrupted (for example, when a power failure occurs while you are updating software), contact 3Com Technical Support. See Appendix B: Technical Support.

6

BASELINING STATISTICS

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

About Setting Baselines

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

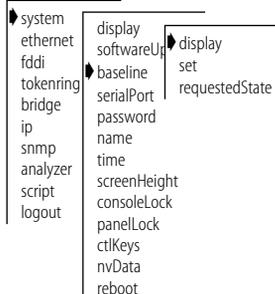
Baselining is maintained across Administration Console sessions. Statistics you view after setting the baseline indicate that they are relative to the baseline. To view statistics relative to the most recent power up, you must disable the baseline.

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

Displaying the Current Baseline

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

Top-Level Menu



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

```
system baseline display
```

See the example below:

```
Baseline was set at 08/24/94 04:43:38 PM.
```

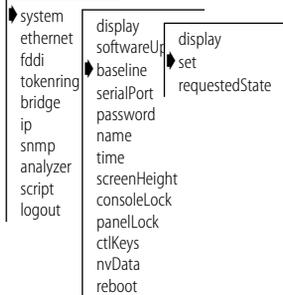
If a baseline has not been set on the system, the following message is displayed:

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

Setting Baselines

Setting a baseline resets the counters to zero (the accumulated totals since power up are still maintained by the system). The baseline is timestamped.

Top-Level Menu



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

```
system baseline set
```

A message similar to the following appears:

```
Baseline set at 08/24/94 04:43:38 PM.
```

Baselining is automatically enabled once 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:

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

```
system baseline requestedState
```

You are prompted to enter a new baseline state, as shown in the example below:

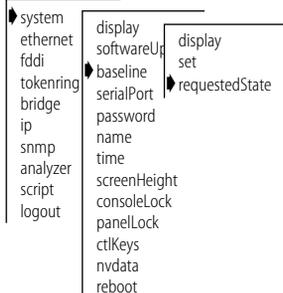
```
Enter new value (disabled,enabled) [enabled]:
```

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

The new value is confirmed as shown below:

```
Baseline set at 10/24/94 04:44:45 PM has been disabled.
```

Top-Level Menu



7

SAVING, RESTORING, AND RESETTING NONVOLATILE DATA

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

About Working with Nonvolatile Data

If you want to transfer NV data from one set of modules to another set or from one module to another, save the system's NV data and restore it as appropriate. You may 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 on each module 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 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 the NV data is restored:

- Software version number
- System ID
- Date and time of creation
- Type of configuration
- Data checksums
- List of stored modules

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 upon a restore.

To save NV data:

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

```
system nvData save
```

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

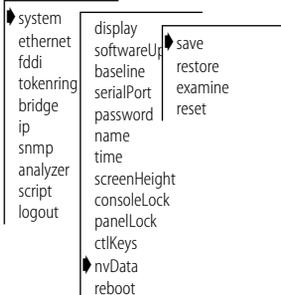
- 2 Enter the IP address of the station to which you want to save the NV data.
- 3 Enter the file path name where you want to save the file.
- 4 Enter your user name on the host system.
- 5 Enter your password on the host system.
- 6 Enter a name of the file (this is optional).

See the following example:

```
Host IP Address [158.101.100.1]: 158.101.112.34
Select file name [usr/jones/systemdata]:
User name [<none>]: Tom
Password [<none>]:
Enter an optional file label [<none>]: Labdata
```

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

Top-Level Menu



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

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

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

Failure Error - Configuration not stored.

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

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

*What if the module
configuration is
altered during a save?*

During the save procedure, the current module configuration could be altered. To detect this event, the software checksums the NVRAM before and after the save. If the checksum is different, you are notified and prompted to save the configuration again. In abnormal situations, this could go on for an extended period, so you are given the option to terminate the save. You are also prompted for a retry request upon network (ftp) I/O failures. See the following example:

```
Error - Checksum failure, configuration of slot 4 changed  
during save.  
Error - Configuration not stored.  
Do you wish to retry the save using the same parameters?  
(y/n):
```

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 below:

Rule 1 *Exact Match* — An exact match is one where the system IDs, module types, slot assignments, and module revisions (if applicable) all match between the saved configuration and the system on which you are restoring the image.

Rule 2 System ID Mismatch — System IDs do not match between the saved NV file and the target system. Mismatches in system IDs are allowed. Before restoring the NV data to a system with a different system ID, you should be aware of the following NV data that may cause problems when restored:

- Management IP addresses (defined in IP interface configurations) are saved as NV data and restored. Before connecting the restored system to the network, you may need to change the IP address of defined interfaces to avoid duplicate IP address problems. Modifying IP interface definitions is described on page 3-9.
- Statically configured Ethernet addresses on switching modules are saved as NV data. You must ensure not to have duplicate addresses when you restore the NV data. Listing statically configured addresses is described on page 13-12.

Rule 3 Subset or Superset — The module configuration is a superset or subset of the original system, and the software proposes a mapping. If one or more exact matches are found, the proposed mapping consists *only* of the exact matches. This would be useful if modules had been removed from a system or modules were added. You would have to individually configure the modules not in the saved set.

In a mismatch situation, if you decide not to use the proposal, the software allows you to selectively restore individual modules. This is useful if you have hot-swapped a module and you want to restore just the configuration of that module. The data that you restore on a module is only the data on that module. It does not include information stored on the LMM+ that relates to that module.

It is possible to cause inconsistencies in the system during manual assignment. Such an action must be made carefully.

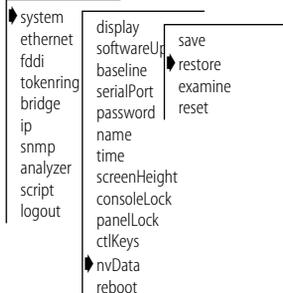
If none of these rules succeed, it is illegal to apply the saved configuration to the system.

To restore the NV data:

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

system nvData restore

Top-Level Menu



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

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

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

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

If a session is successfully opened, the system reads the header information, compares the stored configuration to the current system configuration, and proposes a method of restoration based on one of the restoration rules described on page 7-3.

You are prompted to load the proposal.

```
CAUTION - Restoring nonvolatile data may leave the system
in an inconsistent state and therefore a reboot is
necessary after each restore.
Do you wish to continue? (y/n):
```

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

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

If you enter **n**, the entire saved configuration is displayed for you to manually load. The system prompts you as shown in the example below:

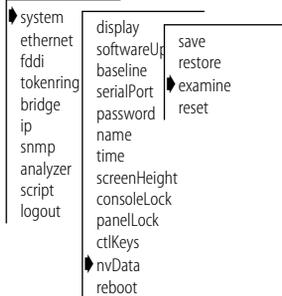
```
There is an ESM in slot 3 of the system which may be
loaded with saved NV data.
Do you wish to perform this load [yes, no, quit]:
```

Entering **quit** terminates the manual selections, but it restores any selections you have already confirmed.

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

Examining a Saved NV Data File

Top-Level Menu



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

To examine the file:

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

```
system nvData examine
```

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

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

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

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

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

```
Product ID #, Product Type #
System ID 102, 12 slots
Saved October 8, 1994 10:24:12. Configuration version 3.
Slot Contents
  1 LANplex Management Module
  2 FDDI Concentrator Module
  3 No configuration saved
  4 Ethernet Switching Module
  5 Ethernet/FDDI Switching Module
  6 No configuration saved
  7 Ethernet/FDDI Switching Module
  8 FDDI Concentrator Module
  9 Ethernet Switching Module
 10 No configuration saved
 11 No configuration saved
 12 No configuration saved
```

You are returned to the NV data menu options.

Resetting NV Data to Defaults

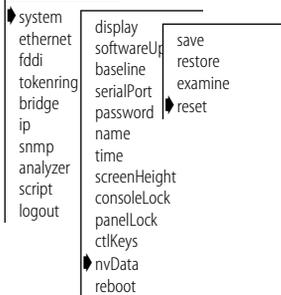
At times you may not want to restore the system NV data. Instead, you may want to reset the values to the factory defaults so that you can start configuring the system from the original settings.



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

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

Top-Level Menu



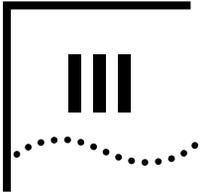
- 1 From the top level of the Administration Console, enter:

```
system nvData reset
```

You see the following prompt:

```
Resetting NV data may leave the system in an inconsistent
state and therefore a reboot is necessary after each reset.
Do you wish to continue (n,y) [y]:
```

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



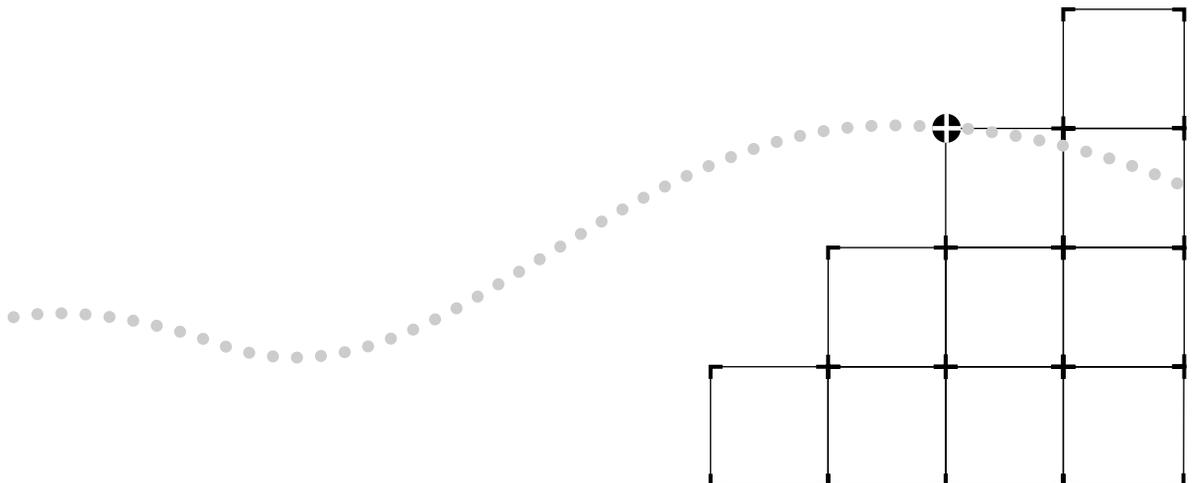
ETHERNET, FDDI, & TOKEN RING

Chapter 8 Administering Ethernet Ports

Chapter 9 Administering FDDI Resources

Chapter 10 Administering Token Ring Ports

Chapter 11 Setting up the System for Roving Analysis



8

ADMINISTERING ETHERNET PORTS

This chapter describes how to:

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

Displaying Ethernet Port Information

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

If you want to display Ethernet port statistics relative to a baseline, see Chapter 6: *Baselining Statistics* for more information.

To display information about the Ethernet ports:

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

```
ethernet summary
```

OR

```
ethernet detail
```

- 2 At the prompt, select the slot(s) related to the port(s) about which you want to view information.
- 3 Enter the port(s) for which you want to view information.

The port information is displayed in the format you specified. .

Top-Level Menu

```
system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout
summary
detail
label
portState
```

The following example shows a detailed display for Ethernet ports on an EFSM:

```

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
  1          Office113_SPARCstation5
 12          Office322_Quadra900

```

An example of a summary display for Ethernet ports on an EFSM is shown below:

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 no	noRxBuffers	txQOverflows
1	0	0	0	0
12	0	0	0	0

Table 8-1 describes the type of information provided about an Ethernet port.

Table 8-1 Description of Fields for Ethernet Port Attributes

Field	Description
alignmentErrs	Number of frames received by this port that are not an integral number of octets in length and do not pass the FCS check
carrierSenseErr	Number of frames discarded because the carrier sense condition was lost while attempting to transmit a frame from this port
collisions	Number of collisions detected on this port
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

(continued)

Table 8-1 Description of Fields for Ethernet Port Attributes (continued)

Field	Description
noRxBuffers	Number of frames discarded because there was no available buffer space
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. Values for each port type are 10BASE-T (RJ-21), 10BASE-T (RJ-45), 10BASE-5 (AUI), 10BASE-FL, and 10BASE-2 (BNC).
requestedState	Configurable parameter used to enable/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
rxErrs	Sum of all receive errors associated with this port (field only appears in the summary option)
rxFrameRate	Average number of frames received per second by this port during the most recent sampling period
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 (non-multicast) 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 frames discarded because the port was disabled
txErrs	Sum of all transmit errors associated with this port (field only appears in the summary option)

(continued)

Table 8-1 Description of Fields for Ethernet Port Attributes (continued)

Field	Description
txFrameRate	Average number of frames transmitted per second by this port during the most recent sampling period
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 (non-multicast) 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. Frames may be discarded, however, for the following reasons:

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

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

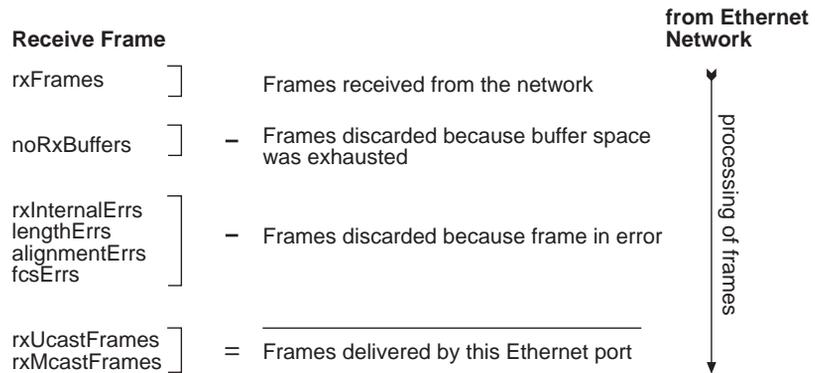


Figure 8-1 How Frame Processing Affects Ethernet Receive Frame Statistics

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

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

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

Transmit Frame Statistics

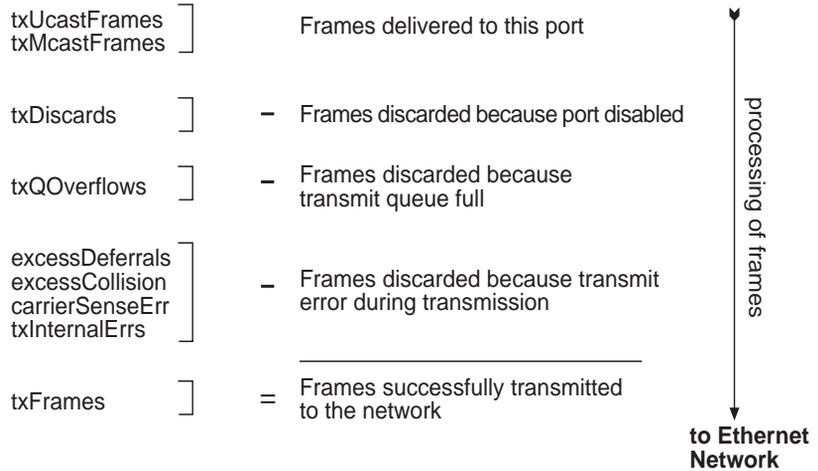


Figure 8-2 How Frame Processing Affects Ethernet Transmit Frame Statistics

Labeling a Port

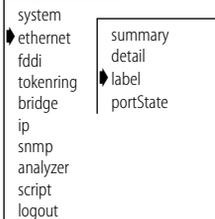
Port labels serve as a useful reference point and as an accurate means of identifying your ports for management. You may want to label your Ethernet ports so that you can easily identify the device specifically attached to each port (for example, LAN, workstation, or server).

To label an Ethernet port:

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

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

Top-Level Menu



Setting the Port State

You can enable (place on-line) 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 At the prompt, enter the module(s) for which you want to set port states.
- 3 Enter the number(s) of the port(s) you want to label.
- 4 Enter **enable** or **disable** for each Ethernet port.

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

Top-Level Menu

```
system
└─ ethernet
└─ fddi
└─ tokenring
└─ bridge
└─ ip
└─ snmp
└─ analyzer
└─ script
└─ logout
```

```
summary
detail
label
└─ portState
```

9

ADMINISTERING FDDI RESOURCES

This chapter describes how to display information about and configure:

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



This chapter covers advanced FDDI topics and 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 6000 Operation Guide*.

About Configuring FDDI Resources

Before you configure the FDDI resources for your LANplex system, you must first decide if you want the three FDDI backplane paths to operate as part of a single FDDI station or part of three FDDI stations, and configure your system for that mode of operation. As the default, the FDDI backplane paths are set to operate as part of a single FDDI station (single station mode).

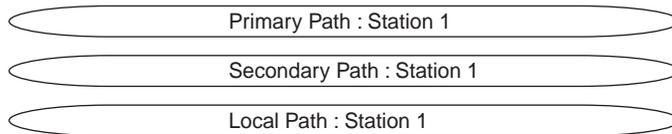
Because the LANplex 6000 is a slot-based system, configuring stations, paths, MACs, and ports is performed based on the slot of the module to which these FDDI resources are related.

The Backplane Path Mode

Figure 9-1 shows how the LANplex 6000 FDDI backplane paths change from single station mode to multi-station mode.

- In **single station mode**, all MACs and ports that connect to the backplane paths belong to Station 1. A MAC or port can be moved to a new backplane path by changing the requested path configuration, which allows these resources to be assigned to the primary, secondary, or local path within Station 1.
- In **multi-station mode**, all MACs and ports that connect to the backplane paths initially belong to Station 1 (which corresponds to the first backplane path). A MAC or port can be moved to a new backplane path by first assigning it to a new station (Station 2 or 3), and then selecting the desired path within that station.

Single Station Mode



Multi-station Mode

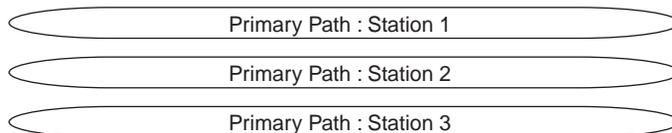


Figure 9-1 LANplex 6000 Backplane Paths in Single Station and Multi-station Modes

See the *LANplex 6000 Operation Guide* for more information about setting up the LANplex system's backplane paths.

Changing to multi-station

In general, if you want to change the mode of operation to multi-station, you should:

- 1 Set the mode to multi-station.
- 2 Reboot the system.
- 3 Assign each backplane MAC and port to the station corresponding to the desired backplane path. See Figure 9-1.
- 4 Reboot the system again.

Modules and FDDI Resources

The FDDI resources available in your LANplex system depend on the type of modules installed in your system. Available modules include:

- LANplex Management Module (LMM+)
- FDDI Concentrator Module (FCM)
- Ethernet/FDDI Switching Module (EFSM)
- Ethernet Switching Module (ESM)
- Token Ring Switching Module (TRSM)

When configuring FDDI resources, you must access these resources through the slot containing the module related to that resource.

Table 9-1 shows an example of modules in a LANplex 6004 system and the possible FDDI resources available with each module.

Table 9-1 Example of Modules and Their Possible FDDI Resources

Slot Number	Module Type	FDDI Resources Available
Slot 1	LMM+	0–3 MACs, 2 A/B ports
Slot 2	FCM	6–12 Master (M) ports
Slot 3	EFSM	2 MACs, 0–1 external port
Slot 4	ESM	1 MAC

For example, if you wanted to change the station assignment of the EFSM MACs, you would select Slot 3 when prompted.

Configuring the Backplane Path Mode

The LANplex 6000 has three FDDI backplane paths that can be configured in one of two ways:

- Single station mode — The three rings are used as the primary, secondary, and local paths for a single FDDI station.
- Multi-station mode — The three rings are used as the primary paths for three different FDDI stations.



You should configure the FDDI backplane path mode prior to performing additional FDDI configurations for your LANplex 6000 system.



CAUTION: *When you change the FDDI backplane path mode and reboot the system, you reset the station and path assignments of all MACs and ports to the default values.*

To set the FDDI backplane path mode:

Top-Level Menu

```

system
ethernet
fddi
tokenring
bridge
ip
snmp
analyzer
script
logout
station
path
mac
port
backplane

```

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

fddi backplane

The following description and prompt appears:

The FDDI backplane paths can be assigned in one of two ways:

singleStation: The three paths are used as the primary, secondary, and local paths for a single FDDI station (station 1).

multiStation: The three paths are used as the primary paths for three different FDDI stations (stations 1, 2, and 3).

Enter new value (singleStation,multiStation)
[singleStation]:

- 2 At the prompt, enter **multiStation** or **singleStation**.

If you are changing to multi-station mode, the following description and prompt are displayed:

The system must be rebooted for the new mode to take effect.

When the system has rebooted, each FDDI MAC and port that connects to a backplane path will belong to FDDI station 1 (whose primary path corresponds to backplane path 1). If you wish to configure any one of these MACs or ports to backplane path 2 or 3, you will have to assign it to the corresponding FDDI station using the "station" menu option in the "mac" or "port" menu.

If you are changing to single station mode, the following description is displayed:

The system must be rebooted for the new mode to take effect.

In either case, you are prompted to reboot the system:

Do you want to change the mode and reboot the system (n,y)
[y]:

- 3 At the reboot system prompt, enter **y**.
The system reboots and the new mode takes effect.

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).

If the FDDI backplane paths are set to single station mode, then they are all part of a single FDDI station. If the paths are set to multiple station mode, then each backplane path is part of a different FDDI station.

You can display station information and set the following parameters in either mode:

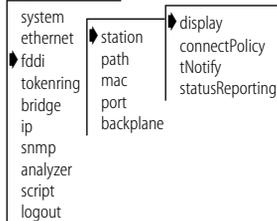
- Connection policies
- Neighbor notification timer
- Status reporting

Displaying Station Information

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

To display FDDI station information:

Top-Level Menu



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

fddi station display

You are prompted for a slot related to the FDDI station.

- 2 Enter the slot(s) related to the FDDI station(s).

You are prompted for a station.



If the system is in multi-station mode, stations only appear if MACs or ports of the module in the slot you select is assigned to one of the FDDI backplane path stations.

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

See the following example of station information:

```

configuration      tNotify statusReporting  connectPolicy
  isolated          30          enabled          0x8000

  ecmState remoteDisconnect  traceMaxExp
    in           false          87500000

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

```

This station display was generated from an LMM+ MAC attached to Station 1 of the system. Table 9-2 describes these statistics.

Table 9-2 Description of Fields for FDDI Station Attributes

Field	Description
configuration	Attachment configuration for the station or concentrator. Values can be Thru, Isolated, Wrap_A, and Wrap_B.
connectPolicy	Bit string representing the connection policies in effect on a station. How connection policies translate into bits is described in Table 9-3. This value can be user-defined.
ecmState	Current state of the ECM state machine
remoteDisconnect	Flag indicating that the station was remotely disconnected from the network as a result of receiving an fddiSMTAction with the value of <i>disconnect</i> in a Parameter Management Frame (PMF). A station requires a Connect Action to rejoin and clear the flag.
station ID	Unique identifier for an FDDI station
statusReporting	Value indicating whether <i>statusReporting</i> is enabled or disabled for the station. This attribute controls whether a station generates Status Report Frames (SRFs) to report events and conditions to network management stations. This value can be user-defined.
tNotify	Timer used in the Neighbor Notification protocol to indicate the interval of time between the generation of Neighbor Information Frames (NIF). This value can be user-defined.
traceMaxExp	Maximum propagation time for a Trace on an FDDI topology. Places a lower bound on the detection time for an unrecovering ring.

Setting the Connection Policies

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

The LANplex 6000 modules with FDDI media options have for the following FDDI connection types:

- LMM+ — A, B, or M
- EFSM — M or S
- FCM — M

By default, all connections to the LANplex 6000 FDDI ports are valid, except for M-M connections. The possible connections to reject and their corresponding bits are listed in Table 9-3.

Table 9-3 Bit to Set for Rejecting a Station Connection

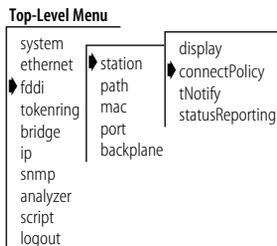
Connection Rejected... (port - Remote port)	If This Bit Is Set	Connection Rules
A-A	0	Undesirable peer connection that creates twisted primary and secondary rings; notify SMT.
A-B	1	Normal trunk ring peer connection.
A-S	2	Undesirable peer connection that creates a wrapped ring; notify SMT.
A-M	3	Tree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

(continued)

Table 9-3 Bit to Set for Rejecting a Station Connection (continued)

Connection Rejected... (port - Remote port)	If This Bit Is Set	Connection Rules
B-M	7	Tree connection with possible redundancy. The node shall not go to Thru state in CFM. In a single MAC node, Port B shall have precedence (with defaults) for connecting to a Port M.
S-A	8	Undesirable peer connection that creates a wrapped ring; notify SMT.
S-B	9	Undesirable peer connection that creates a wrapped ring; notify SMT.
S-S	10	Connection that creates a single ring of two slave stations.
S-M	11	Normal tree connection.
M-A	12	Tree connection that provides possible redundancy.
M-B	13	Tree connection that provides 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:



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

```
fddi station connectPolicy
```

You are prompted for a slot related to the FDDI station.

- 2 Enter the slot(s) related to the FDDI station(s).

You are prompted for a station.

- 3 Enter the station for which you want to set the connection policies.

- 4 Enter the value of the connection policy for that station.

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

See the following example:

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

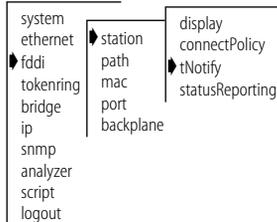
Setting Neighbor Notification Timer

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

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

To set the *T-notify* timer:

Top-Level Menu



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

```
fddi station tNotify
```

You are prompted for a slot related to the FDDI station.

- 2 Enter the slot(s) related to the FDDI station(s).

You are prompted for a station.

- 3 Enter the station.

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

Valid values are 2–30 seconds.

See the following example:

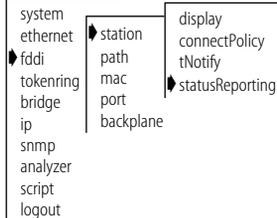
```
Select station [1]:
Station 1 - Enter new value [30]:
```

Enabling/Disabling Status Reporting

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

To enable or disable status reporting for a station:

Top-Level Menu



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

```
fddi station statusReporting
```

You are prompted for a slot related to the FDDI station. For example, if you enter slot 1, you could choose the stations related to the FDDI MACs of the LMM+.

- 2 Enter the slot(s) related to the FDDI station(s).

You are prompted for a station.

- 3 Enter the station number.

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

See the following example:

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

Administering FDDI Paths

FDDI's dual, counter-rotating ring consists of a primary 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 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:

- tvxLowerBound
- tmaxLowerBound
- maxTreq

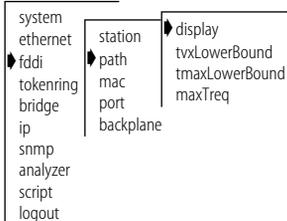
These values are used by all MACs configured in a path.

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:

Top-Level Menu



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

fddi path display

You are prompted for a slot related to the FDDI station about which you want to view path information.

- 2 Enter the slot(s).

You are prompted for a station and path.

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

If you are in multi-station mode, the stations available for you to choose are those for which you have assigned MACs and ports.

4 Enter the path (**p** = primary, **s** = secondary, **l** = local).

In single station mode, the primary, secondary, and local paths are available for selection. In multi-station mode, the primary and secondary paths are available for stations related only to the LMM+ (the secondary path is local to the LMM+) and the primary path is available for all stations related to the FCM, EFSM, and ESM.

See the following example of path information:

```

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

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

```

Notice in the example that the system is set for single station mode. Table 9-2 describes these statistics.

Table 9-4 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 fddiMACT-vxValue that will be used by any MAC that is configured in this path. This value can be user-defined.

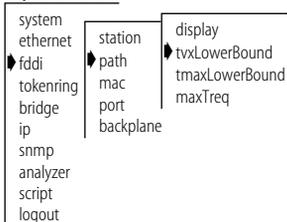
Setting tvxLowerBound

The *tvxLowerBound* attribute specifies the minimum time value of *fddiMAC* *TvxValue* 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 may be reinitialized when there is no problem. The higher you set this value, the less chance of frequent reinitializations, but the network will take longer to recover from errors.

To set *tvxLowerBound*:

Top-Level Menu



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

```
fddi path tvxLowerBound
```

You are prompted for a slot related to the FDDI station.

- 2 Enter the slot(s).

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

- 3 Enter the station.

If you are in multi-station mode, the stations available for you to choose are those for which you have assigned MACs and ports.

- 4 Enter the path (**p** = primary, **s** = secondary, **l** = local).

In single station mode, the primary, secondary, and local paths are available for selection. In multi-station mode, the primary and secondary paths are available for stations related only to the LMM+ (the secondary path is local to the LMM+) and the primary path is available for all stations related to the FCM, EFSM, and ESM.

- 5 Enter the new minimum time value.

The default is 2500 micro-seconds (us).

See the following example:

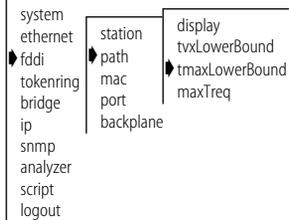
```
Select station [1]:
Select path(s) (p,s|all) [p]:
Station 1 Primary - Enter new value [2500]:
```

Setting tmaxLowerBound

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

To set tmaxLowerBound:

Top-Level Menu



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

```
fddi path tmaxLowerBound
```

You are prompted for a slot related to the FDDI station.

- 2 Enter the slot(s).

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

- 3 Enter the station.

If you are in multi-station mode, the stations available for you to choose are those for which you have assigned MACs and ports.

- 4 Enter the path (**p** = primary, **s** = secondary, **l** = local).

In single station mode, the primary, secondary, and local paths are available for selection. In multi-station mode, the primary and secondary paths are available for stations related only to the LMM+ (the secondary path is local to the LMM+) and the primary path is available for all stations related to the FCM, EFSM, and ESM.

- 5 Enter the new minimum time value.

The default is 165000 micro-seconds (us).

See the example below:

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

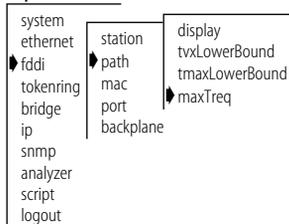
Setting maxT-req

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

When T_Opr is a low value, the token rotates more quickly, so token latency is reduced. However, more of the ring's available bandwidth is used to circulate the token. Higher values of T_Opr use less bandwidth circulating the token, but 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 slot related to the FDDI station.

- 2 Enter the slot(s).

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

- 3 Enter the station.

If you are in multi-station mode, the stations available for you to choose are those for which you have assigned MACs and ports.

- 4 Enter the path (**p** = primary, **s** = secondary, **l** = local).

In single station mode, the primary, secondary, and local paths are available for selection. In multi-station mode, the primary and secondary paths are available for stations related only to the LMM+ (the secondary path is local to the LMM+) and the primary path is available for all stations related to the FCM, EFSM, and ESM.

- 5 Enter the new minimum time value.

The default value is 165000 micro-seconds (us)

See the following example:

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

Administering FDDI MACs

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

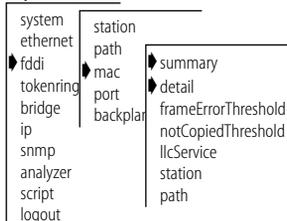
- MAC FrameErrorThreshold
- NotCopiedThreshold
- Logical Link Control (LLC) service

Displaying MAC Information

FDDI MAC information can be viewed in a summary or in a detailed format. 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 of FDDI MAC information includes the information in the summary and additional FDDI MAC statistics.

To view FDDI MAC summary or detail 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 slot related to the MAC.

- 2 Enter the slot(s).
You are prompted for a MAC number.
- 3 Enter the MAC number.

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

slot	mac	rxFrames	txFrames	rxBytes	txBytes
4	1	101884	34320	22689080	10257112
4	2	97115	31939	21782677	9546481

slot	mac	Errors	noRxBuffers	txQOverflows
4	1	0	0	0
4	2	0	0	0

slot	mac	upstream	downstream
4	1	00-80-3e-02-95-16	00-80-3e-02-95-41
4	2	00-80-3e-02-95-40	00-80-3e-02-95-01

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

slot	mac	rames	rxBytes	rxFrameRate	rxByteRate
4	1	03666	23089968	36	7582
4	2	98897	22183565	36	7582
slot	mac	eRate	rxPeakByteRate	lostCount	lateCount
4	1	48	10308	0	0
4	2	48	10308	0	0
slot	mac	Count	notCopiedThresh	notCopiedRatio	notCopiedCond
4	1	0	6550	0	inactive
4	2	0	6550	0	inactive
slot	mac	Count	frameErrThresh	frameErrorRatio	frameErrCond
4	1	0	655	0	inactive
4	2	0	655	0	inactive
slot	mac	ffers	tvxExpiredCount	rxInternalErrs	rxDiscards
4	1	0	0	0	32923
4	2	0	0	0	65078
slot	mac	casts	rxMulticasts	txFrames	txBytes
4	1	34621	36158	34921	10437189
4	2	78	33778	32541	9726635
slot	mac	eRate	txByteRate	txPeakFrameRate	txPeakByteRate
4	1	15	4511	23	6911
4	2	15	4511	16	4800
slot	mac	lErrs	txQOverflows	txDiscards	txUnicasts
4	1	0	0	0	34861
4	2	0	0	15	32493
slot	mac	casts	frameCount	tokenCount	ringOpCount
4	1	94	280867	1331364113	4
4	2	96	262339	1230834352	3
slot	mac	tPath	dupAddrTest	duplicateAddr	upstreamDupAddr
4	1	imary	passed	false	false
4	2	imary	passed	false	false
slot	mac	lable	llcService		smtAddress
4	1	true	enabled		00-80-3e-02-95-40
4	2	true	enabled		00-80-3e-02-95-41
slot	mac		upstream		downstream
4	1		00-80-3e-02-95-16		00-80-3e-02-95-41
4	2		00-80-3e-02-95-40		00-80-3e-02-95-01
slot	mac		oldUpstream		oldDownstream
4	1		unknown		00-80-3e-02-95-01
4	2		unknown		unknown
slot	mac	mType	rmtState	tMaxCapab	tvxCapab
4	1	known	ring op	1342200 us	1342200 us
4	2	known	ring op	1342200 us	1342200 us
slot	mac	tReq	tNeg	tMax	tvxValue
4	1	86 us	164986 us	167770 us	2621 us
4	2	86 us	164986 us	167770 us	2621 us

Table 9-5 describes the type of information provided for the FDDI MAC.

Table 9-5 Description of Fields for FDDI MAC Attributes

Field	Description
currentPath	Path on which this MAC is currently located (primary or secondary)
downstream	MAC address of this MAC's downstream neighbor
downstreamType	Indicates the PC type
dupAddrTest	Pass/fail test for a duplicate address
duplicateAddr	Indication of whether this address is duplicated on the FDDI ring
errorCount	Number of reportable frame errors detected by this MAC
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 shall 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 there was no buffer space 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 shall be generated
oldDownstream	Previous value of the MAC address of this MAC's downstream neighbor

(continued)

Table 9-5 Description of Fields for FDDI MAC Attributes (continued)

Field	Description
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 non-operational 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 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 (non-multicast) 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
tvxCapab	Maximum time value of the valid transmission timer this MAC can support

(continued)

Table 9-5 Description of Fields for FDDI MAC Attributes (continued)

Field	Description
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	Indication of 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. A frame may be discarded, however, 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 9-2 shows the order in which these discard tests are made.

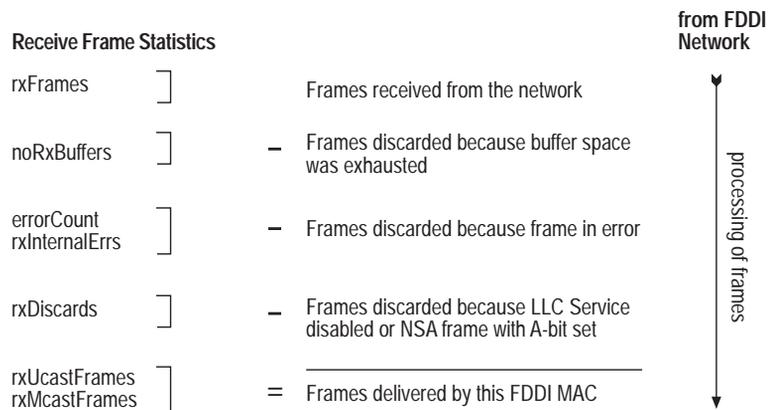


Figure 9-2 How Frame Processing Affects FDDI MAC Receive Frame Statistics

Frames are delivered to an FDDI MAC by bridges and management applications. However, a frame may 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 9-3 shows the order in which the discard tests are made.

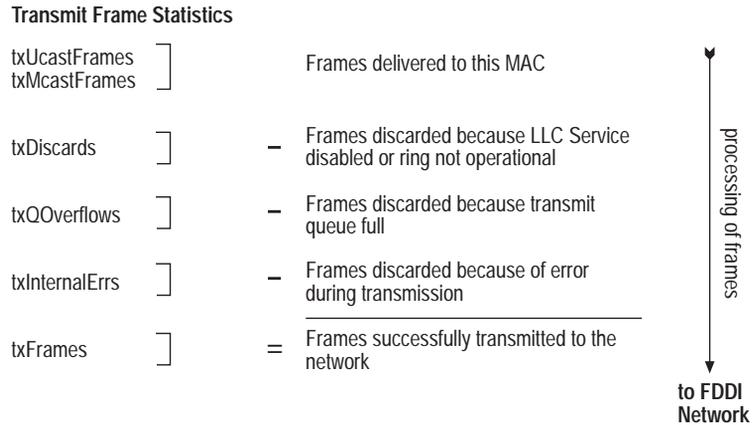


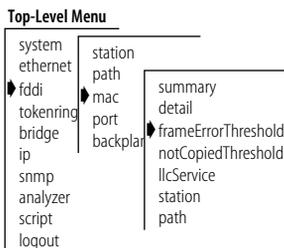
Figure 9-3 How Frame Processing Affects FDDI MAC Transmit Frame Statistics

Setting the Frame Error Threshold

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

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

To set the *FrameErrorThreshold*:



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

fddi mac frameErrorThreshold

You are prompted for a slot related to the MAC.

- 2 Enter the slot(s).

You are prompted for a MAC and new value.

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

See the following example:

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

Setting the Not Copied Threshold

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

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

To set the *NotCopiedThreshold*:

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

```
fddi mac NotCopiedThreshold
```

You are prompted for a slot related to the MAC.

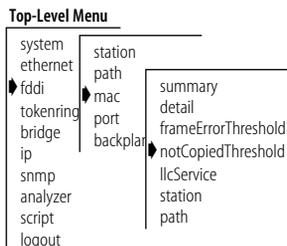
- 2 Enter the slot(s).

You are prompted for a MAC and new threshold value.

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

See the following example:

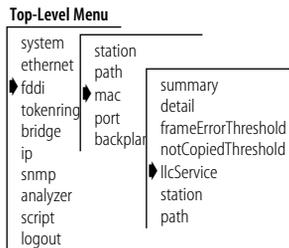
```
Select MAC [1]:
MAC 1 - Enter new value [655]:
```



Enabling/Disabling LLC Service

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

To enable or disable LLC service for the MACs in the LANplex system:



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

```
fddi mac llcService
```

You are prompted for a slot related to the MAC.

- 2 Enter the slot(s).

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

- 3 Enter the MAC number.

- 4 Enter the new MAC value (**enabled** or **disabled**).

See the following example:

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

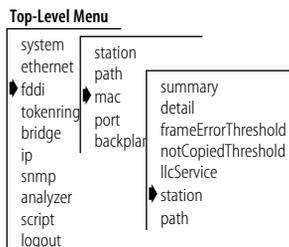
```
MAC 1 - Enter new value (disabled,enabled) [enabled]:
```

```
disabled
```

Assigning MACs to Stations in Multi-station Mode

You can only assign a MAC to a new station when the FDDI backplane paths are set for multi-station mode. If you access this menu item when the FDDI backplane paths are set for single station mode, the following message appears:

```
The FDDI backplane path mode is set to "singleStation" so all three backplane paths are part of the same station. This menu item is only useful if the FDDI backplane path mode is set to "multiStation".
```



To assign MACs to stations:

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

```
fddi mac station
```

You are prompted for a slot related to the MAC.

2 Enter the slot(s).

You are prompted for a station assignment for each MAC in the slot(s) you specified.

3 Enter the station assignment. The possible values are 1, 2, or 3.



Remember that each of these stations uses the corresponding backplane path as its primary path. See Figure 9-1.

See the following example:

```
Select slot(s) (2,4|all) [2,4]: all
Select MAC(s) (1-2|all): all
Slot 2 MAC 1 - Enter new FDDI station number (1-3) [1]:
Slot 4 MAC 1 - Enter new FDDI station number (1-3) [1]: 2
You have requested that this MAC be moved to a different
FDDI station. You must reboot your system for this
request to take effect.
Slot 4 MAC 2 - Enter new FDDI station number (1-3) [1]: 3
You have requested that this MAC be moved to a different
FDDI station. You must reboot your system for this
request to take effect.
```

If you change a station assignment, then you must reboot the system for the change to take effect.

Setting the MAC Paths

In single station mode, the possible backplane path assignments include isolated, primary secondary and local. In multi-station mode, the backplane path assignments include isolated and primary (because all backplane paths become primary paths in this mode).

To assign MACs to paths:

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

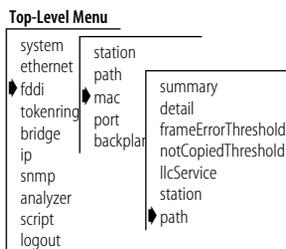
fddi mac path

You are prompted for a path assignment for the MAC.

2 Enter the path.

You are prompted for a slot related to the MAC.

3 Enter the slot(s).



You are prompted for a path assignment for each MAC in the slot(s) you specified.

4 Enter the station assignment. The possible values are 1, 2, or 3.

See the following example for single station mode path assignments:

```
Select slot(s) (1-2,4|all) [1-2,4]: all
Select MAC(s) (1-3|all) [1-3]: all
Slot 1 MAC 1 - Select path (isol,pri,sec,loc) [pri]:
Slot 1 MAC 2 - Select path (isol,pri,sec,loc) [sec]: isol
Slot 1 MAC 3 - Select path (isol,pri,sec,loc) [loc]:
Slot 2 MAC 1 - Select path (isol,pri,sec,loc) [isol]:
Slot 4 MAC 1 - Select path (isol,pri,sec,loc) [pri]:
Slot 4 MAC 2 - Select path (isol,pri,sec,loc) [isol]:loc
```

See the following example for multi-station mode path assignments:

```
Select slot(s) (1-2,4|all) [1-2,4]: all
Select MAC(s) (1-3|all) [1-3]: all
Slot 1 MAC 1 - Select path (isol,pri) [pri]:
Slot 1 MAC 2 - Select path (isol,pri) [pri]: isol
Slot 1 MAC 3 - Select path (isol,pri) [pri]: isol
Slot 2 MAC 1 - Select path (isol,pri) [pri]:
Slot 4 MAC 1 - Select path (isol,pri) [pri]: isol
Slot 4 MAC 2 - Select path (isol,pri) [isol]: pri
```

Administering FDDI Ports

Within an FDDI station, the PHY and PMD entities make up a port. A port (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:

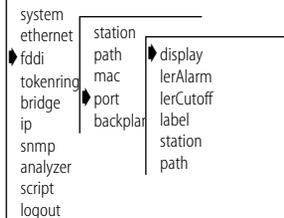
- lerAlarm
- lerCutoff
- port labels
- port paths

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:

Top-Level Menu



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

fddi port display

You are prompted for a slot related to the FDDI port.

- 2 Enter the slot(s).

You are prompted for a port.

- 3 Enter the port about which you want to view information.

See the following example of port information:

slot	port			portLabel	lemCount
1	1			Backbone1	0
1	2			SrvrRm001	0
slot	port	lerEstimate	lerAlarm	lerCutoff	lerCondition
1	1	12	7	4	inactive
1	2	12	7	4	inactive
slot	port	lemRejectCount	lctFailCount	ebErrorCount	ebErrorCond
1	1	0	0	0	inactive
1	2	0	0	0	inactive
slot	port	lineState	currentPath	connectState	pcmState
1	1	qls	isolated	connecting	connect
1	2	qls	isolated	connecting	connect
slot	port	pcWithhold	myType	neighborType	pmdClass
1	1	none	A	unknown	multimode
1	2	none	B	unknown	multimode

Table 9-6 describes the type of information provided for an FDDI port.

Table 9-6 Description of Fields for FDDI Port Attributes

Field	Description
connectState	Connect state of this port (disabled, connecting, standby, or active)
currentPath	Path on which this port is currently located
ebErrorCond	Condition is active when an elasticity buffer error has been detected during the past 2 seconds
ebErrorCount	Number of Elasticity Buffer errors that have been detected
lctFailCount	Number of consecutive times the link confidence test (LCT) has failed during connection management
lemCount	Number of link errors detected by this port
lemRejectCount	Number of times that the link error monitor rejected the link
lerAlarm	The link error rate estimate at which a link connection generates an alarm
lerCondition	Condition is active when the lerEstimate is less than or equal to lerAlarm
lerCutoff	The link error rate estimate at which a link connection is broken
lerEstimate	Average link error rate. It ranges from 10^{-4} to 10^{-15} and is reported as the absolute value of the exponent of the link error estimate
lineState	Line state of this port

(continued)

Table 9-6 Description of Fields for FDDI Port Attributes (continued)

Field	Description
myType	Type of port connector on the port
neighborType	Type of port connector at the other end of the physical connection
pcmState	Current Physical Connection Management (PCM) State defined in SMT
pcWithhold	Reason for withholding the connection
pmdClass	Type of PMD entity associated with this port
portLabel	32-character string containing a user-defined name

Setting lerAlarm

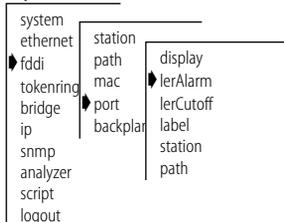
The *lerAlarm* attribute is the link error rate (LER) value at which a link connection generates an alarm. If the LER is greater than the alarm setting, then SMT sends an SRF to the network manager indicating that there is a problem with a port. The *lerAlarm* value is expressed as an exponent (such as 1×10^{-10}). A healthy network would have an LER exponent of 1×10^{-10} to 1×10^{-15} . You should set the *lerAlarm* below these values so that you are only receiving alarms if your network is in poor health. The SMT Standard recommended value is 8.



*The **lerAlarm** value must be higher than the **lerCutoff** value so that the network manager will be made aware of a problem before the PHY (port) is actually removed from the network.*

To set *lerAlarm*:

Top-Level Menu



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

fddi port lerAlarm

You are prompted for a slot related to the FDDI port.

- 2 Enter the slot(s).

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

- 3 Enter the port.

- 4 Enter the estimated link error rate value.

Valid exponent values are -4 through -15. Although these are negative exponents, you should enter the value without the negative symbol. For

example, if you want to express the value 1×10^{-8} , you would enter 8 as the value.

Setting `lerCutoff`

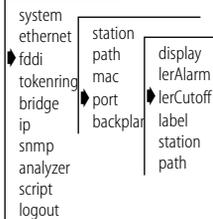
The `lerCutoff` attribute is the link error rate estimate at which a link connection is disabled. Once the `lerCutoff` value is reached, the PHY that detected a problem is disabled. The `lerCutoff` value is expressed as an exponent (such as 1×10^{-10}). A healthy network would have an LER exponent of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



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

To set the `lerCutoff`:

Top-Level Menu



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

```
fddi port lerCutoff
```

You are prompted for a slot related to the FDDI port.

- 2 Enter the slot(s).

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

- 3 Enter the port.

- 4 Enter the estimated link error rate value.

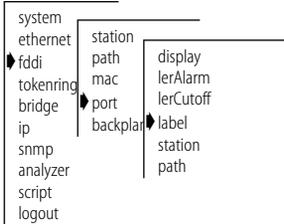
Valid exponent values are -4 through -15. Although these are negative exponents, you should enter the value without the negative symbol. For example, if you want to express the value 1×10^{-7} , you would enter 7 as the value.

Setting Port Labels

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

To label an FDDI port:

Top-Level Menu



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

fddi port label

You are prompted for a slot related to the FDDI port.

- 2 Enter the slot(s).

You are prompted for a port and a label value.

- 3 Enter the port.

- 4 Enter the label value.

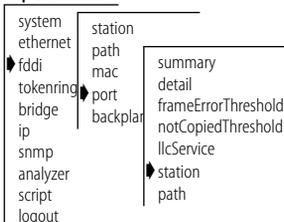
Assigning Ports to Stations in Multi-Station Mode

You can only assign ports to a new station when the FDDI backplane paths are set for multi-station mode. If you access this menu item when the FDDI backplane paths are set for single station mode, the following message appears:

The FDDI backplane path mode is set to "singleStation" so all three backplane paths are part of the same station. This menu item is only useful if the FDDI backplane path mode is set to "multiStation".

To assign ports to stations:

Top-Level Menu



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

fddi port station

You are prompted for a slot related to the port.

- 2 Enter the slot(s).

You are prompted for a station assignment for each port in the slot(s) you specified.

- 3 Enter the station assignment. The possible values are 1, 2, or 3.

If you change a station assignment, then you must reboot the system for the change to take effect.

See the following example:

```
Select slot(s) (2,4|all) [2,4]: 2
Select ports(s) (1-6|all): all
Slot 2 port 1 - Enter new FDDI station number (1-3) [1]:
Slot 2 port 2- Enter new FDDI station number (1-3) [1]: 2
You have requested that this port be moved to a different
FDDI station. You must reboot your system for this
request to take effect.
Slot 2 port 3 - Enter new FDDI station number (1-3) [1]: 3
You have requested that this port be moved to a different
FDDI station. You must reboot your system for this
request to take effect.
Slot 2 port 4 - Enter new FDDI station number (1-3) [1]:
Slot 2 port 5 - Enter new FDDI station number (1-3) [1]:
Slot 2 port 6 - Enter new FDDI station number (1-3) [1]:
```

If you change a station assignment, then you must reboot the system for the change to take effect.

Setting the Port Paths

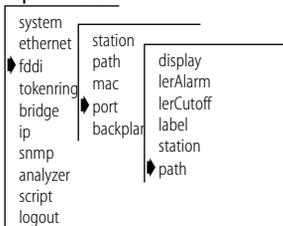
You can configure the FDDI port paths for modules with external FDDI ports: LMM+, FCM, and EFSM.

In the you can assign the A and B ports to either the primary or secondary paths.

To assign ports to paths:

- 1 From the top level of the Administration Console, enter:
fddi port path
You are prompted for a slot related to the FDDI port.
- 2 Enter the slot(s).
You are prompted for a path.
- 3 Enter the port(s) you want to configure.

Top-Level Menu



- 4 Select the DAS configuration **isol** or **thru** for peer mode at the prompt.
- 5 Select the DAS configuration **isol**, **wrap AB**, or **dualHome** for tree mode at the prompt.

The information you enter may vary depending on the type of module you are configuring:

- For an LMM+, you are prompted for the DAS configuration for both the peer and tree connections. Peer connections are formed on the trunk ring, which consists of A-B connections. These connections have no end port of type M. Tree connections connect nodes to the M ports of concentrators. In this type of connection, exactly one end of the connection is a port of type M.

For the peer connection, the possible configurations are isolated or thru. For the tree connection, the possible values are isolated (taking the port off the FDDI ring), wrapAB (the primary path is wrapped to the B port and the secondary path is wrapped to the A port), and dualHome (the primary and secondary paths are concatenated and wrapped to the B port, and the A port is in standby).

See the following example:

```
Select slot [1]:
Select port(s) (1-2|all): all
The A and B ports are configured together.
  Select DAS config for peer mode (isol,thru) [thru]:
  Select DAS config for tree mode (isol,wrapAB,dualHome)
  [isol]:
```

- For an FCM, you are prompted for the path: isolated, primary, secondary, or local in single station mode, and isolated and primary in multi-station mode.

See the following example for single station mode path assignments:

```
Select slot(s) (2,4|all) [2]: 4
Select port(s) (1-6|all) [1-3]: all
Slot 2 port 1 - Select path (isol,pri,sec,loc) [pri]:
Slot 2 port 2 - Select path (isol,pri,sec,loc) [sec]: isol
Slot 2 port 3 - Select path (isol,pri,sec,loc) [loc]:
Slot 2 port 4 - Select path (isol,pri,sec,loc) [isol]:
Slot 2 port 5 - Select path (isol,pri,sec,loc) [pri]:
Slot 2 port 6 - Select path (isol,pri,sec,loc) [isol]: loc
```

See the following example for multi-station mode path assignments:

```
Select slot(s) (2,4|all) [2]: 4
Select port(s) (1-6|all) [1-3]: all
Slot 2 port 1 - Select path (isol,pri) [pri]:
Slot 2 port 2 - Select path (isol,pri) [pri]: isol
Slot 2 port 3 - Select path (isol,pri) [pri]: isol
Slot 2 port 4 - Select path (isol,pri) [pri]:
Slot 2 port 5 - Select path (isol,pri) [pri]: isol
Slot 2 port 6 - Select path (isol,pri) [isol]: pri
```


10

ADMINISTERING TOKEN RING PORTS

This chapter describes how to:

- View Token Ring port information
- Configure Token Ring port labels
- Enable or disable a Token Ring port
- Configure Token Ring port speed
- Configure Token Ring port mode

Displaying Token Ring Port Information

You can display either a summary of Token Ring port information or a detailed report. When you display a summary of Token Ring port information, you receive information about the port, including its label, status (on-line, off-line, etc.), and the most pertinent statistics about general port activity and port errors.

The detailed display of Token Ring port information includes the information in the summary and additional Token Ring port statistics, such as the number of beacon errors.

To display information about the Token Ring ports:

Top-Level Menu

```

system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
  
```

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

```
tokenring summary
```

OR

```
tokenring detail
```

- 2 Select the slot(s) related to the port(s) about which you want to view information.
- 3 Enter the port(s) for which you want to view information.

The port information is displayed in the format you specified. Table 10-1 describes the type of information provided about a Token Ring port.

An example of a summary display for Token Ring ports is shown below:

port	portLabel	portState
1	Server1	on-line
8	Office322_PC	on-line

port	rxFrames	txFrames	rxBytes	txBytes
1	406876	1423733	36377226	234900612
8	242532	1257721	29293858	300479754

port	rxErrs	txErrs	noRxBuffers	txQOverflows
1	0	0	0	n/a
8	0	0	0	n/a

The following example shows a detailed display for Token Ring ports on a TRSM:

```

      rxFrames      rxBytes      rxFrameRate      rxByteRate
          875          65806              0              0
rxPeakFrameRate rxPeakByteRate  noRxBuffers  rxInternalErrs
          114          9741              0              0
      rxUnicasts    rxMulticasts    rxDiscards      fcsErrs
          869           6              0              0
      txFrames      txBytes      txFrameRate      txByteRate
          7535         2485537          0              0
txPeakFrameRate txPeakByteRate  rxLongFrames  txInternalErrs
          158          302241          0              0
      txUnicasts    txMulticasts    txDiscards      txBeacons
          469           6218              0              0
      ringPurges    remRingStation    lineErrs      internalErrs
          1              0              0              0
      burstErrs      a/cErrs    abortDelimErrs  lostFrameErrs
          1              0              0              0
rxCongestion frameCopiedErrs    freqErrs      tokenErrs
          0              0              0              0
txQOverflows requestedState    portState      portMode
          n/a          enabled          on-line      station
insertStatus      macAddress      portSpeed
          inserted      00-01-7c-c8-37-cc      4Mbps
                                portLabel
    
```

Table 10-1 Description of Fields for Token Ring Port Attributes

Field	Description
abortDelimErr	Number of times an abnormally terminated frame was seen
a/cErr	Number times that more than one AMP or SMP frame was seen without the A/C field set.
burstErrs	Number of times a burst error was detected by this station. A burst error is the reception of five consecutive manchester symbols without a transition.
fcsErrs	Number of frames received by this port that are an integral number of octets in length but do not pass the FCS check
frameCopiedErrs	Number of frames sent to this station's specific address which had been previously recognized and accepted by another station
freqErrs	Number of frequency errors detected by this station
internalErrs	Number of errors encountered at this port
insertStatus	Insert or de-insert into a ring
lineErrs	Number of frames or tokens received with a manchester code violation detected
lostFrameErrs	Number of frames sent by this station that did not return to be stripped
macAddress	The MAC address of this port
noRxBuffers	Number of frames discarded because there was no available buffer space
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.
portMode	Specific description of this port mode (station or lobe).
remRingStation	Number of Remove Ring Station frames received, or ignored, by this station
requestedState	Configurable parameter used to enable/disable this port. The default is enabled.
ringPurges	Number of ring purges seen on ring
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 delimiters
rxCongestion	Number of times this port was unable to receive frames due to a lack of buffers

(continued)

Table 10-1 Description of Fields for Token Ring Port Attributes (continued)

Field	Description
rxDiscards	Number of frames discarded during reception
rxErrs	Sum of all receive errors associated with this port (field only appears in the summary option)
rxLongFrame	Number of frames longer than 4500 bytes received by this port
rxFrameRate	Average number of frames received per second by this port during the most recent sampling period
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 (non-multicast) frames delivered by this port to a higher-level protocol or application
tokenErrs	Number of tokens sent by this station that did not return to be stripped
txbeacons	Number of beacons transmitted by this port
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 frames discarded because the port was disabled
txErrs	Sum of all transmit errors associated with this port (field only appears in the summary option)
txFrameRate	Average number of frames transmitted per second by this port during the most recent sampling period
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

(continued)

Table 10-1 Description of Fields for Token Ring Port Attributes (continued)

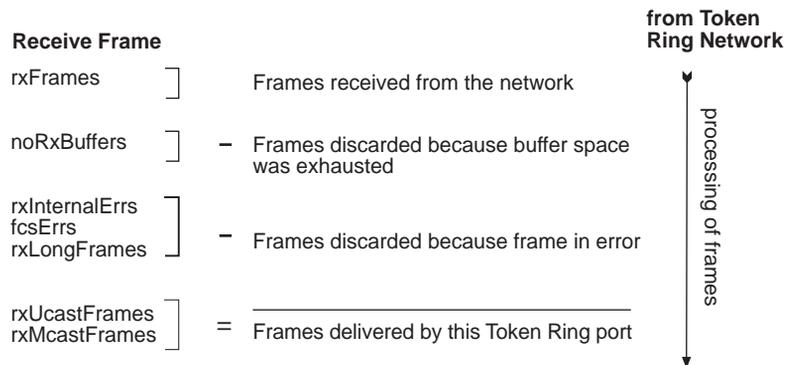
Field	Description
txOverflows	Not applicable
txPeakByteRate	Peak value of ethernetPortByteTransmitRate for this port since the station was last initialized
txPeakFrameRate	Peak value of ethernetPortFrameTransmitRate for this port since the station was last initialized
txUnicasts	Number of unicast (non-multicast) frames queued for transmission by a higher-level protocol or application, including frames not transmitted successfully

Frame Processing and Token Ring Statistics

All frames on the Token Ring network are received promiscuously by a Token Ring port. A frame may be discarded, however, for the following reasons:

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

Figure 10-1 shows the order in which these discard tests are made.

**Figure 10-1** How Frame Processing Affects Token Ring Receive Frame Statistics

Frames are delivered to a Token Ring port by bridges, routers, and management applications. However, a frame may be discarded for the following reasons:

- The Token Ring port is disabled
- There is no room on the transmit queue
- An error occurred during frame transmission
- The frame is too long

Figure 10-2 shows the order in which these discard tests are made.

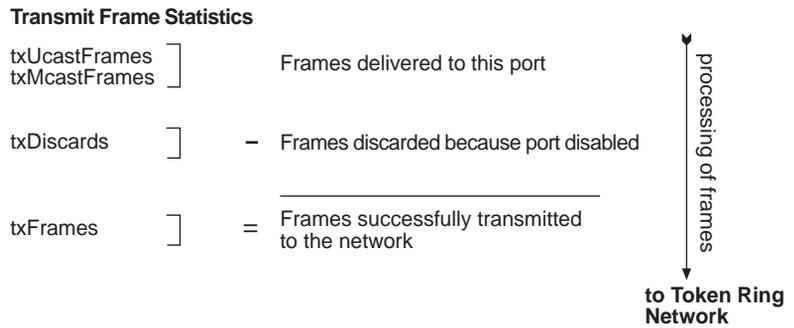


Figure 10-2 How Frame Processing Affects Token Ring Transmit Frame Statistics

Labeling a Port

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

To label a Token Ring port:

Top-Level Menu

```
system
ethernet
fddi
tokenring
bridge
ip
snmp
analyzer
script
```

```
summary
detail
label
portState
portSpeed
portMode
```

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

tokenring label

- 2 Select the slot(s) related to the port(s) you want to label.
- 3 Enter the port(s) you want to label.
- 4 Enter the label of each Token Ring port.

The new port label appears next time you display information for that port.

Setting the Port State

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

To enable or disable a Token Ring port:

Top-Level Menu

```
system
ethernet
fddi
tokenring
bridge
ip
snmp
analyzer
script
```

```
summary
detail
label
portState
portSpeed
portMode
```

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

tokenring portState

- 2 Enter the module(s) for which you want to set port states.
- 3 Enter the number(s) of the port(s) you want to set.
- 4 Enter **enable** or **disable** for each Token Ring port.

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

Setting the Port Speed

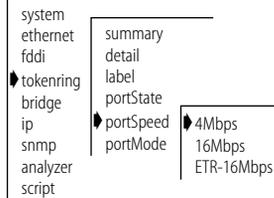
You may set the port speed for any Token Ring port to 4 Mbps, 16Mbps, or 16Mbps Early Token Release. The default port speed for all Token Ring ports is *16 Mbps*.

To set the port speed for a Token Ring port:

- 1 From the top level of the Administration Console, enter:
tokenring portSpeed
- 2 Enter the slot number(s) of the TRSM(s) for which you want to set port speed.
- 3 Enter the number(s) of the port(s) you want to set.
- 4 Enter **4Mbps**, **16Mbps**, or **ETR-16Mbps** for each Token Ring port.

The port speed value (shown in the summary and detail displays) shows on-line for all enabled ports displayed and off-line for all disabled ports displayed.

Top-Level Menu



Setting the Port Mode

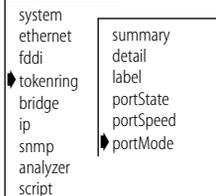
Token Ring ports normally operate as 802.5 compliant stations which are connected to the network through lobe ports on a MAU (Multistation Access Unit). You can set the port mode for ports 1 and 2 to operate as lobe ports for direct attachment of Token Ring stations. The default is *station*.

To set a Token Ring port to lobe mode:

- 1 From the top level of the Administration Console, enter:
tokenring portMode
- 2 Enter the module(s) for which you want to set port mode.
- 3 Enter the number(s) of the port(s) you want to set.
- 4 Enter **station** or **lobe** for each Token Ring port.

The port mode value (shown in the summary and detail displays) reflects the port's current mode.

Top-Level Menu





11

SETTING UP THE SYSTEM FOR ROVING ANALYSIS

This chapter describes how to set up the 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 selectively choose any Ethernet network segment attached to a LANplex 6000 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 port are attached to the same LANplex option module) or remotely (when the analyzer and port are on different modules in the same system or in two separate systems)

You may want to monitor a port to:

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

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

To enable the monitoring of ports on a LANplex, you must do the following:

- 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 or module (ESM

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

Figure 11-2 shows local and remote monitoring on a single LANplex system. Figure 11-3 shows an analyzer configured on one LANplex system and the monitored port configured on a different LANplex system.

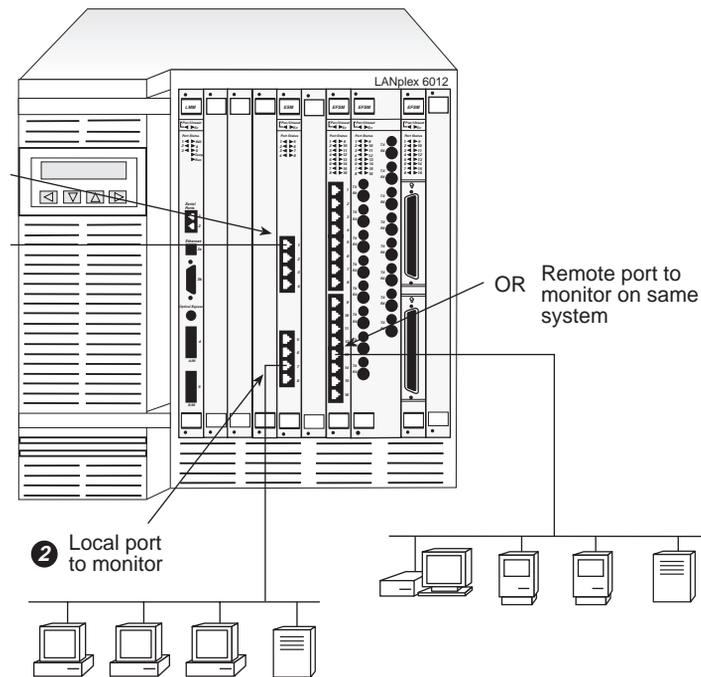


Figure 11-1 Monitoring Locally and Remotely on the Same LANplex System

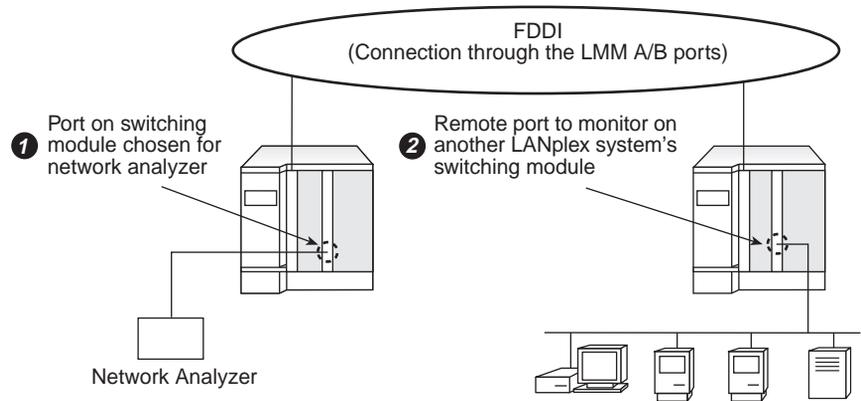


Figure 11-2 Monitoring Remotely Between Two LANplex Systems

Configuration rules You can have a maximum of 16 network analyzers connected to a system and the following maximum number of ports monitored per module:

- ESM — one port monitored
- EFSM — eight ports monitored

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 and from there monitor one Ethernet port at a time.

Displaying the Roving Analysis Configuration

You may want to display the Roving Analysis configuration to see which ports on a module are designated as analyzer ports and which ports are currently being monitored on a specific module.

When you display the Roving Analysis configurations for a system, you receive:

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

To display the Roving Analysis configurations:

- 1 From the top level of the Administration Console, enter:
analyzer display
- 2 Enter the slot number(s) of the module(s) for which you want to display the Roving Analysis configuration.

The configurations are displayed as shown in the following example:

```
Ethernet ports configured as analyzers:
  Ethernet port      Ethernet address
    9                 00-80-3e-0a-3b-02
Ethernet ports being monitored:
  Ethernet port      Ethernet address
    16                00-80-3e-0a-3b-02
```

Top-Level Menu

```
system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout
```

Adding an Analyzer Port

You can have up to 16 network analyzers connected to a system. For a more accurate analysis, you should 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 Enter the slot number of the module to which the network analyzer is attached.
- 3 Enter the number of the Ethernet port to which the network analyzer is attached.

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

```
Select slot(s) (1-4,8-10): 9
Select Ethernet port (1-16): 9
Analyzer port address is 00-80-3e-0a-3b-02
```

Top-Level Menu

```
system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout
```

Port selection errors

If your port selection is not valid, you will 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 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 the address can be learned on remote systems.



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

Removing an Analyzer Port

You may want to 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 Enter the slot number to which the network analyzer is attached.
- 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 prior to setting it as an analyzer port.

Top-Level Menu

```
system
ethernet
fdi
tokenring
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 Enter the slot number of the module on which the port is to be monitored.
- 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 where 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.

See the example below for starting port monitoring:

```
Select slot(s) (1-4,8-10): 4
Select Ethernet port (1-16): 16
Enter the address the Analyzer is located on - type 'q' to
return to the previous menu
Address: 00-80-3e-0a-3b-02
```

Port selection errors

If your port selection is not valid, you will 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 for EFSMs

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

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

Top-Level Menu

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

display
add
remove
start
stop
```



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.*

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

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

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

Stopping Port Monitoring

After analyzing an Ethernet port, you may want to 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 Enter the slot number of the module on which the port is monitored.
- 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
fddi
tokenring
bridge
ip
snmp
▼ analyzer
script
logout
```



IV

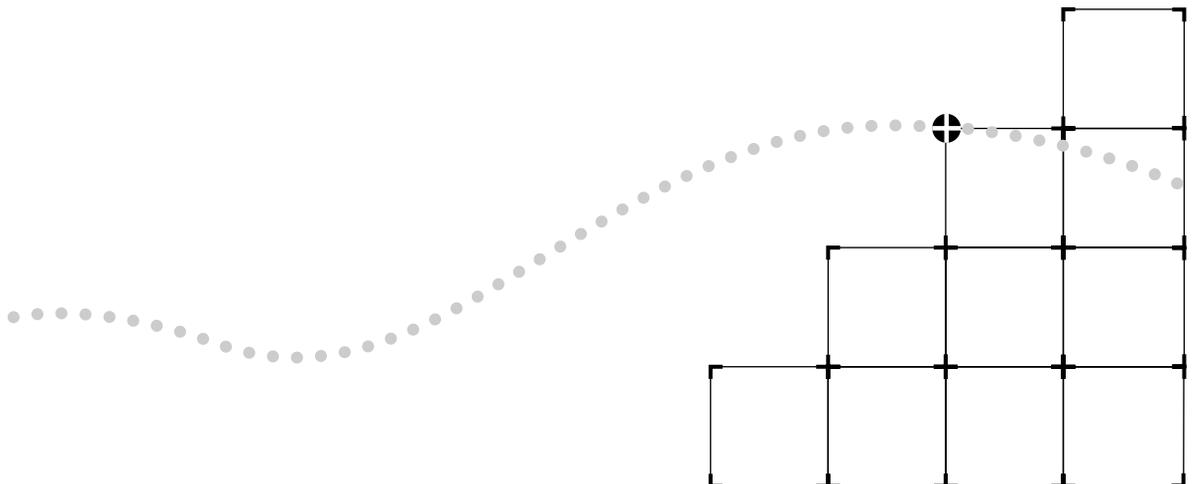
BRIDGING

Chapter 1012 Administering the Bridge

Chapter 1113 Administering Bridge Ports

Chapter 1214 Creating and Using Packet Filters

Chapter 1315 Configuring Address and Port Groups to Use in Packet Filters



12

ADMINISTERING THE BRIDGE

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

- Bridging mode (Transparent or Express Switching)
- IP fragmentation
- IPX snap translation
- Address threshold
- Address aging time
- Spanning Tree Protocol (STP) parameters

For information about configuring the bridge port, see Chapter 13. For information about creating packet filters for a bridge, see Chapter 14.

Displaying Bridge Information

You can display information about the bridge, which includes bridge statistics (such as topology change information) and configurations for the bridge and Spanning Tree. Each Ethernet Switching Module (ESM), Ethernet/FDDI Switching Module (EFSM), and Token Ring Switching Module (TRSM) is considered a bridge.

Top-Level Menu

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

To display bridge information:

- 1 From the top level of the Administration Console, enter:
bridge display
- 2 Enter the slot number of the bridge for which you want to display information.

Information about the bridge is displayed.

The following example shows a display of bridge information.

```

      stpState                timeSinceLastTopologyChange
      enabled                  1 hr 28 mins 31 secs

      topologyChangeCount    topologyChangeFlag
      2                       false

      BridgeIdentifier        designatedRoot
      8000 00803e003040      7fff 00803e026294

      bridgeMaxAge            maxAge bridgeHelloTime    helloTime
      20                      20      2                2

      bridgeFwdDelay          forwardDelay    holdTime        rootCost
      15                      15        1                220

      rootPort                priority        agingTime       mode
      1                       0x8000    300              802.1d

      addrTableSize           addressCount    peakAddrCount   addrThreshold
      8192                    55          105              8000

      ipFragmentation
      enabled

```

Each item in the bridge parameter list is described in Table 12-1.



Source Route (SR) and Source Route/Transparent (SRT) modes are only available on the LANplex TRSM. The default mode for all bridging modules is IEEE 802.1d Transparent operation.

Table 12-1 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.

(continued)

Table 12-1 Bridge Attributes (continued)

Parameter	Description
agingTime	Time-out period in seconds (between 10 and 32267) for aging out dynamically learned forwarding information. The default value is 300 seconds (or 5 minutes).
bridgeFwdDelay	Forward delay value used when this bridge is the root bridge. This value sets the amount of time a bridge spends in the “listening” and “learning” states. The default value is 15 seconds.
bridgeHelloTime	Hello time value used when this bridge is the root bridge. This value is the time that elapses between the generation of configuration messages by a bridge that assumes itself to be the root. The default value is 2 seconds.
bridgeIdentifier	Bridge identification. It includes the bridge priority value and the MAC address of the lowest numbered port (for example: 8000 00803e003dc0).
bridgeMaxAge	Maximum age value used when this bridge is the root bridge. This value determines when the stored configuration message information is too old and is discarded. The default value is 20 seconds.
designatedRoot	Root bridge identification. It includes the root bridge’s priority value and the MAC address of the lowest numbered port on that bridge (for example: 8000 00803e001520).
forwardDelay	The time a bridge spends in the “listening” and “learning” states.
helloTime	The time that elapses between the generation of configuration messages by a bridge that assumes itself to be the root.
holdTime	Minimum delay time between sending BPDUs (topology change Bridge Notification Protocol Data Units)
ipFragmentation	Configurable parameter that controls whether IP fragmentation is enabled or disabled. The default value is enabled.
ipxSnapTranslation	Configurable parameter that controls whether IPX 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, <i>express</i> for Express Switching, <i>SR</i> for Source Routing only, or <i>SRT</i> for 802.1d Source Route/Transparent.

(continued)

Table 12-1 Bridge Attributes (continued)

Parameter	Description
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
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 last reconfigured the network topology
topologyChange-Flag	Indicates whether a topology change is currently occurring on the bridge. A value of <i>true</i> means that it is. A value of <i>false</i> means that no topology change is occurring.
topologyChange-Count	Number of times Spanning Tree has reconfigured the network topology
trFddiMode	Indicates whether embedded protocol addresses are being translated between Token Ring and FDDI ports. <i>Native</i> mode indicates that translations occurring while <i>backbone</i> indicates that translation is off.

Setting the Bridging Mode

You can run a bridge (ESM and EFSM) in either Transparent bridging mode or in Express Switching mode. You can run a Token Ring (TRSM) bridge in Transparent, Source Route/Transparent (SRT) or Source Route (SR) modes. The advantages and disadvantages of using Express Switching are defined in Chapter 6: *Express Switching* or Chapter 12: *Token Ring in the LANplex System* in the *LANplex 6000 Operation Guide*. The advantages and disadvantages of using Transparent, Source Route/Transparent or Source Route modes are defined in Chapter 12: *Token Ring in the LANplex System* in the *LANplex 6000 Operation Guide*.

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



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

You can mix bridging modes within your system.

To set the bridging mode for the ESM, EFSM, or TRSM:

Top-Level Menu

system	display
ethernet	mode
fddi	ipFragmentation
tokenring	ipxSnapTranslation
bridge	trFDLMode
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	srBridgeNumber
	port
	packetFilter

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

bridge mode

- 2 Enter the number(s) of the slot(s) for the bridges in the system.

- a If the slot number you entered contains an ESM or EFSM you are prompted for a mode value as shown below:

Enter new value (express,transparent) [transparent]:

- b If the slot number you entered contains a TRSM you are prompted for a mode value as shown below:

Enter new value (transparent,srt,sr) [transparent]:

- 3 To turn on Express Switching mode for the ESM or EFSM, enter:

express

To turn on transparent bridging mode for the ESM or EFSM, enter:

transparent

- 4 To turn on Source Route/Transparent mode for the TRSM, enter:

srt

To turn on Source Route bridging mode for the TRSM, enter:

sr

To turn on transparent bridging mode for the TRSM, enter:

transparent

If you are enabling Express Switching, you are prompted for a backbone port through which to forward the packets received on Ethernet ports.

- 5 Enter the backbone port type (Ethernet or FDDI).
- 6 Enter the port number through which Express packets will be forwarded.

Enabling/ Disabling IP Fragmentation

When IP fragmentation is enabled, large FDDI packets are “fragmented” into smaller packets. This allows FDDI and Ethernet stations connected to the LANplex 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*.



This function is not available for Token Ring bridges because Token Ring and FDDI maximum packet sizes are equal.

To enable or disable IP fragmentation for a bridge:

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

bridge ipfragmentation

You are prompted for slots.

- 2 Enter the number(s) of the slot(s) or **all** for all bridges in the system.

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

enable

To disable IP fragmentation on a bridge, enter:

disable

Top-Level Menu

system	display
ethernet	mode
fddi	ipFragmentation
tokenring	ipxSnapTranslation
bridge	trFDDIMode
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	srBridgeNumber
	port
	packetFilter

Enabling/ 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*.

Top-Level Menu

system	display
ethernet	mode
fddi	ipFragmentation
tokenring	ipxSnapTranslation
bridge	trFDDIMode
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	srBridgeNumber
	port
	packetFilter

To enable or disable IPX Snap Translation for a bridge:

- 1 From the top level of the Administration Console, enter:
bridge ipxSnapTranslation
- 2 Enter the number(s) of the slot(s) or **all** for all bridges in the system.
- 3 To enable IPX Snap Translation on a bridge, enter:

enable

To disable IP Snap Translation on a bridge, enter:

disable

Setting Protocol Address Translation from Token Ring to FDDI

Embedded protocol addresses for Token Ring use the non-canonical format while FDDI uses the canonical format. To provide interoperability, these embedded addresses can be found and translated to their native formats on TRSM for IP, IPX, NetBIOS and SNA protocols. It is necessary to do this for a client on Token Ring to talk to a server on FDDI. It is unnecessary if FDDI is used purely as a transport backbone. The default value for this parameter is *native address conversion*.



This parameter is only available on Token Ring bridges.

To configure bridge protocol translation:

- 1 From the top level of the Administration Console, enter:
bridge trFddiMode
- 2 Enter the number(s) of the slot(s) or **all** for all bridges in the system.
- 3 Enter the protocol translation value (**native** or **backbone**).

Top-Level Menu

system	display
ethernet	mode
fddi	ipFragmentation
tokenring	ipxSnapTranslation
bridge	trFddiMode
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	srBridgeNumber
	port
	packetFilter

Setting the Address Threshold

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

Address threshold values

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* since the limit will never be reached. The default value is 8000.



This value may be exceeded if you change the bridging mode.

To set the address threshold:

- 1 From the top level of the Administration Console, enter:
bridge addressThreshold
- 2 Enter the number(s) of the slot(s) or **all** for all bridges in the system.
- 3 Enter the value of the threshold.

Top-Level Menu

system	display
ethernet	mode
fddi	ipFragmentation
tokenring	ipxSnapTranslation
bridge	trFDDIMode
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	srBridgeNumber
	port
	packetFilter

Setting the Aging Time

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

Aging time values

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

Top-Level Menu

```

system
ethernet
fddi
tokenring
bridge
ip
snmp
analyzer
script
logout
display
mode
ipFragmentation
ipxSnapTranslation
trFDDIMode
addressThreshold
agingTime
stpState
stpPriority
stpMaxAge
stpHelloTime
stpForwardDelay
srBridgeNumber
port
packetFilter

```

To set the bridge aging time:

- 1 From the top level of the Administration Console, enter:
bridge agingTime
- 2 Enter the number(s) of the slot(s) or **all** for all bridges in the system.
- 3 Enter the aging time value.

Administering STP Bridge Parameters

You can enable or disable Spanning Tree on one or more bridges and set the following STP bridge parameters: priority, maximum age, hello time, and forward delay. For more information about how the Spanning Tree parameters interact at the bridge-level to create a loopless network, see Chapter 5: *Transparent Bridging* in the *LANplex 6000 Operation Guide*.

Enabling/Disabling STP on a Bridge

You can enable and disable Spanning Tree on any bridge in the system. When a bridge is disabled, it does not participate in the Spanning Tree algorithm.

Default value The default value is *disabled*.

To enable or disable Spanning Tree on one or more bridges:

- 1 From the top level of the Administration Console, enter:
bridge stpState
- 2 Enter the number(s) of the slot(s) or **all** for all bridges in the system.
- 3 Enter **enabled** or **disabled** at the prompt.

Top-Level Menu

```

system
ethernet
fddi
tokenring
bridge
ip
snmp
analyzer
script
logout
display
mode
ipFragmentation
ipxSnapTranslation
trFDDIMode
addressThreshold
agingTime
stpState
stpPriority
stpMaxAge
stpHelloTime
stpForwardDelay
srBridgeNumber
port
packetFilter

```

Setting the Bridge Priority

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

Bridge priority values

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

To configure the STP bridge priority:

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

```
bridge stpPriority
```

- 2 Enter the number(s) of the slot(s) or **all** if you want to configure the priority of all bridges in the system.

- 3 Enter the priority value for each bridge you specified at the prompt.

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

Top-Level Menu

```
system
ethernet
fdi
tokenring
bridge
ip
snmp
analyzer
script
logout
display
mode
ipFragmentation
ipxSnapTranslation
trfDDIMode
addressThreshold
agingTime
stpState
stpPriority
stpMaxAge
stpHelloTime
stpForwardDelay
srBridgeNumber
port
packetFilter
```

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 Spanning Tree is configured properly, the maximum age value should ideally never be reached. If the value is too small, then Spanning Tree 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.

Top-Level Menu

system	display
ethernet	mode
fddi	ipFragmentation
tokenring	ipxSnapTranslation
▶ bridge	trFDDIMode
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	▶ stpMaxAge
	stpHelloTime
	stpForwardDelay
	srBridgeNumber
	port
	packetFilter

To configure the bridge max age:

- 1 From the top level of the Administration Console, enter:
bridge stpMaxAge
- 2 Enter the number(s) of the slot(s) or **all** to configure the max age of all bridges in the system.
- 3 Enter the bridge max age value for each bridge you selected.

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 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:

- 1 From the top level of the Administration Console, enter:
bridge stpHelloTime
- 2 Enter the number(s) of the slot(s) or **all** if you want to configure hello time of all bridges in the system.
- 3 Enter the bridge hello time value for each bridge you selected.

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

Top-Level Menu

system	display
ethernet	mode
fddi	ipFragmentation
tokenring	ipxSnapTranslation
▶ bridge	trFDDIMode
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	▶ stpHelloTime
	stpForwardDelay
	srBridgeNumber
	port
	packetFilter

Setting the Bridge Forward Delay

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

Forward delay recommended value

The recommended value is 15 seconds.

To configure the forward delay value:

Top-Level Menu

system	display
ethernet	mode
fddi	ipFragmentation
tokenring	ipxSnapTranslation
bridge	trFDDIMode
ip	addressThreshold
snmp	agingTime
analyzer	stpState
script	stpPriority
logout	stpMaxAge
	stpHelloTime
	stpForwardDelay
	srBridgeNumber
	port
	packetFilter

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

```
bridge stpForwardDelay
```

- 2 Enter the number(s) of the slot(s) or **all** to configure forward delay for all bridges in the system.
- 3 Enter the forward delay value for each bridge you selected.

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.

13

ADMINISTERING BRIDGE PORTS

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

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

Displaying Bridge Port Information

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

To display bridge information:

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

```
bridge port summary
```

OR

```
bridge port detail
```

You are prompted for slot number(s).

- 2 Enter the number(s) of the slot(s) or **all** to view port parameters for all bridges in the system.

You are prompted for the port type.

- 3 Enter **Ethernet**, **FDDI**, **tokenring**, or **all**.

You are prompted for port number(s).

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

Top-Level Menu

```
system
ethernet
fdi
tokenring
bridge
ip
snmp
analyze
script
logout
display
mode
ipFragm
ipxSnap
trFddiM
address
agingTim
stpState
stpPriori
stpMaxA
stpHelloTime
stpForwardDelay
srBridgeNumber
port
packetFilter
summary
detail
multicastLimit
stpState
stpCost
stpPriority
srRingNumber
address
```

The following example shows a bridge port summary display for an EFSM.

	port	rxFrames	rxDiscards	txFrames
	FDDI 1	1680326	1095	654715
	Ethernet 1	411180	0	1353766
	Ethernet 12	243559	0	1184225

	port	portId	stp	state	fwdTransitions
	FDDI 1	0x8001	enabled	forwarding	1
	Ethernet 1	0x8003	enabled	forwarding	1
	Ethernet 12	0x800e	enabled	forwarding	1

The following example shows a bridge port detail display for an EFSM.

	port	rxFrames	rxAllFilters	rxMcastFilters	rxSameSegDiscs
	FDDI 1	1685143	0	0	0
	Ethernet 1	412404	0	0	0
	Ethernet 12	243932	0	0	0

	port	rxBlockedDiscs	rxErrorDiscs	rxSecurityDiscs	rxOtherDiscs
	FDDI 1	1095	0	0	0
	Ethernet 1	0	0	0	0
	Ethernet 12	0	0	0	0

	port	rxForwardUcasts	rxFloodUcasts	rxForwardMcasts	txBlockedDiscs
	FDDI 1	0	0	0	0
	Ethernet 1	0	0	0	0
	Ethernet 12	0	0	0	0

	port	txMtuExcDiscs	txAllFilters	txMcastFilters	txFrames
	FDDI 1	0	0	0	656312
	Ethernet 1	0	0	0	1357939
	Ethernet 12	0	0	0	1187369

	port	portId	stp	state	fwdTransitions
	FDDI 1	0x8001	enabled	forwarding	1
	Ethernet 1	0x8003	enabled	forwarding	1
	Ethernet 12	0x800e	enabled	forwarding	1

	port	priority	pathCost	designatedCost	designatedPort
	FDDI 1	0x80	10	0	0x8001
	Ethernet 1	0x80	100	10	0x8003
	Ethernet 12	0x80	100	10	0x800e

	port	designatedRoot	designatedBridge
	FDDI 1	7fff 00803e028e02	7fff 00803e028e02
	Ethernet 1	7fff 00803e028e02	8000 00803e0b4800
	Ethernet 12	7fff 00803e028e02	8000 00803e0b4800

Table 13-1 describes the type of information provided for the bridge port.

Table 13-1 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 would be 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
filtered	Total number of frames discarded by this port due to user-defined packet filters
forwarded	Total number of frames forwarded from this port to all other ports of this bridge. The total forwarded count is NOT equal to the receive count minus the discard count, as a single received multicast/broadcast frame can be forwarded to multiple bridge ports. Also, frames that are flooded because of unknown destination addresses account for multiple forwarding. The total forwarding count is the instantaneous sum of the "forwarded to counts" to all other ports of the bridge.
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 means the more unstable the topology.
pathCost	Cost to be added to the total path cost when this port is the root port
port	Either Ethernet or FDDI (maximum possible count for EFSM: 1 = FDDI and 2 – 16 = Ethernet)
portId	Identification of the port, which includes the port priority and the port number (for example: 8002)
priority	First factor to determine if a port is to be the designated port when more than one bridge port is attached to the same LAN. If all ports in a bridge have the same priority, then the port number is used as the determining factor.
rxAllFilters	Number of frames discarded because of a user-defined packet filter on the receive all path of this bridge port
rxBlockedDiscs	Number of frames discarded by this port because the receiving bridge port was not in the forwarding state

(continued)

Table 13-1 Bridge Port Attributes (continued)

Parameter	Description
rxErrorDiscs	Number of frames discarded by this port because of internal bridge system errors (such as hardware and software address table discrepancies)
rxFloodUcasts	Number of unicast frames received on this port that were flooded to one or more ports
rxForwardMcasts	Number of multicast frames received on this bridge port
rxForwardUcasts	Number of unicast frames received on this bridge port
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.
rxMcastFilters	Number of frames discarded because of a user-defined packet filter on the receive multicast path of this port
rxOtherDiscs	Number of frames discarded by this port because they contained either invalid (group) source addresses or source addresses belonging to this bridge (indicative of 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)
srRingNumber	Configurable parameter which defines the ring number to which this port is attached. This parameter along with srBridgeNumber is used to determine forwarding of source routed traffic.

(continued)

Table 13-1 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 Spanning Tree 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 or not the port is <i>enabled</i> or <i>disabled</i> for Spanning Tree.
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, FDDI, or Token Ring) 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 may 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 13-1 shows the order in which the discard decisions are made.

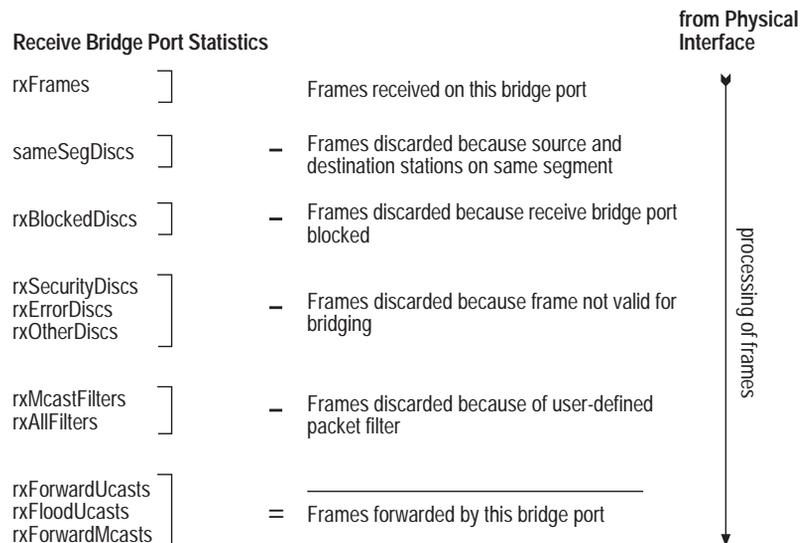


Figure 13-1 How Frame Processing Affects Receive Bridge Port Statistics

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

- The transmit bridge port is blocked
- The frame is too large for the corresponding physical interface

- 3 Enter **Ethernet**, **FDDI**, **tokenring**, or **all**.
You are prompted for port number(s).
- 4 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.
- 5 Enter the new multicast threshold value for the port(s).
See the example below:

```
Slot 3 Ethernet port 4 - Enter new value [0]: 400
Slot 3 Ethernet port 5 - Enter new value [0]: 400
```

Administering STP Bridge Port Parameters

You can enable or disable Spanning Tree for one or more ports on the system. This only affects the operation of the port if Spanning Tree 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 6000 Operation Guide*.

Enabling/Disabling STP on a Port

You can enable and disable Spanning Tree for any port on any switching module in the system. When disabled, a port does not forward frames or participate in the Spanning Tree algorithm.

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	display	summary
ethernet	mode	detail
fddi	ipFragm	multicastLimit
tokenring	ipxSnap	stpState
bridge	trFddiM	stpCost
ip	address	stpPriority
snmp	agingTim	srRingNumber
analyze	stpState	address
script	stpPrior	
logout	stpMaxAge	
	stpHelloTime	
	stpForwardDelay	
	srBridgeNumber	
	port	
	packetFilter	

- 1 From the top level of the Administration Console, enter:
bridge port stpState
You are prompted for slot number(s).
- 2 Enter the number(s) of the slot(s) or **all** to enable or disable ports for Spanning Tree on all ports in the system.
You are prompted for the port type.
- 3 Enter **Ethernet**, **FDDI**, **tokenring**, or **all**.

You are prompted for the port number(s).

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

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

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

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

```
Slot 5 Ethernet port 4 - Enter new value (enabled,disabled)
[enabled]: disabled
Slot 5 Ethernet port 5 - Enter new value (enabled,disabled)
[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 means the less through traffic it will carry. For this reason, you may 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:

- From the top level of the Administration Console, enter:
bridge port stpCost
You are prompted for the slot number(s).
- Enter the number(s) of the slot(s) or **all** if you want to configure path cost on ports for all bridges in the system.
You are prompted for the port type.
- Enter **Ethernet, FDDI, tokenring, or all**.
You are prompted for the port number(s).
- Enter the number(s) of the port(s) or **all** to configure path cost for all ports on each bridge.

Top-Level Menu

system	display	summary
ethernet	mode	detail
fddi	ipFragm	multicastLimit
tokenring	ipxSnap	stpState
bridge	trFddiM	stpCost
ip	address	stpPriority
snmp	agingTim	srRingNumber
analyze	stpState	address
script	stpPrior	
logout	stpMaxA	
	stpHelloTime	
	stpForwardDelay	
	srBridgeNumber	
	port	
	packetFilter	

You are prompted for the path cost for each port you specified.

5 Enter the path cost for the port(s).

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

```
Slot 2 FDDI port 1 - Enter a new value [100]: 50
Slot 5 Ethernet port 3 - Enter a new value [100]: 200
Slot 5 Ethernet port 4 - Enter a new value [100]: 200
```

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

Setting the Port Priority

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

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 slot number(s).

2 Enter the number(s) of the slot(s) or **all** if you want to configure the port priority on ports for all bridges in the system.

You are prompted for the port type.

3 Enter **Ethernet**, **FDDI**, **tokenring**, or **all**.

You are prompted for the port number(s).

4 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.

5 Enter the port priority for the port(s).

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

Top-Level Menu

```
system
ethernet
fddi
tokenring
bridge
ip
snmp
analyze
script
logout
  display
  mode
  ipFragm
  ipxSnap
  trFddiM
  address
  agingTim
  stpState
  stpPriori
  stpMaxAge
  stpHelloTime
  stpForwardDelay
  srBridgeNumber
  port
  packetFilter
  summary
  detail
  multicastLimit
  stpState
  stpCost
  stpPriority
  srRingNumber
  address
```

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

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

Setting the Source Route Ring Number

The Source Route Ring Number is used by the token ring bridge when Source Route traffic is enabled in SR and SRT modes. When combined with the bridge number for the module and the ring number of the output port, it serves to uniquely identify this bridge for forwarding of packets.

Values for the Source Route Ring Number range from 1 to 4095 and is entered as a decimal number. The default value for this parameter is:

Default = slot number x 100 + port number



This parameter is only applicable to Token Ring and FDDI ports.

To set the Source Route Ring Number:

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

```
bridge port srRingNumber
```

You are prompted for the slot number(s) or **a11** to access all Token Ring and FDDI ports.

- 2 Enter new value [1 through 4095].

Top-Level Menu

```
system
ethernet
fddi
tokenring
bridge
ip
snmp
analyze
script
logout
  display
  mode
  ipFrag
  ipxSnaps
  trFddi
  address
  agingTime
  stpState
  stpCost
  stpPriority
  srRingNumber
  address
  stpMaxAge
  stpHelloTime
  stpForwardDelay
  srBridgeNumber
  port
  packetFilter
```

Administering Port Addresses

You can administer the MAC addresses of stations connected to Ethernet and FDDI 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

```

system
ethernet
fddi
tokenring
bridge
ip
snmp
analyze
script
logout
  display
  mode
  ipFragr
  ipxSna
  trFddi
  address
  agingT
  stpStat
  stpPrio
  stpMax
  stpHelloTime
  stpForwardDelay
  srBridgeNumber
  port
  packetFilter
    summar
    detail
    multicas
    stpState
    stpCost
    stpPriori
    srRingNu
    address
    list
    add
    remove
    find
    flushAll
    flushDynamic
    freeze
  
```

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

```
bridge port address list
```

You are prompted for slot number(s).

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

You are prompted for the port type.

- 3 Enter **Ethernet**, **FDDI**, **tokenring**, or **all**.

You are prompted for the port number(s).

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

An example of an address list follows.

Addresses for slot 8, Ethernet port 1:

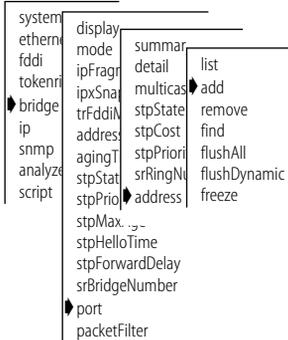
Ethernet address	Type	Age (secs.)
08-00-20-1d-67-e2	Dynamic	219
00-80-3e-02-68-00	Dynamic	219
00-20-af-29-7b-74	Dynamic	219
08-00-02-05-91-c1	Dynamic	219
00-80-3e-02-6d-00	Dynamic	219
00-80-3e-08-5f-00	Dynamic	219
00-80-3e-00-3d-00	Dynamic	219

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 slot number(s).

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

You are prompted for the port type.

- 3 Enter **Ethernet**, or **FDDI**, **tokenring**, or **all**.

You are prompted for the port number(s).

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

You are prompted to add an address.

- 5 Add the MAC addresses, pressing [Return] after each entry.

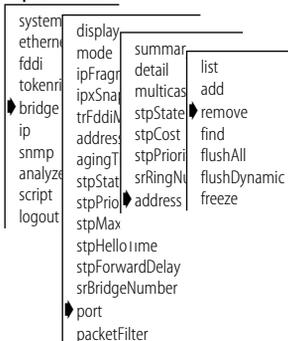
When you finish adding addresses, enter **q** to return to the previous menu.

Removing Addresses

You can remove individual MAC addresses from the 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 slot number(s).

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

You are prompted for the port type.

- 3 Enter **Ethernet**, or **FDDI**, **tokenring**, or **all**.

You are prompted for the port number(s).

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

You are prompted for the address to be removed.

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

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

Flushing All Addresses

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



When a switching module is moved to a different slot or chassis, all of its addresses are flushed automatically.

To flush *all* addresses:

Top-Level Menu

```

system
ethernet
fddi
tokenring
bridge
ip
snmp
analyze
script
logout
  display
  mode
  ipFragr
  ipxSna
  trFddi
  address
  agingT
  stpStat
  stpPrio
  stpMax
  stpHelloTime
  stpForwardDelay
  srBridgeNumber
  port
  packetFilter
  summar
  detail
  multicas
  stpState
  stpCost
  stpPrio
  srRingN
  address
  list
  add
  remove
  find
  flushAll
  flushDynamic
  freeze

```

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

```
bridge port address flushAll
```

You are prompted for slot number(s).

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

You are prompted for the port type.

- 3 Enter **Ethernet**, **FDDI**, **tokenring**, or **all**.

You are prompted for the port number(s).

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

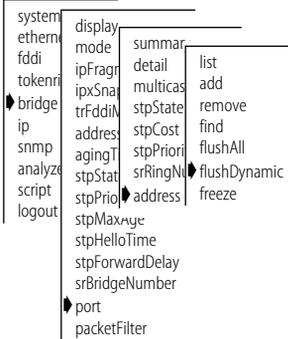
Static addresses are flushed from the ports you specified.

Flushing Dynamic Addresses

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

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 slot number(s).

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

You are prompted for the port type.

- 3 Enter **Ethernet, FDDI, tokenring, or all**.

You are prompted for the port number(s).

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

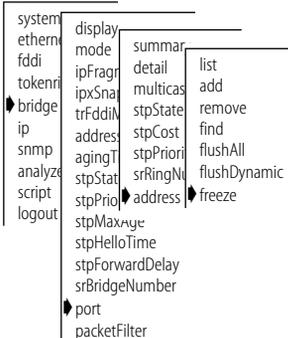
The addresses are flushed from the address table.

Freezing Dynamic Addresses

You can convert all the dynamic addresses associated with the selected port(s) on the module(s) you select into static addresses. This 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 slot number(s).

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

You are prompted for the port type.

- 3 Enter **Ethernet, FDDI, tokenring, or all**.

You are prompted for the port number(s).

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

The dynamic addresses become static.

14

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, and assign filters to ports on a switching module.

About Packet Filtering

Independently configurable packet filtering is provided for the various packet processing paths on each Ethernet and FDDI port of a switching module. The packet processing paths are defined in Table 14-1.

Table 14-1 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 and/or receive paths of each port.



Filters are stored on their assigned switching module(s), which means that all operations on filters are module specific. You can specify multiple switching modules when administering filters. If a switching module is moved to a different slot or chassis, the filters on that module are automatically removed.

For detailed explanations of packet filter concepts, see Chapter 7: User-defined Packet Filtering in the LANplex 6000 Operation Guide.

Listing Packet Filters

Top-Level Menu

```

system
ethernet
fddi
tokenring
bridge
ip
snmp
analyze
script
logout
display
mode
ipFragme
ipxSnapT
addressTr
agingTim
stpState
stpPriority
stpMaxAc
stpHelloT
stpForwal
port
packetFilter
list
display
create
delete
edit
load
copy
assign
unassign
addressGroup
portGroup

```

When you list the packet filters for switching modules in the system, the filter identification, filter name (if any), and filter assignments are displayed.

To list the currently defined packet filters, enter the following from the top level of the Administration Console:

```
bridge packetFilter list
```

The listing of packet filters is displayed. An example of the output follows:

```

Ethernet Packet Filters
  Packet Filter 1 - Receive OUI 08-00-1E
    Slot 3, Port 4, Transmit Multicast
    Slot 5, Port 3, Transmit Multicast
    Slot 5, Port 3, Receive Multicast
    Slot 7, Port 5, Receive Multicast
  Packet Filter 2 - Type > 900 or Multicast
    Slot 3, No port assignments
    Slot 4, Port 6, Receive All
    Slot 4, Port 8, Transmit All
    Slot 4, Port 8, Receive All
  Packet Filter 3 - Forward IP packets only
  No port assignments

```

In this example, there are two packet filters on the system. The first packet filter has a filter id of 1 and a user-defined name of "Receive OUI 08-00-1E." This filter is loaded on three ESMS in the system (slots 3, 5, and 7). On slot 3, the filter is assigned to the *transmit multicast* path of port 4. On slot 5, the filter is assigned to both the *transmit multicast* and *receive multicast* paths of port 3. Slot 7, port 5 is assigned to the *receive multicast* path.

The second filter (filter id 2, user name "Type > 900 or Multicast") is loaded on slots 3 and 4. No port assignments have been made on slot 3. The filter is assigned to both the *receive all* and *transmit all* paths of port 8 on slot 4. The final filter (filter id 3, no user-defined name) is loaded on slot 4 with no port assignments.

Displaying Packet Filters

When displaying the contents of a single packet filter, you select the packet filter using the filter id (which you can obtain by listing the packet filters as described in the previous section). The packet filter instructions are displayed; however, any comments in the original packet filter definition file are not displayed since they are not saved with the packet filter.

To display the contents of a packet filter:

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

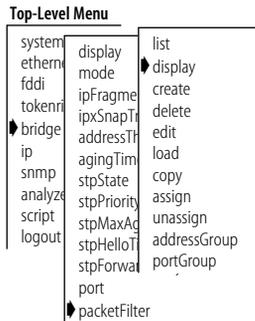
```
bridge packetFilter display
```

You are prompted for the number of the packet filter you want to display.

- 2 Enter the packet filter id number.

The contents of the packet filter are displayed. An example of the output generated by this command is shown below. The packet filter id and name are displayed, followed by a listing of the packet filter instructions.

```
Select packet filter to be displayed [1-n]: 2
Packet filter 2 - Type > 900 or Multicast
  name "Type > 900 or Multicast"
  pushLiteral.w                0x900
  pushField.w                   12
  gt
  reject
  pushField.b                    0
  pushLiteral.b                 0x01
  and
  not
```



Creating Packet Filters

You create custom packet filters by writing a *packet filter definition*. Each packet processing path on a port may have a unique packet filter definition or may share a definition with other ports on the module. Packet filter definitions consist of 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 of the system. This ensures that the packet filter configuration for each system is saved across system reboots or power failures.

Concepts for Writing a Filter

Before writing a packet filter, you should understand the basic concepts listed below:

- How the packet filter language works
- What the basic elements of a packet filter are
- How to implement sequential tests in a packet filter
- What the pre-processed and run-time storage requirements are

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) type of 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 14-3 and Table 14-4 on “Packet Filter Operators”). 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, and *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 14-2 describes the instructions and stacks of a packet filter.

Table 14-2 Packet Filter Instructions and Stacks — Descriptions and Guidelines

Element	Descriptions and Guidelines
Instructions	<p>Each instruction in a packet filter definition must be on a separate line in the packet filter definition file.</p>
<i>Instruction format</i>	<p>An instruction consists of an opcode followed by explicit operands and a comment. Although comments are optional, it is recommended that you use them throughout the packet filter for easier administration of the filters. The opcode includes an explicit operand size specification.</p>
	<p>The general format of an instruction is:</p>
	<pre><opcode>[.<size>] [<operand>...] [# <comment>]</pre>
	<p>For example:</p>
	<pre>pushliteral.1 0xffffffff #load the type field mask</pre>
	<p>You can use any combination of upper and lower case letters for the opcode and size. The contents of a line following the first # (which is outside a quoted string) are ignored.</p>
<i>Operand sizes</i>	<p>The following operand sizes are supported:</p>
	<ul style="list-style-type: none"> ■ 1 byte = .b ■ 2 bytes = .w ■ 4 bytes = .l ■ 6 bytes = .a
	<p>The six-byte operand size is included primarily for use with 48-bit, IEEE, globally-assigned MAC addresses.</p>
<i>Maximum length</i>	<p>The maximum length for a filter definition is 4096 bytes.</p>
Stack	<p>The packet filter language uses a stack to store the operands that will be used by an instruction and the results of the instruction.</p>
	<p>Operands are popped from the stack as required by the instructions. An instruction using two or more operands takes the first operand from the top of the stack with subsequent operands taken in order from succeeding levels of the stack.</p>
	<p>The stack is a maximum of 64 bytes long with space in the stack allocated in multiples of 4 bytes. This provides for a maximum of 16 operands on the stack.</p>
	<p>An address size operand (.a) consumes 8 bytes on the stack, decreasing the maximum number of operands on the stack.</p>

Basic Elements of a Packet Filter

Before creating a packet filter, you must decide which part of the packet you want to filter. This can be the destination address, source address, type/length, 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 14-1.

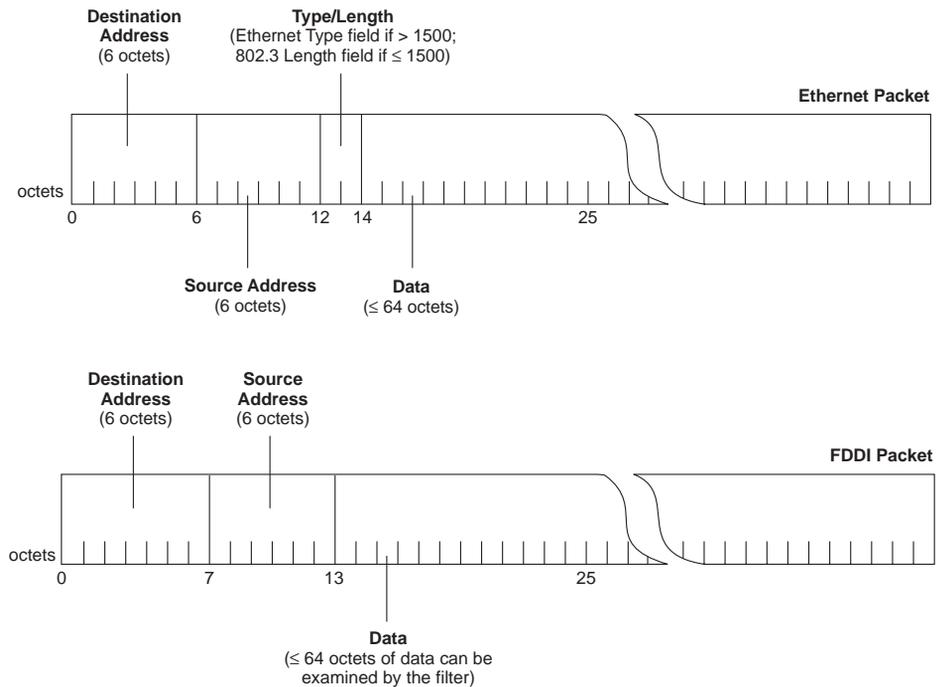


Figure 14-1 Ethernet and FDDI Packet Fields

The Ethernet and FDDI packet fields in Figure 14-1 are used as operands in the packet filter. The two simplest operands are described in Table 14-3.

Table 14-3 Packet Filter Operands

Operand	Description	Opcode to Use
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.	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 14-4.

Table 14-4 Packet Filter Operators

Operator	Result	Opcode to Use
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. This means that each bit of the operands are logically compared to produce the resulting bit.

Implementing Sequential Tests in a Packet Filter

Filter language expressions are normally evaluated to completion — a packet is accepted if the value remaining on the top of the stack is non-zero. Frequently, however, a single test is insufficient to filter packets effectively. Where more tests are warranted, you want to accept a packet that either:

- Satisfies at least one criterion specified in two or more tests (ORs the results of the tests), or
- Satisfies all criteria specified in two or more tests (ANDs the results of the tests)

The *accept* and *reject* instructions are used to implement sequential tests, as shown in Figure 14-2. 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.

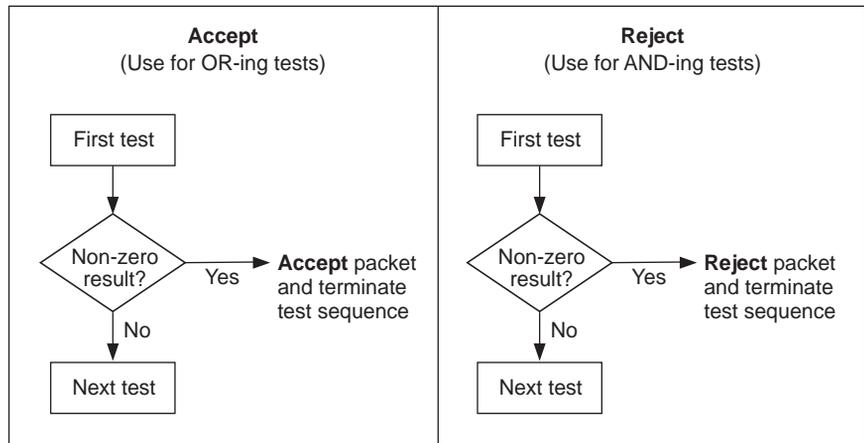


Figure 14-2 Accept and Reject Instructions

The following example shows the use of both accept and reject in a packet filter. This packet filter was created for a network running both Phase I and Phase II AppleTalk.™ The goal of the filter is to eliminate the AppleTalk traffic.

```
Name      "Filter AppleTalk datagrams"
pushField.w      12      # get the type field
pushTop          # make a copy
pushLiteral      0x809b  # EtherTalk Phase I type
eq               # test if the packet type is
                 # equal to the AppleTalk type
reject          # reject the packet and end
                 # otherwise
pushLiteral.w    0x5dc   # largest 802.3 packet size
lt               # if this value is less than the
                 # value in the packet's
                 # type/length field, then this
                 # is an Ethernet frame, so
accept          # accept the packet if it is not
                 # 802.3, otherwise...
pushField.a      16      # get the SNAP OUI and Ethertype
pushLiteral.a    0x03080007809b # value to compare
ne               # if not equal then forward the
                 # packet, otherwise drop it
```

Pre-processed 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 switching module is limited to a maximum of 16 packet filter programs.

Preprocessed packet filters

Each switching module 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 along with 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 is the only packet filter 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 switching module provides a maximum of 8192 bytes. There is no explicit system or per packet filter overhead; however, performance considerations may result in unused areas of the run-time storage.

The run-time format is approximately eight times the size of the stored format. Thus a 200-byte packet filter program in stored format expands to approximately 1600 bytes in the run-time format. A single packet filter program cannot exceed 2048 bytes in the run-time format.

Procedure for Writing a Filter

The following steps show the process of writing a packet filter. Detailed examples are provided in the section “Examples of Creating Filters” on page 14-12.

Write the instructions for the packet filter using the following format:

```
<opcode>[.<size>]  [<operand>...] [# <comment>]
```

You can find the opcode descriptions in the section Appendix A: *Packet Filter Opcodes, Examples, and Syntax Errors*. The description of the operand sizes supported can be found in Table 14-2 on “Packet Filter Instructions and Stacks — Descriptions and Guidelines”. 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 the program is loaded into the system.

When writing a packet filter, ensure that you use comments (preceded by #) to describe each step in the filter. This will help you to revise filters in the future and enable others to understand and use the filters you create.

To write a packet filter:

- 1 Assign a unique, descriptive name to the filter using the `NAME` opcode.
- 2 Specify what to test. For example, use the `PUSHFIELD` opcode to select a field in the packet.
- 3 Specify what to compare to the value in step 2. For example, use the `PUSHLITERAL` opcode to select a constant value.

- 4 Apply a logic operation to the values in steps 2 and 3. The operator you use depends on what comparison you want to make.

Variations to these four basic steps of writing packet filters include:

- Using `pushTop` for each additional comparison you intend to make with the `pushField` value. This 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 just used a second `pushfield` instruction.
- Using `accept` or `reject` with `and` and `or` operators when you have sequential tests, and you would like the filter to accept or reject a packet before the entire expression has been evaluated. Using `accept` and `reject` can significantly improve the performance of certain types of filters. See the section “Implementing Sequential Tests in a Packet Filter” on page 14-9 for more information.
- Using `pushSAGM`, `pushDAGM`, `pushSPGM`, or `pushDPGM` for filtering by address or port groups. See Chapter 14: *Configuring Address and Port Groups to Use in Packet Filters* for more information.

Examples of Creating Filters

The following is an example of 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 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, it is necessary to filter the broadcasts.

Packet Filter Solution

The solution 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 either on TCP/IP or on XNS. 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 the case of IP datagrams, and 30 bytes into the packet in the case of XNS datagrams.

The above information can be used to create pseudocode that simplifies the process of writing the actual filter. The pseudocode may be written in outline form, as shown below:

- 1 Determine if the packet has a broadcast address. (This is done through 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
 - 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
 - b The socket value is less than 0x898
- 6 End the filter.

The pseudocode translates into the following packet filter:

```

Name      "IP XNS ticker bcast filter"
          # Assign this filter in the multicast path
          # of a port only--this is very important
          #
          # XNS FILTERING SECTION
          #
pushField.w      12      # get the type field of the packet and
                    # place it on top of the stack
pushLiteral.w    0x0600  # put the type value for XNS 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 highest socket value on top of
                    # the stack
pushField.w      30      # 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 lowest socket value on top of
                    # the stack
pushField.w      30      # 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, a "one"
                    # will be placed on the stack.
and             # compare if XNS & in range
                    #
                    # IP FILTERING SECTION
                    #
pushField.w      12      # get the type field of the packet and
                    # place it on top of the stack
pushLiteral.w    0x0800  # put the type value for IP on top of
                    # the stack
eq              # if the two values on the top of the
                    # stack are equal, then return a non-zero
                    # value
pushLiteral.w    0x76c   # put the lowest socket value on top of
                    # the stack (1900)
pushField.w      24      # put the value of the socket from the
                    # packet on top of the stack
ge              # compare if the value of the socket is
                    # greater than or equal to the lower bound
pushLiteral.w    0x0898  # put the highest socket value on top of
                    # the stack (2200)
pushField.w      24      # put the value of the socket from the
                    # packet on top of the stack
lt              # compare if the value of the socket is
                    # less than the upper bound
and             # "and" together with "ge" and "lt" test
                    # to determine if the socket value is
                    # "within" the range. If it is in range, a
                    # "one" will be placed on the stack.
and             # compare if IP and in range
or              # determine if the type field is either
                    # XNS or IP
not             # discard if (IP & in range) & (XNS & in
                    # range)

```

The rest of this section concentrates on the parts of the filter, showing you how to translate the pseudocode's requirements into filter language. The large filter on page 14-14 is broken down into subsets to show how you can create small filters that perform one or two tasks, then combine them for more sophisticated filtering. Table 14-5 shows how the purpose of each pseudocode step is accomplished in the following series of packet filters.

Table 14-5 Pseudocode Requirements Mapped to the Packet Filter

Step	Accomplished Through...
1	The path to which you assign the packet filter. For administrative purposes, this is specified in the first two comment lines in the filter definition. The filter has to be assigned to a multicast path to filter packets with broadcast addresses.
2	Packet Filter One — Forwarding XNS packets
3	Packet Filter Two — Checking for specified socket range
4 & 5	Combining a Subset of Filters — Forwarding IP packets within specified socket range

Packet Filter One. This filter is designed to forward XNS packets. The steps below show how to create this filter.

- 1 Name the filter:

```
"Forward only XNS packets"
```

This is important to distinguish the function of each filter when it is loaded onto a switching module that has more than one filter stored in memory. This is also useful for archiving filters on an ftp server so that the filters may 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
```

By itself, this filter looks like the following:

```
Name      "Forward only XNS packets"
pushField.w  12      # get the type field of the packet and
                  # place it on top of the stack
pushLiteral.w 0x0600 # put the type value for XNS on top of
                  # the stack
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. The steps below show how to create this filter.

1 Name the filter:

```
"Socket range filter"
```

2 Enter executable instruction #1:

```
pushLiteral.w 0x76c # put the highest 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 lowest socket value on top
of the stack
```

6 Enter executable instruction #5:

```
pushField.w 30 # put the value of the socket from the
packet on top of the stack
```

7 Enter executable instruction #6:

```
lt # compare if the value of the socket is less than the
upper bound
```

8 Enter executable instruction #7:

```
and # "and" together with "ge" and "lt" test to determine
if the socket value is "within" the range. If it is, a
"one" will be placed on the stack.
```

By itself, this filter looks like the following:

```
Name          "Socket range filter"
pushLiteral.w  0x76c          # put the lowest socket value on top of
# the stack (1900)
pushField.w    30           # 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    30           # 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, a "one"
# will be placed on the stack.
```

Combining a Subset of the Filters. This filter accepts IP packets with a socket range of 0x76c (1900) and 0x898 (2200). This filter combines packet filters one and two, modifying them for IP. The steps below show how to create this filter.

1 Name the filter:

```
"Only IP pkts w/in socket range"
```

2 Go through steps 2 through 4 as described in "Packet Filter One" on page 14-15, except the pushLiteral instruction (in step 3) should have a value of 0x0800 for IP.**3** Go through steps 2 through 8 as described in "Packet Filter Two" on page 14-16, except the socket value for IP (in step 3) is located 24 bytes into the packet (instead of 30 as in the case of 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 would look like the following:

```
Name      "Only IP pkts w/in socket range"
pushField.w    12      # get the type field of the packet and
                # place it on top of the stack
pushLiteral.w  0x0800  # put the type value for IP on top of
                # the stack
eq            # if the two values on the top of the
                # stack are equal, then return a non-zero
                # value
pushLiteral.w  0x76c   # put the lowest socket value on top of
                # the stack (1900)
pushField.w    24      # put the value of the socket from the
                # packet on top of the stack
ge            # compare if the value of the socket is
                # greater than or equal to the lower bound
pushLiteral.w  0x0898  # put the highest socket value on top of
                # the stack (2200)
pushField.w    24      # put the value of the socket from the
                # packet on top of the stack
lt            # compare if the value of the socket is
                # less than the upper bound
and           # "and" together with "ge" and "lt" test
                # to determine if the socket value is
                # "within" the range. If it is in range, a
                # "one" will be placed on the stack.
and           # compare if IP and in range
```

Combining All the Filters. Together, the four packet filters work to perform the solution to the problem: filtering the broadcast packets from the market data servers. The steps below show how to create this filter:

- 1 Name the filter:

```
"Discard XNS & IP pkts w/in socket range"
```

- 2 Go through steps 2 through 4 as described in "Packet Filter One" on page 14-15.
- 3 Go through steps 2 through 8 as described in "Packet Filter Two" on page 14-16.
- 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 Go through steps 2 through 4 as described in “Combining a Subset of the Filters” on page 14-17.

6 Add an *or* statement:

```
or # determine if the type field is either XNS or IP
```

7 Add a *not* statement to cause any matching packets to be discarded:

```
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 14-14.

```
Select slot(s) [1-n, all]: 3,6,7
```

Tools for Writing a Filter

You can create a new packet filter using either an ASCII-based text editor (such as *EMACS* or *vi*) or the line editor built into the Administration Console. Using an ASCII-based text editor allows you to create multiple copies of the packet filter definition, which you can then copy onto one or more LANplex systems from a networked workstation. This method also allows you to archive copies of filter definitions.

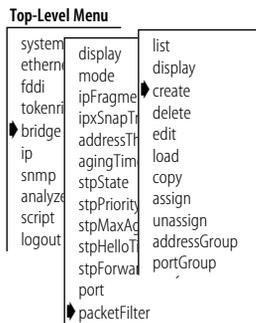
Using the Built-in Line Editor

The built-in text editor provides a minimal set of editing functions that you can use to edit a packet filter definition one line at a time. A single line is limited to no more than 79 characters, while the number of lines is limited only by available memory.



The maximum length of a packet filter definition is 4096 bytes.

The editor assumes a terminal capability no higher than a glass tty (that is, an addressable screen is not assumed). You can place any ASCII printable character into the editing buffer at the cursor position. If a character exceeds the maximum line length, the character is discarded and a bell sounds. The editor initially operates in *insert mode*. The commands supported by the editor are summarized in Table 14-6.



To use the built-in line editor to create a packet filter definition:

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

```
bridge packetFilter create
```

The packet filter line editor appears.

- 2 Enter the definition for the packet filter. See the text editor in Table 14-6.

- 3 Save the packet filter by pressing Ctrl-W.

The syntax of the filter definition is checked. If any errors are detected, the errors are displayed and the editor is re-entered at the line containing the first error. After correcting the errors, 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 stored on the selected modules.

Table 14-6 Packet Filter Editor Commands

Command	Keys	Description
List buffer	Ctrl-l	Displays each of the lines in the editing buffer and then re-displays the current line being edited
Next Line	Ctrl-n	Moves cursor between lines, positions cursor at start of line
Previous Line	Ctrl-p	Moves cursor between lines, positions cursor at start of line
Start of Line	Ctrl-a	Moves cursor within a line to the start of that line
End of Line	Ctrl-e	Moves cursor within a line to the end of that line
Left 1 Character	Ctrl-b	Moves cursor within a line <i>left</i> one character
Right 1 Character	Ctrl-f	Moves cursor within a line <i>right</i> one character
Insert Line	Enter	Inserts a new line. The new line becomes the current line with the cursor positioned at the start of the line. If the cursor is positioned over the first character on a line, a blank new line is inserted prior to the current line. Otherwise the new line is inserted after the current line. In this case, the current line is split at the cursor position with the current line retaining the characters before the cursor and the new line containing the remainder of the characters.
Delete Previous Character	Ctrl-h or BSP	Deletes a single character preceding the cursor and shifts the remainder of the line <i>left</i> one position
Delete Current Character	Ctrl-d or DEL	Deletes a single character under the cursor and shifts the remainder of the line <i>left</i> one position

(continued)

Table 14-6 Packet Filter Editor Commands (continued)

Command	Keys	Description
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 removes the line from the edit buffer.
Insert/Overstrike Toggle	Ctrl-o	Toggles between the insert and overstrike modes
Write Changes	Ctrl-w	Writes 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, Error: Edit buffer exceeds maximum length , is displayed. The contents of the editing buffer are unaffected; however, the packet filter definition only contains those lines that fit entirely within the length limitation.
Exit Editor	ESC	Allows you to leave the editor. A warning is issued if the edit buffer has not been successfully written since the last modification. You have the option of either discarding the changes or returning 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 Enter the definition for the filter into a text file.
- 2 From a networked workstation, ftp the file to the LANplex system where you want to load the filter.
- 3 Load the filter as described in “Loading Packet Filters” on page 14-23.

Deleting Packet Filters

Deleting a packet filter removes the filter from the switching module where it is loaded.

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 get the id of the filter, list the filters as described in the section “Listing Packet Filters” on page 14-2.

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	display	list
ethernet	mode	display
fddi	ipFragme	create
tokenr	ipxSnapT	delete
bridge	addressTh	edit
ip	agingTim	load
snmp	stpState	copy
analyz	stpPriority	assign
script	stpMaxAg	unassign
logout	stpHelloT	addressGroup
	stpForwar	portGroup
	port	
	packetFilter	

Editing Packet Filters

You can use the LANplex system line editor to edit packet filters. Once you save the packet filter, it is checked for syntax errors. The LANplex system software will not allow you to assign the packet filter to any slots 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 onto the LANplex system from a networked workstation.

To edit a packet filter:

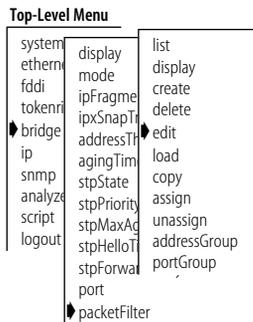
- 1 From the top level of the Administration Console, enter:
bridge packetFilter edit
- 2 Enter the packet filter id number.
Specifying a filter id loads that filter into the edit buffer.
- 3 Edit the filter. For more information see the section "Using the Built-in Line Editor" on page 14-19.
- 4 Press [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 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 syntax of the filter definition is checked. If any errors are detected, the errors are displayed and the editor is re-entered at the line containing the first error. After correcting the errors, you must exit the editor and attempt to save the packet filter again.



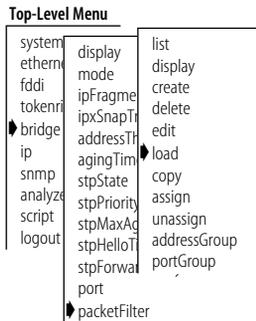
After you have corrected all errors and saved the packet filter, it is converted to internal form and updated on the switching modules to which it is assigned.

For filters edited with an external text editor, filter errors are checked when the filter is loaded onto the module (see “Loading Packet Filters”). If errors are detected when loading the filter, you must make corrections in the external editor and try to load the filter again.

Loading Packet Filters

When you create packet filters using an external text editor, you must load the filters onto a module 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 switching module.

To load a packet filter:



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

```
bridge packetFilter load
```

- 2 Enter the packet filter id number.

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.

- 3 Enter the host IP address.
- 4 Enter the path name.
- 5 Enter your user name.
- 6 Enter your password.
- 7 Enter the slot number(s).

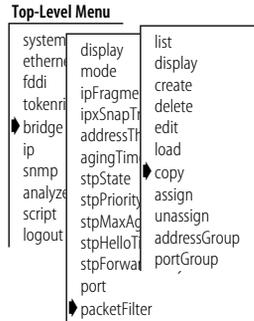
The packet filter is loaded onto the LANplex system.

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.

Copying Packet Filters

You can copy a packet filter from one module to another.

To copy a packet filter:



- 1 From the top level of the Administration Console, enter:
bridge packetFilter copy
- 2 Enter the number of the packet filter you want to copy.
- 3 Enter the number of the slot where you want to place the copied filter.

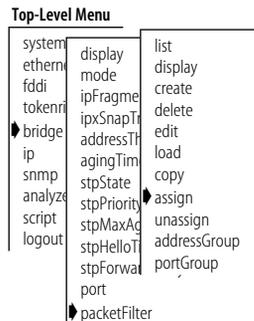
Assigning Packet Filters to Ports

To assign a packet filter to one or more ports, the packet filter must be resident on at least one switching module. 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/or ports.

Packet filter path assignments

Placing a filter on the transmit path confines the packet to the segment it originated from if it does not meet the forwarding criteria. Placing a filter on the receive path prohibits a packet from accessing certain segments unless it meets the forwarding criteria. A packet that does not meet the forwarding criteria defined in the filter is discarded.

To assign a packet filter:



- 1 From the top level of the Administration Console, enter:
bridge packetFilter assign
- 2 Enter the id number of the packet filter to be assigned. To get the id of the packet filter, you can list all packet filters as described in "Listing Packet Filters" on page 14-2.
- 3 Enter the port type (**Ethernet, FDDI, all**).
- 4 Enter the port(s) to assign the filter.
- 5 Enter the path(s) you want to place the filter (**txA, txM, rxA, rxM, all**).

In the example below, the assignment is to the transmit all (txA) and receive all (rxA) paths on port 1 of slots 3 and 4.

```
Select filter [1-n]: 1
Select slot(s) [3-5, all]: 3,4
Select port type(s) (Ethernet,FDDI|all) [Ethernet,FDDI]:
FDDI
Select port(s) (1-16|all) [1-16]: 1
Select path(s) (txA,txM,rxA,rxM|all): txA,rxA
```

The slots are limited to those that have the filter loaded and have at least one port/path unassigned. 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 slots, ports, and paths taken in combination.



One or more assignments may fail due to a previous assignment.

Unassigning Packet Filters from Ports

To unassign a packet filter from one or more ports, the packet filter must have been previously assigned to at least one port.

To unassign a packet filter:

- 1 From the top level of the Administration Console, enter:
bridge packetFilter unassign
- 2 Enter the id number of the packet filter to unassign.
- 3 Enter the slot number of the packet filter to unassign.
- 4 Enter the port type (**Ethernet, FDDI, all**).
- 5 Enter the port number of the packet filter to unassign.
- 6 Enter the path of the packet filter to unassign.

An example of unassigning a packet filter is shown below. In this example, the unassignment is from the transmit all (txA) paths on port 1.

```
Select filter [1-n]: 1
Select slot(s) [3-5, all]: 3,4
```

Top-Level Menu

system	display	list
ethernet	mode	display
fddi	ipFragme	create
tokenr	ipxSnapT	delete
bridge	addressT	edit
ip	agingTim	load
snmp	stpState	copy
analyze	stpPriority	assign
script	stpMaxAd	unassign
logout	stpHelloT	addressGroup
	stpForwa	portGroup
	port	
	packetFilter	

```
Select port type(s) (Ethernet,FDDI|all) [Ethernet,FDDI]: FDDI  
Select port(s) (1-16|all) [1-16]: 1  
Select path(s) (txA,txM,rxA,rxM|all) [txA,rxA]: txA
```

The slots are limited to those that have the filter loaded and have at least one port/path assigned to the filter. The ports are limited to those that have at least one path assigned to the filter, while the paths are limited to those that are assigned to the filter. 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 slots, ports, and paths taken in combination.



One or more of the unassignments may fail if the filter is not assigned.

15

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 LANplex 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 6000 Operation Guide.

A packet filter uses groups to make filtering decisions by accessing the source group mask and the destination group mask of a group. 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). Below are 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.l    0      # Put a zero on the stack
ne               # If not equal, returns a "one" to
                  # stack, resulting in packet
                  # forwarded
```

Port group packet filter example

In this example, packets are not forwarded to ports included in groups 3 and 8.

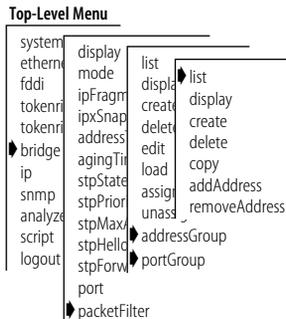
```
Name      "Discard Groups 3 and 8"
pushSPGM          # Get source port group mask
pushLiteral.1    0x0084 # Select bits 3 and 8
and              # If port group bits 3 & 8 are common
                # with SPGM, then non-zero value is
                # pushed onto stack
pushLiteral.1    0      # 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/remove addresses and ports to or from a group

Listing Groups

You can list the address and port groups currently defined on the switching modules of the LANplex system. The group id, group name (if any), group mask, and the slots where the group is loaded are displayed.



Address group example

To list currently defined *address* groups, enter the following from the top level of the Administration Console:

```
bridge packetFilter addressGroup list
```

To list currently defined *port* groups, enter the following from the top level of the Administration Console:

```
bridge packetFilter portGroup list
```

The listing of address or port groups is displayed.

The listing of address or port groups is displayed, as shown below. 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) and the address group mask is loaded into slots 3, 6, and 11.

```

Address Groups
Address Group 1 - Accounting
    Address group mask - bit 1
    Slot 3, 6, 11
Address Group 2 - Development
    Address group mask - bit 6
Slot 4
Address Group 3 - Sales
    Address group mask - bit 3
    Slot 5, 7, 10
    
```

Port group example

An example of listing port groups is shown below. In this example, 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 and is loaded into slots 4, 5, and 9 (the bit set in the mask).

```

Port Groups
Port Group 1 - Sales
    Port group mask - bit 7
    Slot 4, 5, 9
Port Group 2 - Manufacturing
    Port group mask - bit 23
    Slot 4
    
```

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 From the top level of the Administration Console, enter the following to display an *address* group:

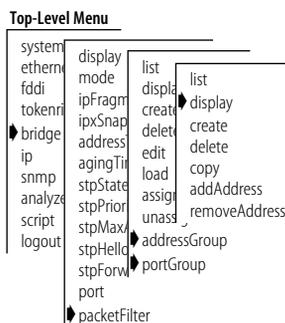
```
bridge packetFilter addressGroup display
```

Or enter the following to display a *port* group:

```
bridge packetFilter portGroup display
```

- 2 Enter the id number of the address or port group you want to display.

The address or port group you selected is displayed.



Address group example

An example of displaying an address group is shown below. In this example, address group 2 is displayed. The address group id and the name (if any) are displayed, followed by Ethernet addresses that are members of the group. The name of the address group in this example is *Development*, and the group has five members.

```
Select address group to be displayed [1-n]: 2
Address Group 2 - Development
05-39-24-56-ab-ee      08-29-34-fd-32-14      08-29-34-dd-ee-01
09-34-56-32-12-e3     00-14-32-54-fd-4e
```

Port group example

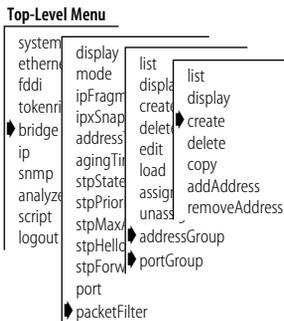
An example of displaying a port group is shown below. In this example, port group 2 is displayed. The port group id and the name (if any) are displayed, followed by the ports that are members of the group. The name of the port group in this example is *Manufacturing* and the group has three members — Ethernet ports 1 and 5 and FDDI port 1.

```
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, one slot with an unused address or port group must be available. A port group is limited to the number of ports on a module.

To create an address or port group:



- 1 From the top level of the Administration Console, enter the following to create a new *address* group:

```
bridge packetFilter addressGroup create
```

OR enter the following to create a new *port* group:

```
bridge packetFilter portGroup create
```

- 2 Enter the slot number(s).

For port groups, your selection is limited to slots that have a switching module and that have at least one unused port group.

- 3 For address groups, enter the address group mask.

For port groups, enter the port group mask.

- 4 Enter the address or port group name.
- 5 Enter the addresses or ports to add to the new group. Type **q** after entering all the addresses or ports.

Enter the addresses in MAC format as:

```
xx-xx-xx-xx-xx-xx
```

Enter the ports in the syntax:

```
< Ethernet | E | FDDI | F > [port] < port number >
```

As you enter each address or port, the system attempts to add it to each of the modules on which the group is loaded. If the address or port you enter is already a member of the group, a message is displayed, as shown below, 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 any of the modules fail to accept the additional address, the address is not added to the group and an error message is displayed as follows:

```
Error: No room in group for an additional address.
```

When this message occurs, the specified address is ignored and creation of the address group stops. All addresses entered up to the last address are added to the group and the group is loaded on the selected module(s).

If you enter an invalid port name, the port is not added to the group, and you receive one of the following error messages:

```
Error: No port type specified for the port.
```

```
Error: No port number specified for the port.
```

```
The correct format is < Ethernet | E | FDDI | F > [port] <
port number >
```

```
Specified port number is invalid.
```

```
Valid FDDI port for this group is 1 or 2.
```

Address group example

An example of creating a new address group is shown below. In this example, a new address group is created and loaded on the modules in slots 8 and 9. 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 slot(s) [4, 7-9, all]: 8-9
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 into slot 8
Address Group 4 - Marketing - has been loaded into slot 9

```

Port group example

An example of creating a new port group is shown below. In this example, a new port group is created and loaded on the module in slot 3. 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 slot(s) [3-5, all]: 3
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 into slot 3

```

Deleting Groups

When you delete address or port groups from a module, those groups are no longer available for use in packet filters. If you have saved the groups to an ASCII file, then you can use the groups in the future.

To delete an address or port group:

- 1 From the top level of the Administration Console, enter the following to delete an *address* group:

```
bridge packetFilter addressGroup delete
```

OR enter the following to delete a port group:

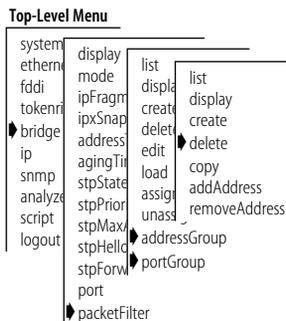
```
bridge packetFilter portGroup delete
```

You are prompted for the slot on which the address or port group resides.

- 2 Enter the slot number(s).

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.



Copying Groups

Copying from one module to another is faster than loading a group to multiple modules.

To copy an address or port group:

- 1 From the top level of the Administration Console, enter the following for the address group menu:

```
bridge packetFilter addressGroup copy
```

Or enter the following for the port group menu:

```
bridge packetFilter portGroup copy
```

You are prompted for the slot on which the address or port group resides.

- 2 Enter the slot number(s).

You are prompted for the slot number where you want to copy the address or port group.

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 or port group must exist before you can add addresses or ports. An address may be in multiple address groups.

Address group size

An address group for a switching module supports a maximum of 8192 addresses in Transparent bridging mode and 1024 in Express switching mode LANplex. 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 9, which is the maximum number of ports on a switching module.

To add addresses or ports to an existing group:

For port groups, entering an invalid port specification results in error messages, similar to those described on page 15-5.

Address group example

An example of adding addresses to a group follows. In the example, two additional addresses are added to the *Development* address group.

```
Select address group to be modified [1-4]: 2
Adding addresses to group 2 - Development
Enter the addresses to be added - type q to return to the menu:
Address: 08-21-42-62-98-ab
Address: 08-37-21-65-78-c4
Address: q
```

Port group example

The example below shows a port successfully added to the *Manufacturing* port group.

```
Select port group to be modified [1-4]: 2
Adding ports to group 2 - Manufacturing
Enter the ports to be added - type q to return to the menu:
Port: Ethernet 3
Port: q
```

Removing Addresses or Ports from a Group

When removing addresses or ports from a group, you can either enter the addresses or groups at the prompts, or import them from a file. At least one group must exist to remove an address or port.

To remove addresses or ports from an address group:

- 1 From the top level of the Administration Console, enter the following to remove addresses from an *address* group:

```
bridge packetFilter addressGroup removeAddress
```

Or enter the following to remove ports from a *port* group:

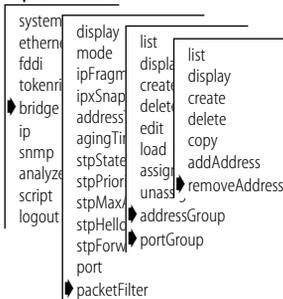
```
bridge packetFilter portGroup removePort
```

- 2 Enter the number of the group to modify.
- 3 Enter the addresses or ports to remove from the new group. Type **q** after entering all the addresses or ports.

Enter the addresses in MAC format as:

```
xx-xx-xx-xx-xx-xx
```

Top-Level Menu



Enter the ports in the syntax:

```
< Ethernet | E | FDDI | F > [port] < port number >
```

As you enter addresses and ports, the system attempts to remove them from each of the modules on which the group is loaded.

If the address or port is not found in the group, a warning message is displayed, as shown below:

```
Warning: Specified address was not a member of the address
group.
```

OR

```
Warning: Specified port was not a member of the port
group.
```

When this message occurs, the specified address or port is ignored and you are prompted for the next one to be removed.

*Address group
example*

An example of removing addresses from an address group is shown below. In this example, two Ethernet addresses are removed from the *Marketing* address group.

```
Select address group to be modified [1-4]: 4
Removing addresses from group 4 - Marketing
Enter the addresses to be removed - type q to return to the menu:
Address: 08-37-21-65-78-c4
Address: 08-42-21-84-78-f1
Address: q
```

Port group example

An example of removing ports from a port group is shown below. In this example, an Ethernet and an FDDI port are removed from the *Education* port group.

```
Select port group to be modified [1-4]:4
Removing ports from group 4 - Education
Enter the ports to be removed - type q to return to the menu:
Port: FDDI 1
Port: Ethernet 4
Port: q
```

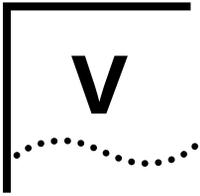
Loading Groups

There is no explicit menu item to load address and port groups that are defined in a file on a remote host. However, you can “load” groups by creating a script on a remote host (which includes your address or port group), and running that script.

The following example shows a script that builds an address group:

```
bridge packetFilter addressGroup create
08-37-21-65-78-c4
08-32-18-55-40-a0
08-22-12-65-78-05
08-18-23-00-82-00
08-52-12-65-5f-22
08-25-43-41-6e-09
08-00-65-23-00-ee
08-5a-42-77-8a-01
08-22-13-66-00-2a
08-8e-54-11-78-3b
08-77-12-65-78-8c
q
```

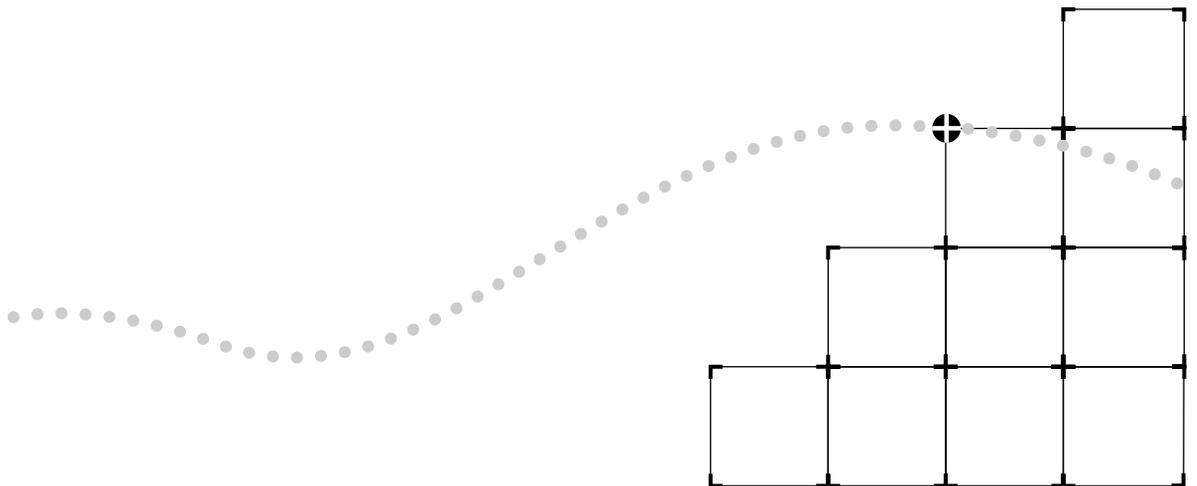
When you run the script, the address group is automatically created and stored on the system. For more information on running scripts, see Chapter 2: *How to Use the Administration Console*, on page 2-20.



APPENDICES

Appendix A Packet Filter Opcodes, Examples, and Sytax Errors

Appendix B Technical Support





PACKET FILTER OPCODES, EXAMPLES, AND SYNTAX ERRORS

This appendix describes the specific opcodes you can use when creating the packet filter and provides numerous examples for commonly used packet filters. It also describes the possible syntax errors you may receive when loading a packet filter.

Opcodes

Opcodes are instructions used in packet filter definitions. The available opcodes are described below:

name "<name>"

Description:

Assigns a user-defined <name> to the packet filter. The name may be any sequence of ASCII characters other than quotation marks. The name is limited to 32 characters. Only a single name statement can be included in a packet filter program.

Storage Needed:

2 + n bytes of packet filter storage where n is the length of the <name>

pushField.size <offset>

Description:

Pushes a field from the target packet onto the stack. Packet data starting at <offset> is copied onto the stack. The most significant byte of the field is the byte at the specified offset. The number of bytes pushed is determined by the size field of the instruction. The pushField instruction provides direct access to any 1, 2, 4, or 6 byte field contained within the first 65535 bytes of the target packet.

Certain implementations of the packet filter language further limit the maximum offset, based on the packet lengths supported by the underlying network. Ethernet based packet filters are limited to accessing fields in the first 1518 bytes of the target packet.

Specify the offset as either an octal, decimal, or hexadecimal number. An octal number must be preceded by a "0". A hexadecimal number must be preceded by either "0x" or "0X". Either upper or lower case letters can be used for the hexadecimal digits "a" through "f".

Storage Needed:

3 bytes

pushLiteral.size <value>

Description:

Pushes a literal constant <value> onto the stack. The most significant byte of the <value> is the first byte of the literal. Bytes are copied directly from the instruction stream onto the stack. The number of bytes pushed is determined by the size field of the instruction.

Specify the value as either an octal, decimal, or hexadecimal number. An octal number must be preceded by a "0". A hexadecimal number must be preceded by either "0x" or "0X". Either upper or lower case letters can be used for the hexadecimal digits "a" through "f".

Storage Needed:

2 (.b), 3 (.w), 5 (.l), or 7 (.a) bytes—depending on the size of the operand

pushTop

Description:

Pushes the current top of the stack onto the stack (that is, it reads the top of the stack and pushes the value onto the stack). The size of the push is determined by the size of the contents of the stack.

Storage Needed:

1 byte

pushSAGM

Description:

Pushes the source address group mask (SAGM) onto the top of the stack. The SAGM is a bitmap representing the groups to which the source address of a packet belongs. This instruction pushes 4 bytes onto the stack.

Each address group is represented by a single bit in the SAGM.

Multicast addresses (including broadcast addresses) are in all groups.

Storage Needed:

1 byte

pushDAGM

Description:

Pushes the destination address group mask (DAGM) onto the top of the stack. The DAGM is a bitmap representing the groups to which the destination address of a packet belongs. This instruction pushes 4 bytes onto the stack.

Each address group is represented by a single bit in the DAGM.

Multicast addresses (including broadcast addresses) are in all groups.

Storage Needed:

1 byte

pushSPGM

Description:

Pushes the source port group mask (SPGM) onto the top of the stack. The SPGM is a bitmap representing the groups to which the source port of a packet belongs. This instruction pushes 4 bytes on to the stack.

Each port group mask is represented by a single bit in the SPGM bitmap. Port group masks are assigned to the bitmap in sequence, starting with port group mask 1 as the least significant bit through port group mask 32 as the most significant bit.

Storage Needed:

1 byte

pushDPGM

Description:

Pushes the destination port group mask (DPGM) onto the top of the stack. The DPGM is a bitmap representing the groups to which the destination port of a packet belongs. This instruction pushes 4 bytes on to the stack.

Each port group mask is represented by a single bit in the DPGM bitmap. Port group masks are assigned to the bitmap in sequence, starting with port group mask 1 as the least significant bit through port group mask 32 as the most significant bit.

Storage Needed:

1 byte

eq (equal)

Description:

Pops two values from the stack and compares them. If they are equal, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

ne (not equal)

Description:

Pops two values from the stack and compares them. If they are not equal, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

lt (less than)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is less than the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

le (less than or equal to)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is less than or equal to the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

gt (greater than)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is greater than the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

ge (greater than or equal to)

Description:

Pops two values from the stack and performs an unsigned comparison. If the first is greater than or equal to the second, a byte containing the value non-zero is pushed onto the stack; otherwise, a byte containing 0 is pushed. The size of the operands is determined by the contents of the stack.

Storage Needed:

1 byte

and (bit-wise AND)

Description:

Pops two values from the stack and pushes the bit-wise *AND* of these values back onto the stack. The size of the operands and the result are determined by the contents of the stack.

Storage Needed:

1 byte

or (bit-wise OR)

Description:

Pops two values from the stack and pushes the bit-wise *OR* of these values back onto the stack. The size of the operands and the result are determined by the contents of the stack.

Storage Needed:

1 byte

xor (bit-wise exclusive-OR)

Description:

Pops two values from the stack and pushes the bit-wise *exclusive-OR* of these values back onto the stack. The size of the operands and the result are determined by the contents of the stack.

Storage Needed:

1 byte

not

Description:

A byte is popped from the stack; if it is non-zero, a zero byte is pushed back onto the stack. Otherwise, a non-zero byte is pushed back onto the stack.

Storage Needed:

1 byte

accept

Description:

Conditionally accepts the packet being examined. A byte is popped from the stack. If it is non-zero, the packet is accepted, and evaluation of the filter ends immediately; otherwise, filter evaluation continues with the next instruction.

Storage Needed:

1 byte

reject

Description:

Conditionally rejects the packet being examined. A byte is popped from the stack. If it is non-zero, the packet is rejected and evaluation of the filter ends immediately; otherwise, filter evaluation continues with the next instruction.

Storage Needed:

1 byte

shifl (shift left)

Description:

Pops two values from the stack and shifts the first operand left by the number of bits specified by the second operand. Bits shifted out the left side of the operand are discarded and zeros are shifted in from the right. The resulting value is pushed back onto the stack. The size of the first operand and the size of the result are determined by the contents of the top of the stack. The second operand is always one byte and only the low 5 bits of the byte are used as the shift count.

Storage Needed:

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 the right side of the operand are discarded and zeros are shifted in from the left. The resulting value is pushed back onto the stack. The size of the first operand and the size of the result are determined by the contents of the top of the stack. The second operand is always one byte and only the low 5 bits of the byte are used as the shift count.

Storage Needed:

1 byte

Packet Filter Examples

The following are examples of using the packet filter language. They start with basic packet filter concepts.

Destination Address Filter

This filter operates on the destination address field of a frame. It allows packets to be forwarded that are destined for stations with an 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.1   0           # Get first 4 bytes of
destination   # address
pushLiteral.1 0xffffffff # Set up mask to isolate first
              # 3 bytes
and           # Top of stack now has OUI
pushLiteral.1 0x08000200 # 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.1   0           # Get first 4 bytes of source
              # address
pushLiteral.1 0xffffffff # Set up mask to isolate first
              # 3 bytes
and           # Top of stack now has OUI
pushLiteral.1 0x08000200 # Load OUI value
eq           # Check for match
```

Length Filter

This filter operates on the length field of a frame. It allows packets to be forwarded that are less than 400 bytes in length. To customize this filter to another length value, change the literal value loaded in the **pushLiteral.w** instruction.

```
name          "Forward < 400"
pushField.w   12          # Get length field
pushLiteral.w 400         # Load length limit
lt           # Check for frame length < limit
```

Type Filter This filter operates on the type field of a frame. It allows packets to be forwarded that are IP frames. To customize this filter to another type value, change the literal value loaded in the **pushLiteral.w** instruction.

```

name                "Forward IP frames"
pushField.w         12                # Get type field
pushLiteral.w       0x0800           # Load IP type value
eq                  # Check for match

```

Ethernet Type IPX and Multicast Filter This filter *rejects* frames that have either a Novell IPX Ethernet type field (8134 hex) or a multicast destination address.

```

name                "Type > 900 or Multicast"
pushField.w         12                # Get type field
pushLiteral.w       0x900            # Push type value to test
against
gt                  # Is type field > 900 (hex)?
reject              # If yes: reject frame (done)
pushLiteral.b       0x01             # Multicast bit is low-order bit
pushField.b         0                # Get 1st byte of destination
and                 # Isolate multicast bit
not                 # Top of stack 1 to accept,
                   # 0 to reject

```

Multiple Destination Address Filter This filter operates on the destination address field of a frame. It allows packets to be forwarded that are destined for one of four different stations. To customize this filter to other destination stations, change the literal values.

```

name                "Forward to four stations"
pushField.a         0                # Get destination address
pushTop             # Make 3 copies of address
pushTop
pushTop
pushLiteral.a       0x367002010203 # Load allowed destination
address
eq                  # Check for match
accept              # Forward if valid address
pushLiteral.a       0x468462236526 # Load allowed destination
address
eq                  # Check for match
accept              # Forward if valid address
pushLiteral.a       0x347872927352 # Load allowed destination
address
eq                  # Check for match
accept              # Forward if valid address
pushLiteral.a       0x080239572897 # Load allowed destination
address
eq                  # Check for match

```

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 addr. group
mask
and                 # Compare if source and
destination         # groups are common members of
                    # an address group (result is
                    # either zero or non-zero)
                    # address group masks
pushLiteral.1      0 # Put a zero on the stack
ne                  # If not equal, returns a "one"
                    # to stack, resulting in packet
                    # forwarded

```

Port Group Filter This filter discards all frames sourced from a port in either group three or eight.

```

name                "Discard Port Groups 3 and 8"
pushSPGM            # Get source port group mask
pushLiteral.1      0x0084 # Select bits 3 and 8
and                 # If port group bits 3 & 8 are
                    # common with SPGM, then
non-zero            # value is pushed onto stack
pushLiteral.1      0 # Push zero
eq                  # Only if SPGM is not in port
                    # groups corresponding to bits
                    # 3 & 8, then packet is
forwarded

```

Common Syntax Errors

When a packet filter definition is loaded, the definition is checked for syntax errors. The syntax errors and their causes are listed in Table A-1.

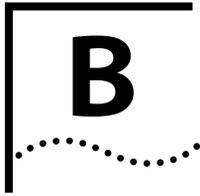
Table A-1 Syntax Errors When Loading Packet Filters

Syntax Error	Description
Opcode not found	An opcode was expected on the line and was not found. The opcode must be one of those identified above and must include the size, if any. The opcode and size must be separated by a single "." with no intervening spaces. Any mix of upper and lower case characters is permitted.
Unknown opcode	An opcode was expected on the line and was not found. The opcode must be one of those identified above and must include the size, if any. The opcode and size must be separated by a single "." with no intervening spaces. Any mix of upper and lower case characters is permitted.
Operands are not the same size	The opcode requires two operands of the same size. The top two operands currently on the stack are of different sizes.
Stack underflow	The opcode requires one or more operands. An insufficient number of operands are currently on the stack.
Stack overflow	The opcode pushes an operand on the stack. The stack does not have sufficient room for the operand.
No result found on top of stack	The program must end with a byte operand on the top of the stack. After the last instruction in the program is executed, the stack is either empty or contains an operand other than a byte.
Extra characters on line	The source line contains extraneous characters that are not part of the instruction and are not preceded by a comment character (#).
Expected a byte operand	The opcode requires a byte operand as one of its parameters. The operand is of a size other than a byte.
Offset not found	The opcode requires an offset to be specified. None was found on the line.
Literal not found	The opcode requires a literal value to be specified. None was found on the line.
String not found	The opcode requires a quoted string to be specified. None was found on the line.

(continued)

Table A-1 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> <p>A number with a leading "0x" or "0X" is treated as hexadecimal; a number with a leading 0 is treated as octal; all other numbers are treated as decimal.</p>
Number is too large	<p>The number specified as an offset or literal is too large. An offset is limited to 1518 less 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.

On-line Technical Services

3Com offers worldwide product support seven days a week, 24 hours a day, through the following on-line systems:

- 3Com Bulletin Board Service (3ComBBS)
- World Wide Web site
- 3ComForum on CompuServe®
- 3ComFactsSM automated fax service

3Com Bulletin Board Service

3ComBBS contains patches, software, and drivers for all 3Com products, as well as technical articles. This service is available via modem or ISDN seven days a week, 24 hours a day.

Access by Modem

To reach the service by modem, set your modem to 8 data bits, no parity, and 1 stop bit. Call the telephone number nearest you:

Country	Data Rate	Telephone Number
Australia	up to 14400 bps	(61) (2) 9955 2073
France	up to 14400 bps	(33) (1) 69 86 69 54
Germany	up to 9600 bps	(49) (89) 627 32 188 or (49) (89) 627 32 189
Hong Kong	up to 14400 bps	(852) 2537 5608
Italy (fee required)	up to 14400 bps	(39) (2) 273 00680
Japan	up to 14400 bps	(81) (3) 3345 7266
Singapore	up to 14400 bps	(65) 534 5693
Taiwan	up to 14400 bps	(886) (2) 377 5838
U.K.	up to 28800 bps	(44) (1442) 278278
U.S.	up to 28800 bps	(1) (408) 980 8204

Access by ISDN

ISDN users can dial-in to 3ComBBS using a digital modem for fast access up to 56 Kbps. To access 3ComBBS using ISDN, dial the following number:

(408) 654-2703

World Wide Web Site

Access the latest networking information on 3Com's World Wide Web site by entering our URL into your Internet browser:

<http://www.3Com.com/>

This service features news and information about 3Com products, customer service and support, 3Com's latest news releases, selected articles from 3TECH™ (3Com's award-winning technical journal) and more.

3ComForum on CompuServe

3ComForum is a CompuServe-based service containing patches, software, drivers, and technical articles about all 3Com products, as well as a messaging section for peer support. To use 3ComForum, you need a CompuServe account.

To use 3ComForum:

- 1** Log on to CompuServe.
- 2** Enter **go threecom** .
- 3** Press [Return] to see the 3ComForum main menu.

3ComFacts Automated Fax Service

3Com Corporation's interactive fax service, 3ComFacts, provides data sheets, technical articles, diagrams, and troubleshooting instructions on 3Com products 24 hours a day, seven days a week.

Call 3ComFacts using your touch-tone telephone. International access numbers are:

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

Local access numbers are available within the following countries:

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

Support from Your Network Supplier

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

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

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

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

Support from 3Com

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

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

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

Country	Telephone Number	Country	Telephone Number
Australia (Sydney)	(61) (2) 959 3020	Japan	(81) (3) 3345 7251
(Melbourne)	(61) (3) 653 9515	Mexico	(525) 531 0591
Belgium*	0800 71429	Netherlands*	06 0227788
Brazil	(55) (11) 546 0869	Norway*	800 13376
Canada	(905) 882 9964	Singapore	(65) 538 9368
Denmark*	800 17309	South Africa	(27) (11) 803 7404
Finland*	0800 113153	Spain*	900 983125
France*	05 917959	Sweden*	120 795482
Germany*	0130 821502	Taiwan	(886) (2) 577 4352
Hong Kong	(852) 868 9111	United Arab Emirates	(971) (4) 349049
Ireland*	1 800 553117	U.K.*	0800 966197
Italy*	1678 79489	U.S.	(1) (408) 492 1790

* These numbers are toll-free.

Returning Products for Repair

A product sent directly to 3Com for repair must first be assigned a Return Materials Authorization (RMA) number. A product sent to 3Com without an RMA number will be returned to the sender unopened, at the sender's expense.

To obtain an RMA number, call or fax:

Country	Telephone Number	Fax Number
U.S. and Canada	(800) 876 3266, option 2	(408) 764 7120
Europe	31 30 60 29900, option 5	(44) (1442) 275822
Outside Europe, U.S., and Canada	(1) (408) 492 1790	(1) (408) 764 7290

INDEX

Numerics

- 3Com Bulletin Board Service (3ComBBS) B-1
- 3Com sales offices B-4
- 3ComFacts B-3
- 802.1d bridging
 - enabling mode 1-7, 12-4
 - See also* bridge and STP

A

- abort
 - at prompts 2-10
 - enabling CTL-C 1-3, 2-12
- accept opcode 14-8, A-7
- access levels 2-1
- address
 - adding static 13-13
 - aging timer 1-7
 - filters A-9
 - flushing 13-14
 - for SNMP trap reporting 3-25
 - freezing 13-15
 - in routing table 3-10
 - IP 3-5
 - IP to MAC, translating 3-14
 - maximum number in group 15-7
 - removing static 13-13
- address group
 - adding addresses 15-7 to 15-9
 - as filtering criteria 15-1
 - copying 15-7
 - creating 1-8, 15-4
 - deleting 15-6
 - displaying contents 15-3
 - listing 15-2
 - loading on ESM 15-11
 - removing addresses 15-9
 - used in packet filter 15-1
- address reporting threshold
 - setting 1-7
- Address Resolution Protocol. *See* ARP
- address threshold
 - setting for bridge 12-8
 - values 12-8
- addressThresholdEvent 12-8
- administer access example 2-2
- Administration Console
 - command strings, using 2-9
 - Control keys, enabling 1-3, 2-12
 - exiting 2-17
 - initial user access 2-2
 - interface parameters 2-11, 2-13
 - locking 1-3, 2-12
 - menu descriptions 2-3 to 2-8
 - menu hierarchy, moving up 2-10
 - menu options, selecting 2-9
 - password access 1-2, 2-1, 4-2
 - preventing disconnections 2-12
 - purpose 1-2
 - screen height, setting 1-3, 2-11
 - scripts 2-14
 - top-level menu 2-3
 - values, entering 2-10
- Administration Console, restart 2-13
- agent
 - multiple 1-5
 - single 1-5
- aging time
 - setting for bridge 12-8
 - values 12-8
- analyzer
 - connecting 1-7, 11-4
 - MAC address display 11-4
 - removing port 11-5
 - setting up monitored port 1-7, 11-6
- and (bit-wise AND) opcode A-6
- Appletalk
 - packet filter 14-9
- ARP
 - defined 3-14
 - See also* ARP cache 3-14
- ARP cache
 - displaying contents 3-14
 - flushing 3-15
 - maintaining 1-4
 - removing entry 3-14

ASCII-based editor
 and scripts 2-14
 for packet filters 14-18

B

backplane paths mode
 assigning MACs to stations 9-25
 assigning ports to stations 9-32
 changing to multi-station 9-2
 configuring 9-3
 defining 9-1, 9-2
 backup
 saving NV data 7-2
 baseline
 displaying current 6-2
 enabling and disabling 6-2
 setting 1-2, 6-2
 why set one 6-1
 baud rate
 serial port 1-4
 bell, warning 4-2
 blocking state 13-5
 bridge
 address threshold, setting 1-7, 12-8
 aging time, setting 1-7, 12-8
 designated 13-3
 IP fragmentation, enabling 12-6
 IPX Snap Translation, enabling 12-6
 menus 2-6
 mode, defined 12-3
 mode, setting 1-7, 12-4
 Spanning Tree
 bridge priority, setting 12-10
 enabling 12-9
 forward delay, setting 12-12
 hello time, setting 12-11
 maximum age, setting 12-10
 Spanning Tree configurations 1-8
 statistics, displaying 12-1
See also packet filter, ESM, and EFSM
 bridge port
 MAC addresses
 adding 13-13
 flushing 13-14
 freezing 13-15
 listing 13-12
 removing 13-13
 multicast limit, setting 13-7
 Spanning Tree
 configuring 1-8
 enabling 13-8

path cost, setting 13-9
 port priority, setting 13-10
 states defined 13-5
 static address configuration 1-8
 statistics, displaying 13-1 to 13-3
 broadcast address 3-6
 bulletin board service B-1

C

code. *See* scripting *and* packet filter
 command
 and entering values 2-10
 using 2-9
 community strings
 and multiple agent mode 3-20, 3-22
 setting 1-5, 3-22
 values 3-21
 CompuServe B-2
 connectPolicy
 configuring 9-8
 Control keys
 enabling 1-3, 2-12
 Control Panel
 locking 1-3
 write access 2-13
 conventions
 notice icons 3
 cost
 of IP interface 3-6
 Spanning Tree settings 12-4, 13-3, 13-9
See also metric
 CTL-C (abort) 1-3, 2-12
 CTL-X (reboot) 1-3, 2-12

D

DAS configuration 9-34
 datagrams, statistics 3-17
 date
 formats 4-5
 setting 1-3, 4-4
 default route
 configuring 1-4
 defined 3-10
 removing 3-13
 setting 3-13
 destination address
 for SNMP trap reporting 3-26
 destination address group mask (DAGM) 15-1
 destination IP address
 in routing table 3-10

destination port group mask (DPGM) 15-1
 direct, route status 3-10
 documentation
 for the LANplex system 4
 DOS
 copying software to 5-3
 software media 5-1
 dual homing configuration 9-34

E

editor
 for packet filters 14-18
 for scripts 2-14
 EFSM
 and FDDI 9-3
 packet filter storage 14-1
 EMACS 2-14, 14-18
 eq opcode A-4
 ESM
 and FDDI 9-3
 packet filter storage 14-1
 Ethernet
 analyzing segments 11-1 to 11-7
 fragmenting packets 12-6
 menus 2-4
 packet fields 14-6
 port configurations 1-5
 portState 8-8
 station MAC addresses 13-12
 Ethernet address
 and restoring NV data 7-4
 for the monitored port 11-6
 Ethernet port
 analyzer attached 11-4
 displaying information 8-1
 label 8-4
 labeling 1-5, 8-7
 monitoring activity 1-7
 setting state (on-line or off-line) 8-8
 static MAC addresses 13-13
 statistics 8-3
 See also Roving Analysis
 Ethernet Switching Module. *See* ESM
 Ethernet/FDDI Switching Module. *See* EFSM

F

fan, warning 4-2
 fax service. *See* 3ComFacts
 FCM
 and FDDI 9-3

FDDI
 about configuring resources 9-1
 backplane paths mode 9-1, 9-2, 9-3
 configurations 1-5, 1-6, ?? to 1-7
 dual homing configuration 9-34
 fragmenting packets 12-6
 MAC configurations 1-6
 management 9-1
 menus 2-5
 packet fields 14-6
 path configurations 1-6
 port configurations 1-6
 port label 1-6, 9-30
 rings 9-11
 station configurations 1-5
 station MAC addresses 13-12
 thru configuration 9-34
 wrapped ring 9-11, 9-34
 FDDI MAC
 condition report 9-23
 configuring 1-6
 defined 9-16
 FrameErrorThreshold, setting 9-23
 LLC Service, enabling 9-25
 NotCopiedThreshold, setting 9-24
 path assignments 9-26
 station assignments 9-25
 statistics, displaying 9-16
 FDDI path
 defined 9-11
 maxT-Req, setting 9-15
 statistics, displaying 9-11
 tmaxLowerBound, setting 9-14
 tvxLowerBound, setting 9-13
 FDDI port
 and Roving Analysis 11-7
 configuring 1-6
 defined 9-28
 labeling 9-32
 lerAlarm, setting 9-30
 lerCutoff, setting 9-31
 path assignments 9-33
 station assignments 9-32
 statistics, displaying 9-28
 FDDI station
 and SMT 9-5
 and SRFs 9-6, 9-10
 connection policies, setting 9-8
 defined 9-5
 statistics, displaying 9-5
 status reporting, enabling 9-10
 T-notify, setting 9-9
 filter id 14-2

flash memory 5-1
 flushing
 ARP cache 3-15
 learned routes 3-12
 MAC addresses 13-14
 SNMP trap addresses 3-27
 forward delay 12-12
 forwarding state 13-5
 FrameErrorThreshold
 defined 9-23
 setting 9-23
 freezing addresses 13-15
 ftp
 IP address 3-2, 3-5
 server in software load 5-4

G

gateway
 configuring 1-4
 IP address 3-10
 See also route
 ge opcode A-6
 group. *See* address group or port group
 gt opcode A-6

H

hard disk
 copying software to 5-1
 hello time 12-11
 Help
 Administration Console 2-16
 outline form 2-16
 topical 2-16

I

ICMP
 and ping 3-16
 echo (request and reply) 3-16
 in-band management 3-2
 instructions
 opcodes 14-5, A-1
 operands 14-5, 14-7
 operators 14-7
 interface
 Administration Console parameters 2-11, 2-13
 defining 1-4, 3-7
 displaying 3-7
 parts of 3-5, 3-6
 removing definition 3-9

Internet Control Message Protocol. *See* ICMP
 Internet Protocol. *See* IP

IP
 address translation 3-14
 ARP cache 3-14
 interface 3-5
 management access 3-2
 menus 2-7
 pinging 1-4, 3-16
 RIP mode 1-4, 3-15
 routes 3-9
 statistics, displaying 3-17
 IP address
 and restoring NV data 7-4
 and software installation 5-4
 configuring 1-4
 for IP interface 3-5
 in routing table 3-10
 IP fragmentation
 enabling 12-6
 IP interface
 address 3-5
 broadcast address 3-6
 cost 3-6
 defining 1-4, 3-5 to ??, 3-7
 displaying 3-7
 removing definition 3-9
 subnet mask 3-6
 IP packets filter 14-12, 14-16
 IP route
 default 3-10, 3-13
 defining static 1-4, 3-11
 destination address 3-10
 displaying table 3-11
 gateway IP address 3-10
 metric 3-10
 removing from table 3-12
 status 3-10
 subnet mask 3-10
 IPX Snap Translation
 enabling 12-6

L

LANplex
 Administration Console purpose 1-2
 administration overview 1-1
 and network monitoring 11-1
 bell warning 4-2
 documentation 4
 fan warning 4-2
 naming 1-3, 4-4

- NV data restoration 7-3
 - power supply warning 4-2
 - rebooting 4-5
 - resetting to system defaults 1-3, 7-7
 - software installation 1-2
 - system backup 7-2
 - system configuration, displaying 4-1
 - system date and time 1-3, 4-4
 - system-level configurations 1-2 to ??, 1-3 to 1-5
 - temperature warning 4-2
 - user access levels 2-1
 - warning messages 4-2
 - LANplex Management Module. *See* LMM
 - le opcode A-5
 - learned, route status 3-10
 - learning state 12-12, 13-5
 - LER
 - alarm value 9-30
 - cutoff value 9-31
 - lerAlarm
 - and lerCutoff value 9-30
 - defined 9-30
 - setting 9-30
 - lerCutoff
 - and lerAlarm value 9-31
 - defined 9-31
 - Link Error Rate. *See* LER
 - listening state 12-12, 13-5
 - LLC
 - enabling 9-25
 - service description 9-25
 - LMM
 - and FDDI 9-3
 - lobe
 - port mode 10-9
 - Logical Link Control. *See* LLC
 - lt opcode A-5
-
- M**
- MAC (Media Access Control). *See* FDDI MAC
 - MAC address
 - adding 13-13
 - and ARP 3-14
 - configuring 13-12
 - displaying 13-12
 - dynamic to static 13-15
 - flushing 13-14
 - removing static 13-13
 - Roving Analysis configuration 11-3
 - management
 - and naming the system 4-4
 - and port labels 8-7, 9-32, 10-8
 - configuring system access 3-1 to 3-16
 - FDDI 1-5, 9-1
 - in-band 3-2
 - IP interface 1-4, 3-2, 3-5
 - out-of-band 3-2
 - serial connection 1-4
 - SNMP community strings 1-5, 3-21
 - Transcend LANplex Manager 1-2
 - maximum age 12-10
 - maxT-Req
 - defined 9-15
 - setting 9-15
 - media types 5-1
 - menu
 - analyzer (Roving Analysis) 2-8
 - and command strings 2-9
 - bridge 2-6
 - ethernet 2-4
 - fddi 2-5
 - IP 2-7
 - moving up hierarchy 2-10
 - outlining 2-16
 - selecting options 2-9
 - SNMP 2-8
 - system 2-4
 - Token Ring 2-6
 - metric
 - in routing table 3-10
 - modem 3-1
 - connecting to 3-4
 - connection to LMM 3-1
 - escape sequence 3-4
 - serial port speed 3-4
 - modem serial port
 - and RJ-12 cable 3-3, 3-4
 - setting baud rate 3-3
 - modules
 - and FDDI resources 9-3
 - revision numbers 4-2
 - multicast frames
 - and packet filters 14-1
 - multicast limit
 - configuring 13-7
 - defined 13-7
 - multiple agent mode
 - configuration guidelines 3-19
 - setting 3-21
 - multi-station mode
 - and FDDI resources 9-1
 - assigning MACs to stations 9-25
 - assigning ports to stations 9-32
 - configuring 9-3

defined 9-2

N

name opcode A-1
 naming the LANplex 1-3, 4-4
 ne opcode A-5
 neighbor notification
 and LLC Service 9-25
 network monitoring. *See* Roving Analysis *and* analyzer
 network supplier support B-3
 network troubleshooting 11-1
 not opcode A-7
 NotCopiedThreshold
 defined 9-24
 setting 9-24
 Novell
 in packet filter A-10
 NV data
 and packet filters 14-3
 backup 7-1
 contents saved 7-1
 examining a saved file 7-6
 file information 7-2
 resetting 1-3, 7-7
 restoring 1-3
 saving 1-3, 7-2
 transferring 7-1

O

off-line
 port state 8-8
 on-line
 port state 8-8
 on-line Help 2-16
 on-line technical services B-1
 opcode
 and packet filter language 14-4
 and writing packet filters 14-10
 descriptions A-1 to A-8
 operand 14-5
 and opcodes 14-7
 sizes supported 14-5
 operator
 and opcodes 14-7
 purpose 14-7
 or opcode A-7
 OUI
 in packet filter A-11
 out-of-band management 3-2

P

packet
 Ethernet type 14-6
 FDDI type 14-6
 fields for operands 14-7
 packet filter
 address group example 15-1
 assigning to ports 14-23
 basic elements 14-6
 concepts 14-4 to 14-11
 copying 14-23
 correcting errors 14-21
 creating 1-8, 14-3 to 14-18
 definitions 14-3
 deleting 14-20
 displaying contents 14-3
 editing 14-21
 editor
 commands 14-20
 description 14-18
 using 14-19
 example
 accept XNS or IP A-11
 address group A-12
 destination address A-9
 Ethernet Type IPX and Multicast A-10
 frame length field A-9
 frame type field A-10
 multiple destination address A-10
 port group A-12
 source address A-9
 source address and type A-11
 step by step 14-11 to 14-18
 XNS routing A-11
 external editor 14-20
 filter id 14-2
 filtering criteria, groups 15-1
 instructions 14-5
 language description 14-3, 14-4
 listing 14-2
 loading 14-22
 opcodes A-1
 operands 14-5
 port group example 15-2
 pre-processed storage 14-9
 procedure for writing 14-10
 processing paths 14-1, 14-23
 pseudocode 14-12
 run-time storage 14-10
 sequential tests 14-8
 stack 14-5
 storage space 14-9

syntax errors A-13, A-14
 unassigning from ports 14-24
See also address group *and* port group
 password
 configuring 1-2, 4-2
 initial system access 2-2
 levels of user access 2-1
 path cost
 defined 13-9
 setting 13-9
 path. *See* FDDI path *and* backplane paths
 peer connections 9-34
 PHY
 and FDDI ports 9-28
 ping
 IP station 1-4, 3-16
 PMD
 and FDDI ports 9-28
 port
 bridging priority 13-10
 for analyzer 11-4
 including in IP interface 3-6
 label 9-30
 maximum number in group 15-7
 path cost 13-9
 setting the state 8-8
 speed, setting 3-3
 types 9-28
 See also FDDI port
 port group
 adding ports 15-7 to 15-9
 as filtering criteria 15-1
 copying 15-7
 creating 1-8, 15-4
 deleting 15-6
 displaying contents 15-3
 listing 15-2
 loading on ESM 15-11
 removing ports 15-9
 used in packet filter 15-2
 power supply
 warning 4-2
 probe. *See* Roving Analysis *and* analyzer
 pushDAGM 1-8, 15-1, A-3
 pushDPGM 1-8, 15-1, A-4
 pushField.opcode A-2
 pushLiteral.opcode A-2
 pushSAGM 1-8, 15-1, A-3
 pushSPGM 1-8, 15-1, A-4
 pushTop opcode A-3

R

read access example 2-3
 reboot
 enabling CTL-X 1-3, 2-12
 resetting the system 4-5
 reboot system 2-13
 receive all
 packet processing path 14-1
 receive multicast
 packet processing path 14-1
 reject opcode 14-8, A-8
 Restart, Administration Console 2-13
 returning products for repair B-4
 RIP
 and broadcast address 3-6
 default mode 3-15
 displaying state 3-7
 Off mode 3-15
 Passive mode 3-15
 setting mode 1-4, 3-15
 RJ-12 cable
 for modem serial port 3-3, 3-4
 rlogin
 and exiting the Console 2-17
 and rebooting the system 4-5
 disconnection reason 2-12
 management access 3-2
 route
 default 3-10
 defining static 1-4, 3-11
 destination IP address 3-10
 flushing learned routes 3-12
 gateway IP address 3-10
 metric 3-10
 removing default 3-13
 removing from table 3-12
 status 3-10
 subnet mask 3-10
 router. *See* ESM *and* EFSM
 Routing Information Protocol. *See* RIP
 routing table
 contents 3-10
 default route, setting 3-13
 display routes 3-11
 flushing learned routes 3-12
 removing default route 3-13
 removing route 3-12
 Roving Analysis
 adding analyzer port 11-4
 and Spanning Tree 11-5
 configuration rules 11-3
 configuration, displaying 11-3, 11-4

- configuring 1-7
 - defined 11-1
 - menu 2-8
 - process overview 11-1
 - removing analyzer port 11-5
 - starting port monitoring 11-6
 - stopping port monitoring 11-7
 - See also* analyzer
-
- S**
- screen height
 - adjusting 2-11
 - script
 - examples 2-15
 - running 2-14
 - serial port
 - connecting to external
 - serial port speed
 - external connection 1-4
 - setting baud rate 1-4
 - serial port (modem)
 - setting baud rate 3-3
 - serial port (terminal)
 - and rebooting the system 4-5
 - disconnection reason 2-12
 - for management 3-1
 - setting baud rate 3-3
 - Service Access Points (SAPs)
 - and packet filters 14-4
 - shifl opcode A-8
 - shiftr opcode A-8
 - single agent mode
 - configuration guidelines 3-19
 - setting 3-21
 - single station mode
 - and FDDI resources 9-1
 - configuring 9-3
 - defined 9-2
 - SMT
 - and FDDI stations 9-5
 - lerAlarm value 9-30
 - lerCutoff value 9-31
 - SMT event
 - enabling proxying 3-28
 - proxying defined 3-27
 - Sniffer. *See* Roving Analysis *and* analyzer
 - SNMP
 - addressThresholdEvent trap 1-7
 - community strings
 - setting 1-5, 3-22
 - values 3-21
 - displaying configurations 3-18, 3-20
 - management modes 1-5
 - menus 2-8
 - mode
 - configuration guidelines 3-19
 - description 3-18
 - setting 3-21
 - proxying remote SMT events 3-28
 - trap reporting
 - and SMT event proxying 3-27
 - configuring destinations 1-5, 3-25
 - descriptions of traps 3-24
 - displaying configuration 3-23
 - flushing addresses 3-27
 - See also* trap *and* community strings
 - SNMP agent
 - accessing through IP 3-2
 - defined 3-18
 - single or multiple 3-18
 - SNMP trap
 - address reporting threshold 1-7
 - Address Threshold 3-24
 - addressThresholdEvent 12-8
 - Authentication Failure 3-24
 - Coldstart 3-24
 - MAC Duplicate Address Condition 3-24
 - MAC Frame Error Condition 3-24
 - MAC Neighbor Change 3-24
 - MAC Not Copied Condition 3-24
 - MAC Path Change 3-24
 - New Root 3-24
 - Port EB Error Condition 3-24
 - Port LER Condition 3-24
 - Port Path Change 3-24
 - Port Undesired Connection 3-24
 - Slot Extract 3-24
 - Slot Insert 3-24
 - Slot Overtemperature 3-24
 - SMT Hold Condition 3-24
 - SMT Peer Wrap Condition 3-24
 - System Overtemperature 3-24
 - Topology Change 3-24
 - socket values filter 14-12, 14-15
 - software
 - backup NV data 7-1, 7-2
 - build date and time 4-2
 - copying to hard disk 5-1
 - corrupted on install 5-5
 - from factory 1-1
 - installation 1-2, 5-1
 - loading time 5-4
 - version number 4-2
 - source address group mask (SAGM) 15-1

- source port group mask (SPGM) 15-1
 - Spanning Tree Protocol. *See* STP
 - SRF
 - and FDDI stations 9-6, 9-10
 - and IerAlarm 9-30
 - stack 14-5
 - static, route status 3-10
 - station
 - and FDDI backplane paths 9-1
 - port mode 10-9
 - See also* FDDI station
 - Station Management. *See* SMT
 - statistics
 - baselining 1-2, 6-1
 - Ethernet ports 8-3
 - FDDI MAC 9-17
 - FDDI path 9-11
 - FDDI port 9-28
 - FDDI station 9-6
 - IP 3-17
 - Status Report Frames. *See* SRF
 - status reporting
 - configuring 9-10
 - defined 9-10
 - STP
 - bridge priority, setting 12-10
 - designated bridge 13-3
 - designated cost 13-3
 - designated port 13-3
 - designated root 13-3
 - enabling on bridge 12-9
 - enabling on bridge port 13-8
 - forward delay, setting 12-12
 - hello time, setting 12-11
 - maximum age, setting 12-10
 - port priority 13-10
 - states 13-5
 - subnet mask
 - for IP address 3-6
 - in routing table 3-10
 - system configuration
 - displaying 4-1
 - system menus 2-4
-
- T**
- T_Opr
 - defined 9-15
 - TCP/IP
 - management access 3-2
 - technical support B-1
 - telnet
 - and rebooting the system 4-5
 - disconnection reason 2-12
 - management access 3-2
 - temperature, warning 4-2
 - terminal emulation
 - and the serial port 3-1
 - terminal serial port
 - and rebooting the system 4-5
 - disconnection reason 2-12
 - for management 3-1
 - setting baud rate 3-3
 - ThreeComForum B-2
 - thru ring configuration 9-34
 - time
 - formats 4-5
 - setting 1-3, 4-4
 - timing out, route status 3-10
 - tmaxLowerBound
 - defined 9-14
 - setting 9-14
 - T-notify
 - configuring 9-9
 - defined 9-9
 - token
 - and FDDI MAC 9-16
 - Token Ring
 - menus 2-6
 - port configurations 1-7
 - port label 1-7
 - portMode 10-9
 - portSpeed 10-9
 - portState 10-8
 - Token Ring port
 - label 10-4
 - labeling 10-8
 - setting mode (station or lobe) 10-9
 - setting state (on-line or off-line) 10-8
 - setting the speed 10-9
 - transmit all
 - packet processing path 14-1
 - transmit multicast
 - packet processing path 14-1
 - Transparent
 - enabling mode 1-7, 12-4
 - trap reporting
 - configuring destinations 1-5, 3-25
 - descriptions of traps 3-24
 - flushing addresses 3-27
 - tree connections 9-34
 - T-Req
 - defined 9-15

txvLowerBound
 defined 9-13
 setting 9-13

U

UNIX
 and terminal emulation with LANplex 3-1
 copying software to 5-2
 software media 5-1

V

vi 2-14, 14-18

W

warning messages for system 4-2
wrapped ring 9-11
write access example 2-2

X

XNS
 in packet filter 14-12, 14-14
 packet filter A-11
xor opcode A-7