ARTISAN®
TECHNOLOGY GROUP

Your **definitive** source for quality pre-owned equipment.

Artisan Technology Group
(217) 352-9330 | sales@artisantg.com | artisantg.com

Full-service, independent repair center
with experienced engineers and technicians on staff.

We buy your excess, underutilized, and idle equipment
along with credit for buybacks and trade-ins.

Custom engineering
so your equipment works exactly as you specify.

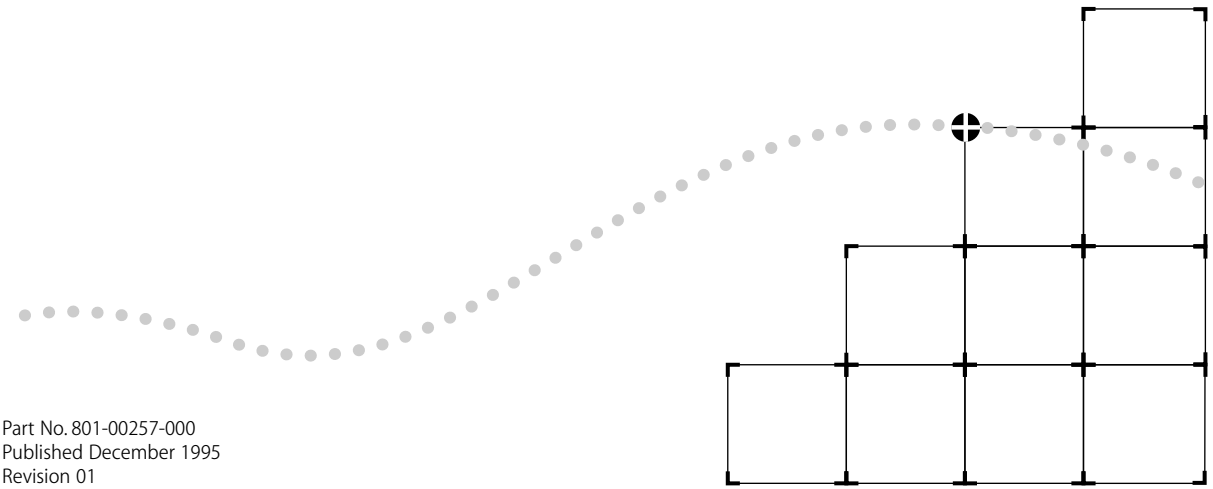
- Critical and expedited services
- In stock / Ready-to-ship
- Leasing / Rentals / Demos
- ITAR-certified secure asset solutions

Expert team | Trust guarantee | 100% satisfaction

All trademarks, brand names, and brands appearing herein are the property of their respective owners.

Find the *3COM LANplex 6012* at our website: ***Click HERE***

LANplex 6000 Extended Switching User Guide



Part No. 801-00257-000
Published December 1995
Revision 01

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's standard commercial agreement for the software. Unpublished rights reserved under the copyright laws of the United States.

3ComFacts, Ask3Com, CardFacts, NetFacts, and CardBoard are service marks of 3Com Corporation.

3Com and NETBuilder II are registered trademarks of 3Com Corporation.

LANplex and Transcend are trademarks of 3Com Corporation.

CompuServe is a registered trademark of CompuServe, Inc.

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

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

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

CONTENTS

ABOUT THIS GUIDE

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

PART I GETTING STARTED

1 LANPLEX EXTENDED SWITCHING FEATURES

About LANplex Extended Switching 1-1
Using Menus to Perform Tasks 1-2
- IP Menu 1-3
- IPX Menu 1-4
- AppleTalk Menu 1-5

2 INSTALLING EXTENDED SWITCHING SOFTWARE

About Installing Software 2-1
Copying Software to a Hard Disk 2-1
- Copying to UNIX 2-2
- Copying to DOS 2-3

Loading Software 2-4

PART II ABOUT ROUTING PROTOCOLS

3 ROUTING AND THE LANPLEX SYSTEM

What is Routing? 3-1
- LANplex in a Subnetted Environment 3-2
- Integrating Bridging and Routing 3-3

Bridging/Routing Models 3-4
- Traditional Bridging/Routing Model 3-4
- LANplex Bridging/Routing Model 3-6

4 ROUTING WITH IP

IP Routing and the OSI Model 4-1
The Elements of IP Routing 4-2
- IP Addresses 4-2
  - Address Classes 4-3
  - The Subnet Part of the IP Address 4-3
- Router Interfaces 4-4
- Routing Table 4-5
  - Static Routes 4-6
  - Dynamic Routes Using RIP 4-6
  - Default Route 4-7
- Address Resolution Protocol (ARP) 4-7

IP Routing Transmission Errors 4-9
IP Routing References 4-10

5 ROUTING WITH IPX

IPX Routing in the NetWare Environment 5-1
- Internet Packet Exchange (IPX) 5-2
- Routing Information Protocol (RIP) 5-3
- Service Advertising Protocol (SAP) 5-3

How IPX Routing Works 5-4
- IPX Packet Format 5-4
- IPX Packet Delivery 5-6
  - Sending Node's Responsibility 5-6
  - Router's Responsibility 5-7

The Elements of IPX Routing 5-8
- Router Interfaces 5-8
- Routing Tables 5-8
  - Generating Routes 5-9
  - Selecting the Best Route 5-10

Service Advertising Protocol (SAP) 5-10
Internetwork Service Information 5-10
SAP Packet Structure 5-11
Server Information Table 5-13
Server Information Maintenance 5-14

6 ROUTING WITH APPLETALK

About AppleTalk 6-1
AppleTalk Network Elements 6-1
AppleTalk Networks 6-2
AppleTalk Nodes 6-2
Named Entities 6-2
AppleTalk Zones 6-3
Seed Routers 6-4
AppleTalk Protocols 6-4
Physical Connectivity 6-5
The Datagram Delivery Protocol (DDP) 6-6
End-to-End Services 6-6
Transport Layer Protocols 6-6
The Session Layer Protocols 6-9
The Presentation Layer 6-10
About AARP 6-10

PART III ADMINISTERING ROUTING PROTOCOLS

7 ADMINISTERING IP ROUTING

Administering Interfaces 7-1
Displaying Interfaces 7-3
Defining an Interface 7-3
Modifying an Interface 7-4
Removing an Interface 7-5
Administering Routes 7-5
Displaying the Routing Table 7-6
Defining a Static Route 7-7
Removing a Route 7-8
Flushing a Route 7-8
Setting the Default Route 7-8
Removing the Default Route 7-9
Administering the ARP Cache 7-9
Displaying the ARP Cache 7-9

Removing an ARP Cache Entry 7-10
Flushing the ARP Cache 7-10
Administering UDP Helper 7-11
Displaying UDP Helper Information 7-11
Defining a Port and IP Forwarding Address 7-12
Removing a Port and IP Forwarding Address 7-12
Setting the Hop Count Limit 7-13
Setting the BOOTP Relay Threshold 7-13
Enabling/ Disabling IP Forwarding 7-13
Setting the RIP Mode 7-14
Pinging an IP Station 7-15
Displaying IP Statistics 7-16

8 ADMINISTERING IPX ROUTING
Administering Interfaces 8-2
Displaying IPX Interfaces 8-3
Defining an Interface 8-3
Modifying an Interface 8-4
Removing an Interface 8-4
Administering Routes 8-5
Displaying the Routing Table 8-6
Defining a Static Route 8-6
Removing a Route 8-7
Flushing Routes 8-7
Administering Servers 8-8
Displaying the Server Table 8-8
Defining a Static Server 8-9
Removing a Server 8-10
Flushing Servers 8-10
Setting IPX Forwarding 8-11
Setting the RIP Mode 8-11
Setting the Enhanced RIP Mode 8-12
Setting the SAP Mode 8-13
Displaying Statistics 8-14
Displaying IPX Summary Statistics 8-14
Displaying IPX RIP Statistics 8-15
Displaying IPX SAP Statistics 8-16
Displaying IPX Forwarding Statistics 8-17

9 ADMINISTERING APPLETALK ROUTING

Administering Interfaces 9-2
- Displaying AppleTalk Interfaces 9-3
- Defining an Interface 9-3
- Removing an Interface 9-4

Administering Routes 9-4
- Displaying the Routing Table 9-5
- Flushing all Routes 9-6

Administering the AARP Cache 9-6
- Displaying the AARP Cache 9-7
- Removing an Entry in the Cache 9-8
- Flushing All Cache Entries 9-8

Displaying the Zone Table 9-8
Configuring Forwarding 9-10
Configuring Checksum 9-10
Pinging an AppleTalk Node 9-11
Viewing Appletalk Statistics 9-11
- Displaying DDP Statistics 9-11
- Displaying RTMP Information 9-13
- Displaying ZIP Information 9-14
- Displaying NBP Information 9-16

PART IV APPENDIX

A TECHNICAL SUPPORT

On-line Technical Services A-1
- 3Com Bulletin Board Service A-1
  - Access by Modem A-1
  - Access by ISDN A-2
- World Wide Web Site A-2
- ThreeComForum on CompuServe A-2
- 3ComFacts Automated Fax Service A-2

Support from Your Network Supplier A-3
Support from 3Com A-4
Returning Products for Repair A-4

INDEX

About This Guide

Introduction

The *LANplex 6000 Extended Switching User Guide* provides information about the features included with the LANplex Extended Switching software. These features include IP, IPX, and AppleTalk routing.

You will use this guide with the *LANplex 6000 Administration Console User Guide* when you work with the Administration Console.

Audience description

This guide is intended for the system or network administrator who is responsible for configuring, using, and managing the LANplex 6000 system. It assumes a working knowledge of local area network (LAN) operations and a familiarity with communications protocols used on interconnected LANs.



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

How to Use This Guide

The following table shows where to find specific information.

If you are looking for...	Turn to...
An overview of Extended Switching features	Chapter 1
Information on how to install Extended Switching software	Chapter 2
An overview of routing in the LANplex system	Chapter 3
An overview of IP routing	Chapter 4
An overview of IPX routing	Chapter 5
An overview of AppleTalk routing	Chapter 6
Information on how to administer IP routing	Chapter 7
Information on how to administer IPX routing	Chapter 8
Information on how to administer AppleTalk routing	Chapter 9
Information on Technical Support	Appendix A

Conventions

Table 1 and Table 2 list conventions that are used throughout this guide.

Table 1 Notice Icons

Icon	Type	Description
	Information Note	Information notes call attention to important features or instructions.
	Caution	Cautions alert you to personal safety risk, system damage, or loss of data.
	Warning	Warnings alert you to the risk of severe personal injury.

Table 2 Text Conventions

Convention	Description
"Enter" vs. "Type"	When the word "enter" is used in this guide, it means type something, then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says "type."
"Syntax" vs. "Command"	When the word "syntax" is used in this guide, it indicates that the general form of a command syntax is provided. You must evaluate the syntax and supply the appropriate port, path, value, address, or string; for example: The following syntax specifies the time and date: `mm/dd/yy hh:mm:ss` When the word "command" is used in this guide, it indicates that all variables in the command have been supplied and you can enter the command as shown in text; for example: The following command enables Spanning Tree: **`bridge stpState enabled`**
Text represented as `screen display`	`This typeface` is used to represent displays that appear on your terminal screen, for example: `Login:`
Text represented as **`commands`**	**`This typeface`** is used to represent commands that you enter, for example: **`bridge stpState disabled`**

(continued)

Table 2 Text Conventions (continued)

Convention	Description
Keys	When specific keys are referred to in the text, they are called out by their labels, such as "the Return key" or "the Escape key," or they may be shown as [Return] or [Esc]. If two or more keys are to be pressed simultaneously, the keys are linked with a plus sign (+), for example: Press [Ctrl]+[Alt]+[Del].
Italics	*Italics* are used to denote *new terms* or *emphasis*.

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, configuring management access to, 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, bridging, FDDI technology, and Token Ring technology. 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 and provides information about using the Administration Console to configure and manage your routing protocols. (shipped with the option package/Part No. 801-257-000)

- *Command Quick Reference for the 6000 Administration Guide*

  Contains all of the Administration Console commands for the LANplex system. (Shipped with the system/Part No. 801-000258-000)

- *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. 801-00255-000)

- *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 our documentation more useful to you. Please email comments about this document to 3Com at: **sdtechpubs_comments@3Mail.3Com.com**

Please include the following information when commenting:

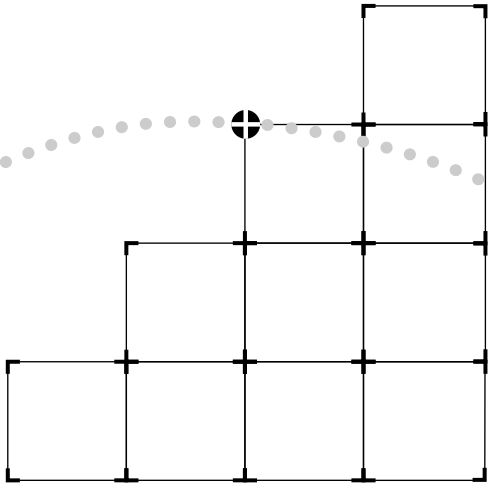
- Document title
- Document part number (listed on back cover of document)
- Page number (if appropriate)

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

I GETTING STARTED

Chapter 1 LANplex Extended Switching Features

Chapter 2 Installing Extended Switching Software



1 LANplex Extended Switching Features

This chapter provides an overview of the Extended Switching software, and describes the new enhanced Administration Console menus.

About LANplex Extended Switching

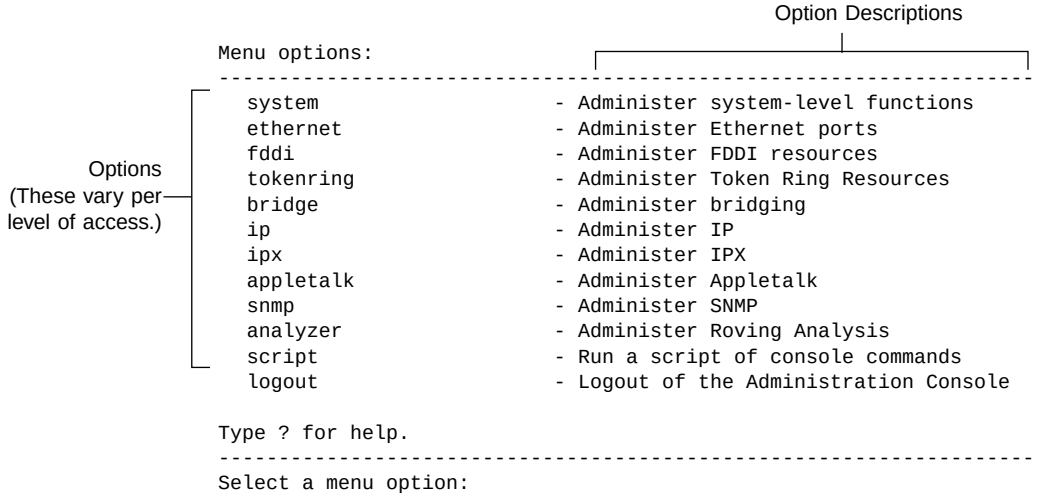
The LANplex Extended Switching software replaces your existing LANplex software and adds new functionality to your system. Extended Switching software contains all the features of standard LANplex software, in addition to routing capabilities with the following protocol support:

- IP Routing (an enhanced version of IP from the standard system software)
- IPX Routing
- AppleTalk Routing

For information on how to gain access to online help, to use scripts, and to exit from the Administration Console, see the *LANplex 6000 Administration Console User Guide*.

Using Menus to Perform Tasks

When you gain access to the Administration Console, the top-level menu appears. The Extended Switching software contains two new top-level menus (IPX and AppleTalk) and enhancements to the IP menu option:



The following sections show the new and enhanced menus provided with Extended Switching software. All other menu items appear in the *LANplex 6000 Administration Console User Guide*.

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), UDP Helper, IP Forwarding, and ping IP stations. See Figure 1-1. 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.

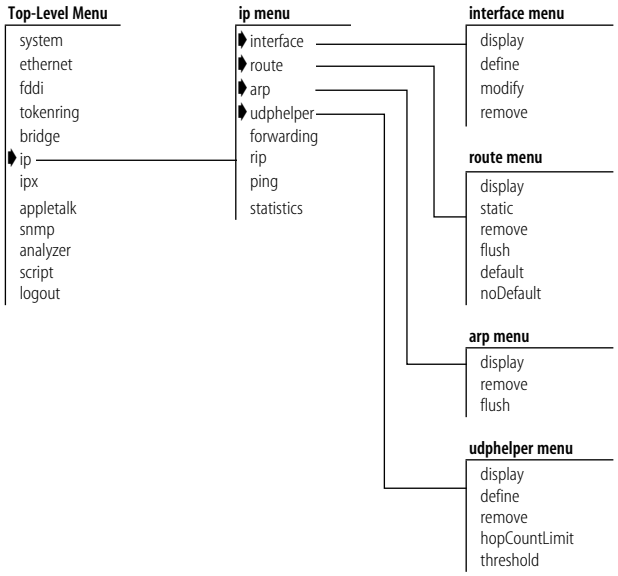


Figure 1-1 IP Menu Hierarchy

IPX Menu

From the **ipx** menu, you can view information about and configure Internet Packet Exchange (IPX) interfaces, routes, and servers. You can also administer the Routing Information Protocol (RIP), Enhanced RIP mode, Service Advertising Protocol (SAP), and statistics. See Figure 1-2. For example, to define a new IPX interface, you would enter **ipx** at the top-level menu, **interface** at the ipx menu, then **define** at the interface menu.

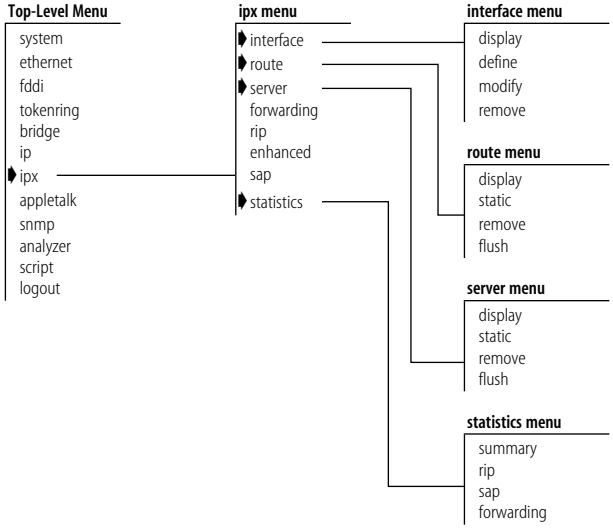


Figure 1-2 IPX Menu Hierarchy

AppleTalk Menu

From the **appletalk** menu, you can view information about and configure Appletalk interfaces, routes, and zones. You can also administer the Appletalk Address Resolution Protocol (AARP), Appletalk forwarding, and statistics. See Figure 1-3. For example, to define a new appletalk interface, you would enter **appletalk** at the top-level menu, **interface** at the AppleTalk menu, then **define** at the interface menu.

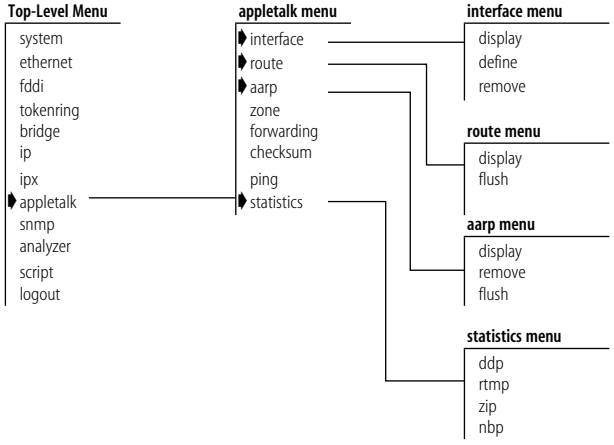


Figure 1-3 AppleTalk Menu Hierarchy

2 Installing Extended Switching Software

This chapter explains how to install Extended Switching software onto your system.

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

About Installing Software

When you upgrade to the Extended Switching Software, all configuration information is preserved. You can install a new version from any host running ftp.

CAUTION: *To run LANplex Extended Switching Software, you must have the LANplex Management Module Plus (LMM+) installed on your system. This new software does not run on the original 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/lp6000R" 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 into a disk drive (these instructions assume drive fd0).
2. Extract the first part of the LANplex software file using the following commands:

```
cd /usr/lp6000R
tar xvf /dev/rfd0
```

3. Remove diskette #1 using the following command:

```
# eject
```

4. Insert diskette #2 into a disk drive and extract the second part of the file using the following commands:

```
tar xvf /dev/rfd0
```

5. Remove diskette #2 using the following command:

```
# eject
```

The following files should be in your current default directory:

- README1
- lp6000R00
- lp6000R01
- lp6000R02
- restore_lpxR

6 Use the supplied script to decompress and restore the split file (`lp6000R00, lp6000R01,` and `lp6000R02`).

```
# ./restore_lpxR
```

See the `README1` 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.

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

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

If the directory "lp6000R" does not exist on your computer, create the directory before proceeding.

1 Insert diskette #1 into a disk drive (these instructions assume drive B:).

2 Copy the system software file to the directory of your computer using the following commands:

```
cd lp6000R
copy b:lp6000R.exe
```

The file lp6000R.exe is a self-extracting archive. It decompresses and creates the loadable lanplex file.

3 Decompress the file using the following command:

```
lp6000R
```

This creates a file called **`lp6000R`**, which you can then load into flash memory.

Loading Software

Before loading the system software on the LMM+, you must 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* in the *LANplex 6000 Administration Console User Guide*.

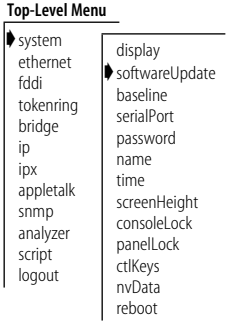
You can load the system software into flash memory while the system is operating. You do not need to bring the system down. After the flash install is completed, a reboot will put the newly-loaded software to use.

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:



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. 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. 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:

   `/c:\lp6000R\lp6000R`

4. Enter your user name.

5. 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/lp6000R/lp6000R]:
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 Flash is corrupted (for example, when a power failure occurs while you are updating software), contact 3Com Technical Support, as described in Appendix A.

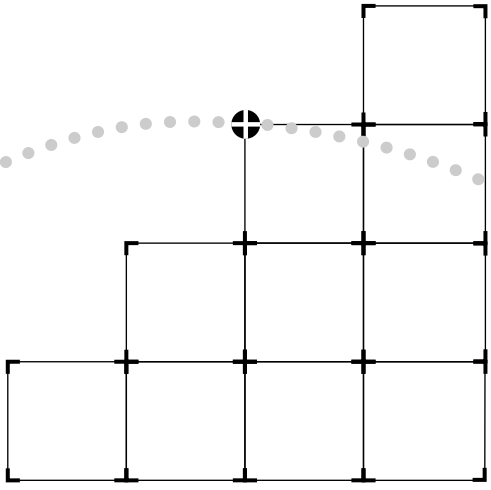
II About Routing Protocols

Chapter 3 Routing and the LANplex System

Chapter 4 Routing with IP

Chapter 5 Routing with IPX

Chapter 6 Routing with AppleTalk



3 Routing and the LANplex System

This chapter shows how the LANplex system operates in a subnetted routing environment and describes the LANplex routing methodology — specifically, how the LANplex bridging and routing model compares with traditional models.

What is Routing?

Routing is the process of distributing packets over potentially dissimilar networks. A router (also called a gateway) is the machine that accomplishes this task. Routers are typically used to:

- Connect enterprise networks together
- Connect subnets (client/server networks) to the enterprise network

Figure 3-1 shows where routes are typically used in a network.

The LANplex system performs routing that connects subnets to the enterprise network, providing connectivity between devices within a workgroup, department, or building.

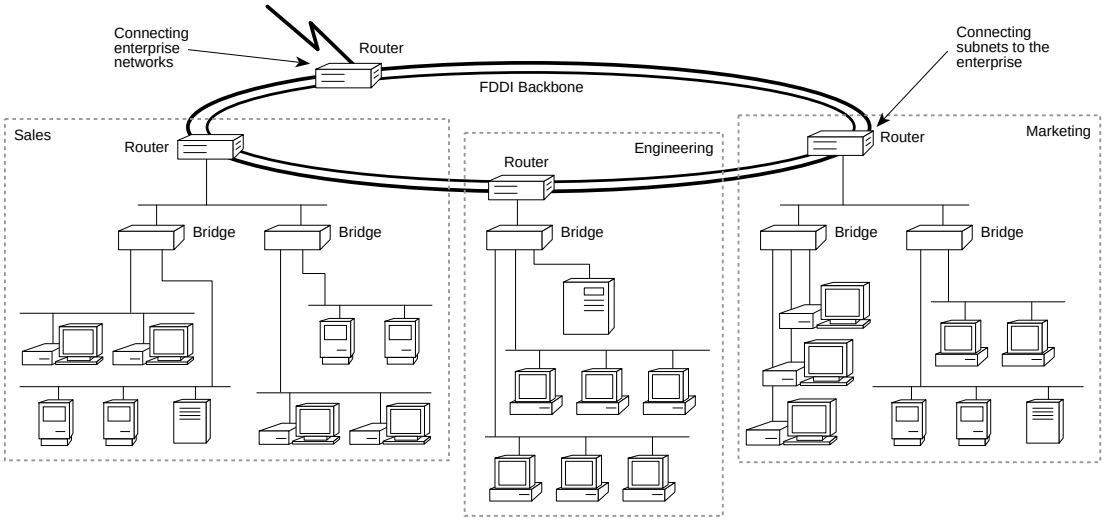


Figure 3-1 Traditional Architecture of a Routed Network

LANplex in a Subnetted Environment

The LANplex system allows you to fit Ethernet switching capability into highly subnetted environments. When you put the LANplex system into a subnetted network, it streamlines your network architecture and easily switches traffic between and within subnets over Ethernet and FDDI. See Figure 3-2.

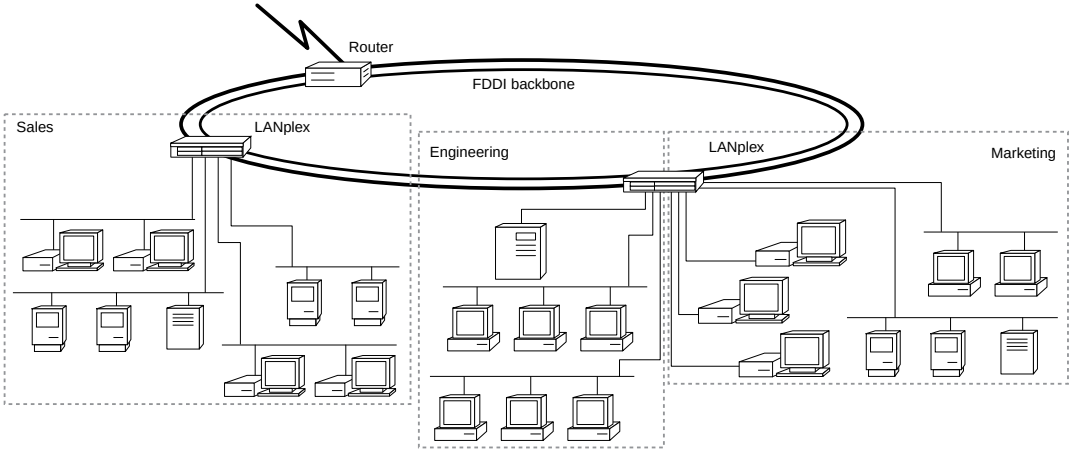


Figure 3-2 Subnetted Architecture with LANplex Switching Hubs

Integrating Bridging and Routing

The LANplex system has bridging and routing integrated into the Ethernet/FDDI Switching Module (EFSM) and the Ethernet Switching Module (ESM). The Token Ring Switching Module (TRSM) supports bridging only.

Multiple switch ports can be assigned to each subnet. See Figure 3-3. Traffic between ports assigned to the same subnet is switched transparently using transparent bridging or Express switching (described in the *LANplex 6000 Operation Guide*). Traffic traveling to different subnets is routed using one of the supported routing protocols.

In the following descriptions of bridging and routing on the LANplex system, the term **MAC address** *refers to a physical hardware address. The term* **network address** *refers to a logical address that applies to a specific protocol.*

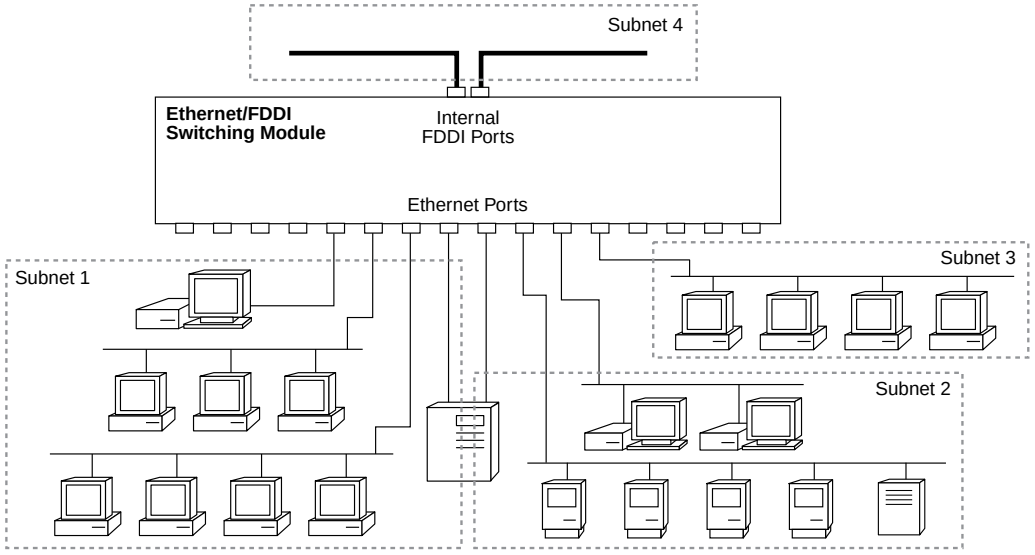


Figure 3-3 Multiple Ports per Subnets with the EFSM

Because the LANplex model of bridging and routing allows several segments to be connected to the same subnet, you can increase the level of segmentation in your network without having to create new subnets or assign network addresses. Instead, you can use additional

Ethernet ports to expand your existing subnets. This is in contrast to more traditional forms of bridging and routing where, at most, one port is connected to any subnet.

In the traditional model, if you want to increase the level of segmentation in your network, you must create additional subnets and assign new network addresses to your existing hosts.

Bridging/Routing Models

The way routing is implemented in the LANplex system differs from how bridging and routing usually coexist in a system.

- **Traditional Bridging/Routing Model** — In this model, bridging and routing are peer entities; either a packet is bridged or routed. Packets belonging to recognized protocols are routed; all others are bridged.
- **LANplex Bridging/Routing Model** — In this model, the bridge and router operate hierarchically on the module — routing over bridging. When a packet enters the system, the module first tries to bridge the packet. If the packet's destination network address is not on the same subnet, then the module routes the packet.

Traditional Bridging/Routing Model

The bridge/router determines whether a packet should be bridged or routed based on the protocol to which the packet belongs. If the packet belongs to a recognized protocol, the packet is routed. Otherwise, it is bridged.

In the traditional bridging/routing model, a packet is *bridged* as follows (see Figure 3-4):

1. The packet enters the bridge/router.
2. The bridge/router determines that the packet does not belong to a recognized routed protocol, so the packet is passed to the bridge.
3. The bridge examines the destination MAC address and forwards the packet to the port on which that address has been learned.

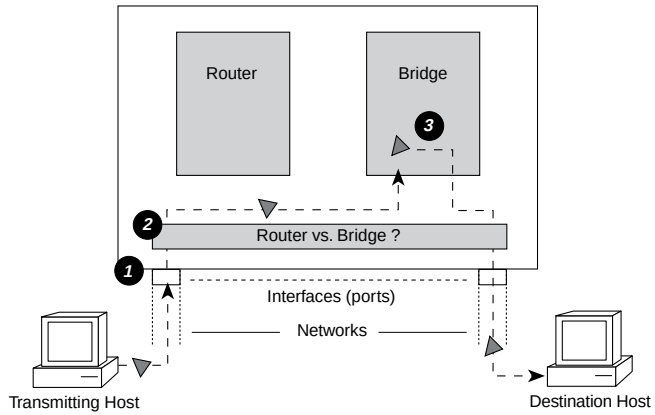


Figure 3-4 Bridging in the Traditional Bridging/Routing Model

In the traditional bridging/routing model, a packet is *routed* as follows (see Figure 3-5):

1 The packet enters the bridge/router.

2 The bridge/router determines that the packet belongs to a recognized routed protocol, so the packet is passed to the router.

3 The router examines the destination network address and forwards the packet to the interface (port) connected to the destination subnet.

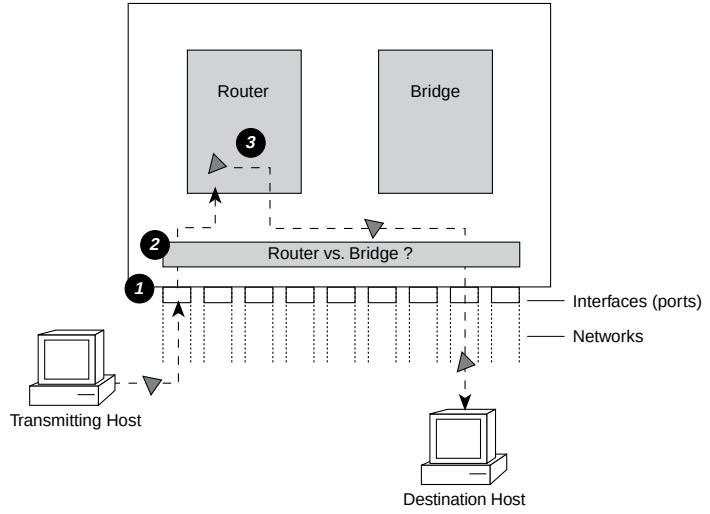


Figure 3-5 Routing in the Traditional Bridging/Routing Model

LANplex Bridging/Routing Model

The LANplex 6000 determines whether a packet should be bridged or routed using the destination MAC address. Before a host sends a packet to another host, it compares its own network address to the network address of the other host as follows:

- If network addresses are on the same subnet, the packet is bridged directly to the destination host's address.
- If network addresses are on different subnets, the packet must be routed from one subnet to the other. In this case, the host transmits the packet to the connecting router's MAC address.

In the LANplex bridging/routing model, a packet is *bridged* as follows (see Figure 3-6):

1. The packet enters the module.
2. The packet's destination MAC address is examined by the bridging layer.
3. The destination MAC address does not correspond to the MAC address of one of the module ports configured for routing. The bridging layer selects a segment (port) based on the destination MAC address and forwards the packet to that segment.

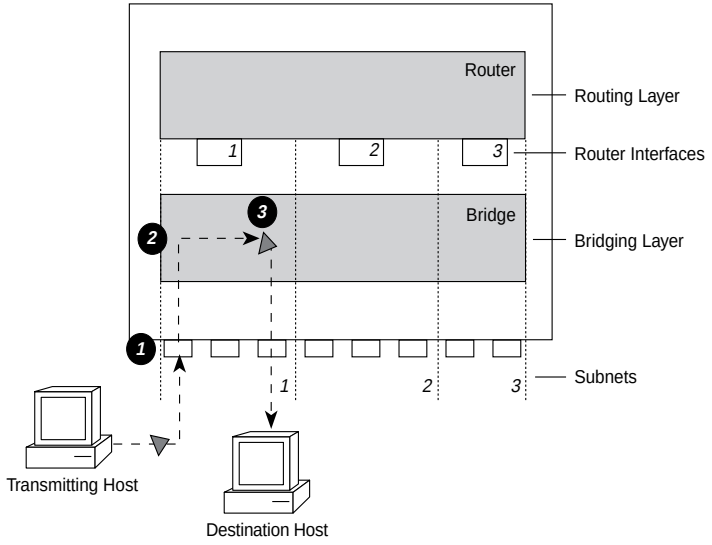


Figure 3-6 Bridging and the LANplex Bridging/Routing Model

In the LANplex bridging/routing model, a packet is *routed* as follows (see Figure 3-7):

1 The packet enters the module.

2 The packet's destination address is examined by the bridging layer.

3 The destination address corresponds to the address of one of the module ports configured for routing (as opposed to a learned end-station address). The packet is passed to the router interface associated with the port on which the packet was received.

4 The routing layer:

- **a** Selects a destination interface based on the destination network address.
- **b** Determines the MAC address of the next hop (either the destination host or another gateway).
- **c** Passes the packet back to the bridging layer.

5 The bridging layer then selects a segment (port) based on the destination MAC address and forwards the packet to that segment.

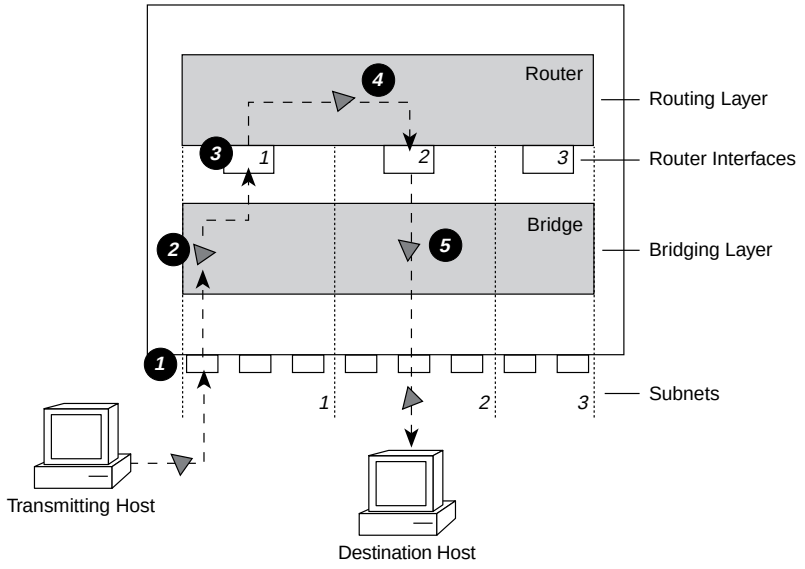


Figure 3-7 Routing in the LANplex Bridging/Routing Model

4 Routing with IP

This chapter gives an overview of IP routing, specifically defining:

- What IP routing involves
- What elements are necessary for IP routers to effectively transmit packets
- How IP routing transmission errors are detected and resolved

IP Routing and the OSI Model

An IP router, unlike a bridge, operates at the network layer of the OSI Reference Model. This means that it routes packets by examining the network layer address (IP address). Bridges use the data-link layer MAC addresses to make forwarding decisions. See Figure 4-1.

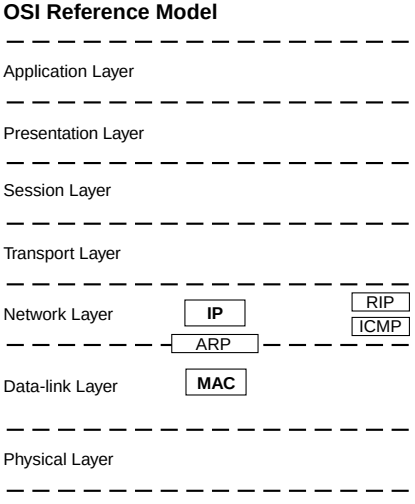


Figure 4-1 OSI Reference Model and IP Routing

When an IP router sends a packet over multiple physical networks, it does not know the complete path to a destination — only the next hop. Each hop involves the following:

- The IP routing algorithm computes the *next hop* IP address (the next router interface) using the routing table entries.
- ARP translates the next hop IP address into a physical MAC address.
- The router sends the packet over the network to the next hop.

These routing elements are described in more detail in the following section.

The Elements of IP Routing

IP Routers use the following elements to transmit packets in a subnetted environment:

- IP addresses
- Router interfaces
- Routing tables
- Address Resolution Protocol (ARP)

IP Addresses

IP addresses are 32-bit addresses composed of a *network part* (network on which the host is located) and a *host part* (the host on that network). See Figure 4-2. They differ from Ethernet and FDDI MAC addresses, which are unique hardware-configured 48-bit addresses.

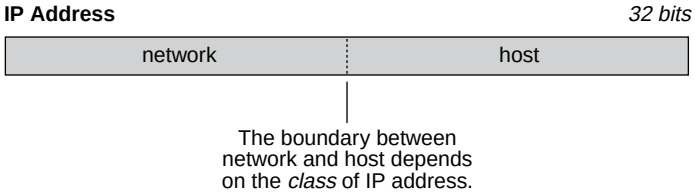


Figure 4-2 IP Address Network Part and Host Part

The IP address network part is assigned by a central agency, and the host part is assigned by each network's administrator. All devices connected to the same network share the same IP address prefix (the network part of the address).

Address Classes

The boundary of the network part and the host part depends on the class of network you are assigned by the central agency. The primary classes of IP addresses are Class A, Class B, and Class C.

- **Class A addresses** — have seven bits for the network part and 24 bits for the host part. Although only a few Class A networks can be created, each can contain a very large number of hosts.
- **Class B addresses** — have 14 bits for the network part and 16 bits for the host part.
- **Class C addresses** — have 21 bits for the network part and eight bits for the host part. Each Class C network can only contain about 250 hosts, but many such networks can be created.

The class of an IP address is designated in the high-order bits of the address.

The Subnet Part of the IP Address

In some environments, the IP address contains a *subnet part*. Subnetting allows a single Class A, B, or C network to be further subdivided internally while still appearing as a single network to other networks. The subnet part of the IP address is only visible to those hosts and gateways on the subnetted network.

When an IP address contains a subnet part, a *subnet mask* is used to identify which bits are the subnet and which are the host. A subnet mask is a 32-bit number that uses the same format and representation as IP addresses. Each IP address bit corresponding to a *one* in the subnet mask is in the network/subnet part of the address. Each IP address bit corresponding to a *zero* is in the host part of the IP address. See Figure 4-3.

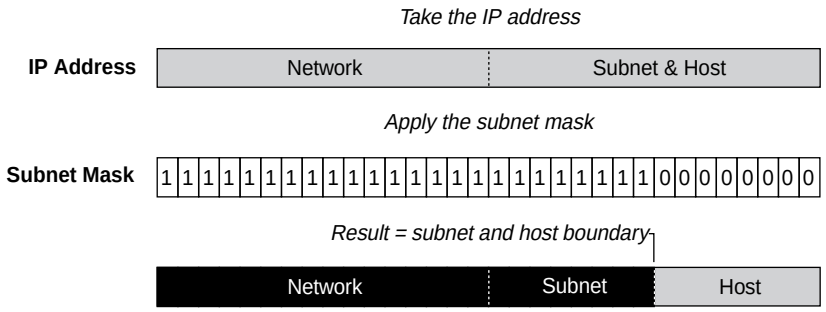


Figure 4-3 How a Subnet Mask is Applied to the IP Address

An example of an IP address that includes the network, subnet, and host parts is *158.101.230.52* with a subnet mask of *255.255.255.0*. This address is divided as follows:

- *158.101* is the network part
- *230* is the subnet part
- *52* is the host part

Router Interfaces

A router interface is the connection between the router and a subnet. In traditional routing models, the interface would be the same as the port, since only one interface could exist per port. In the LANplex system's IP routing, more than one port can be connected to the same subnet. Therefore, the router interface is the *relationship* between the ports and the subnets in your IP network.

Each router interface has an IP address and a subnet mask. This address defines both the number of the network the router interface is attached to and its host number on that network. A router interface's IP address serves two functions:

- The IP address is used when sending IP packets to or from the router itself.
- The IP address defines the network and subnet numbers of the segment connected to that interface. See Figure 4-4.

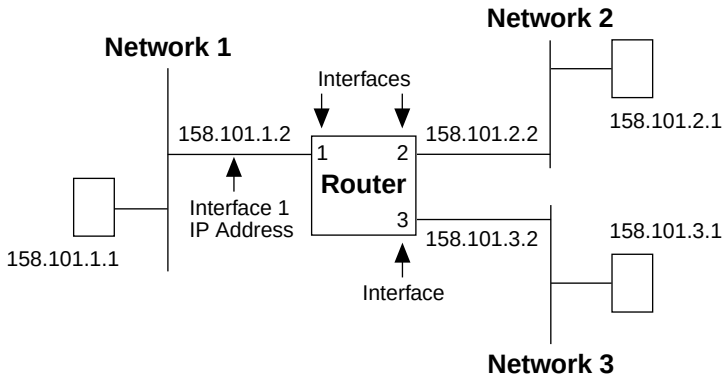


Figure 4-4 Router Interfaces

Routing Table

A routing table allows a router or host to determine how to send a packet toward the packet's ultimate destination. The routing table contains an entry for every destination network, subnet, or host to which the router or host is capable of forwarding packets. A router or host uses the routing table when the destination IP address of the packet it is sending is not on a network or subnet to which it is directly connected. The routing table provides the IP address of a router that can forward the packet toward its destination.

The routing table consists of the following:

- **Destination IP Address** — the destination network, subnet, or host
- **Subnet Mask** — the subnet mask corresponding to the destination IP address
- **Metric** — a measure of the "distance" to the destination (in the Routing Information Protocol (RIP), the metric is the number of hops)
- **Gateway** — the IP address of the next hop router (the IP address of the interface through which the packet travels)
- **Interface** — the interface number through which a packet must travel to reach that router

Figure 4-5 shows the routing table of the router in Figure 4-4.

Routing Table

Destination IP Address	Subnet Mask	Metric	Gateway	Interface
158.101.1.1	255.255.255.0	1	158.101.1.2	1
158.101.2.1	255.255.255.0	1	158.101.2.2	2
158.101.3.1	255.255.255.0	1	158.101.3.2	3
default route	255.255.255.0	1	158.101.1.2	1

Figure 4-5 Example of a Routing Table

Routing table information is generated and updated in the following ways:

- **Statically** — You manually enter routes, which do not change until you change them (that is, they will not time out).
- **Dynamically** — The router uses a routing protocol, such as RIP, to exchange information. Routes are recalculated at regular intervals.

Static Routes

A static route is one that you manually configure in the routing table. Static routes are useful in environments where no routing protocol is used, or where you want to override some of the routes generated with a routing protocol. Because static routes do not automatically change in response to network topology changes, you should manually configure only a small number of reasonably stable routes.

Dynamic Routes Using RIP

Automated methods of configuring routes help you keep up with a changing network environment, allowing routes to be reconfigured quickly and reliably. Interior Gateway Protocols (IGP), protocols that operate within networks, provide this automated method. The LANplex system uses RIP, one of the most widely used IGPs, to configure its routing tables dynamically.

RIP operates in terms of active and passive devices. The active devices, usually routers, broadcast their RIP messages to all devices in a network or subnet; they update their own routing tables when they receive a RIP message. The passive devices, usually hosts, listen for RIP messages and update their routing tables; they do not send RIP messages.

An active router sends a RIP message every 30 seconds. This message contains both the IP address and a metric (the distance to the destination from that router) for each destination. In RIP, each router that a packet must travel through to reach a destination equals one hop.

Default Route

In addition to the routes to specific destinations, the routing table may contain an entry called the *default route*. The router uses the default route to forward packets that do not match any other routing table entry. A default route is often used in place of routes to numerous destinations all having the same gateway IP address and interface number. The default route can be configured statically, or it can be learned dynamically using RIP.

Address Resolution Protocol (ARP)

ARP is a low-level protocol used to locate the MAC address corresponding to a given IP address. This allows a host or router to make its routing decisions using IP addresses while it uses MAC addresses to forward packets from one hop to the next.

Once the host or router knows the IP address of the *next hop* to the destination, the host or router must translate that IP address into a MAC address before the packet can be sent. To do this, the host or router first looks in its ARP cache, a table of IP addresses with their corresponding MAC addresses. Each device participating in IP routing maintains an ARP cache. See Figure 4-6.

ARP Cache

IP Address	MAC Address
158.101.1.1	00308e3d0042
158.101.2.1	0080232b00ab

Figure 4-6 Example of an ARP Cache

If the IP address does not have a corresponding MAC address listed, the host or router broadcasts an *ARP request* packet to all the devices on the network. The ARP request contains information about the hardware and protocol. The two key elements of the ARP request are the target and

source addresses for both the hardware (MAC addresses) and the protocol (IP addresses). See Figure 4-7.

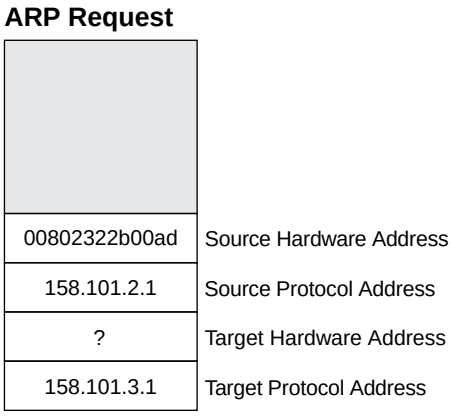


Figure 4-7 Example of an ARP Request Packet

When the devices on the network receive this packet, they examine it, and if their address is not the target protocol address, they discard the packet. When a device receives the packet and confirms that its IP address is the target protocol address, this device places its MAC address in the target hardware address field and sends the packet back to the source hardware address. When the originating host or router receives the *ARP reply*, it takes the new MAC address and places it in its ARP cache next to the corresponding IP address. See Figure 4-8.

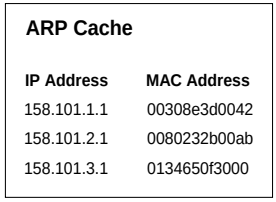


Figure 4-8 Example of ARP Cache Updated with ARP Reply

Once the MAC address is known, the host or router can send the packet directly to the next hop.

IP Routing Transmission Errors

Because each router only knows about the next hop, it is not aware of problems that may be further "down the road" toward the destination. Destinations can be unreachable if:

- Hardware is temporarily out of service
- You inadvertently specify a nonexistent destination address
- The router does not have a route to the destination network

To help routers and hosts know of problems in packet transmission, an error-reporting mechanism called Internet Control Message Protocol (ICMP) provides error reporting back to the source when routing problems arise. ICMP is a required part of IP. Without ICMP, you could not tell if a delivery failure resulted from a local or remote malfunction.

ICMP does the following:

- Tests the reachability of nodes (*ICMP Echo Request* and *ICMP Echo Reply*)

  A host or gateway sends an ICMP echo request to a specified destination. If the destination receives the echo request, it sends an ICMP echo reply back to the original sender. This process tests that the destination is reachable and responding, and verifies that the major pieces of the transport system work. The *ping* command is often used to invoke this process.

- Creates more efficient routing (*ICMP Redirect*)

  Many times the host route configuration specifies the minimal possible routing information needed to communicate (for example, the address of a single router). The host relies on routers to update its routing table. In the process of routing packets, a router may detect a host not using the best route. The router then sends the host an ICMP redirect, requesting that the host use a different gateway when sending packets to that destination. The next time the host sends a packet to that same destination, the host uses the new route.

- Informs sources that a packet has exceeded its allocated time to exist within the network (*ICMP Time Exceeded*)

IP Routing References

Comer, Douglas E. *Internetworking with TCP/IP Volume I: Principles, Protocols, and Architecture*. Englewood Cliffs, New Jersey: Prentice Hall, Inc., 1991.

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

Sterns, Richard. *TCP/IP Illustrated, Volume 1 The Protocols*. Addison-Wesley Professional Computing Services, 1992

RFC 791. *Internet Protocol Specification*.

RFC 792. *Internet Control Message Protocol Specification*.

RFC 1009. *Requirements for Internet Gateways*.

RFC 1042. A *Standard for the Transmission of IP Datagrams over IEEE 802 Networks*.

RFC 1058. *Routing Information Protocol*.

RFC 1122. *Requirements for Internet Hosts*.

5 Routing with IPX

This chapter provides an overview of IPX routing, specifically defining:

- What part IPX plays in the NetWare environment
- How IPX works
- What elements are necessary for IPX routers to transmit packets effectively

IPX Routing in the NetWare Environment

NetWare is a network operating system (NOS) developed and introduced to the market by Novell™, Inc. in the early 1980s. Much of the NetWare networking technology was derived from Xerox Network System (XNS)™, a networking system developed by Xerox Corporation™.

As a network operating system environment, NetWare specifies the upper five layers of the OSI reference model. It provides file and printer sharing and supports various applications such as electronic mail and database access. NetWare is based on a client/server architecture where clients request certain services from servers such as file and printer access.

Figure 5-1 illustrates a simplified view of NetWare's better-known protocols and their relationship to the OSI reference model.

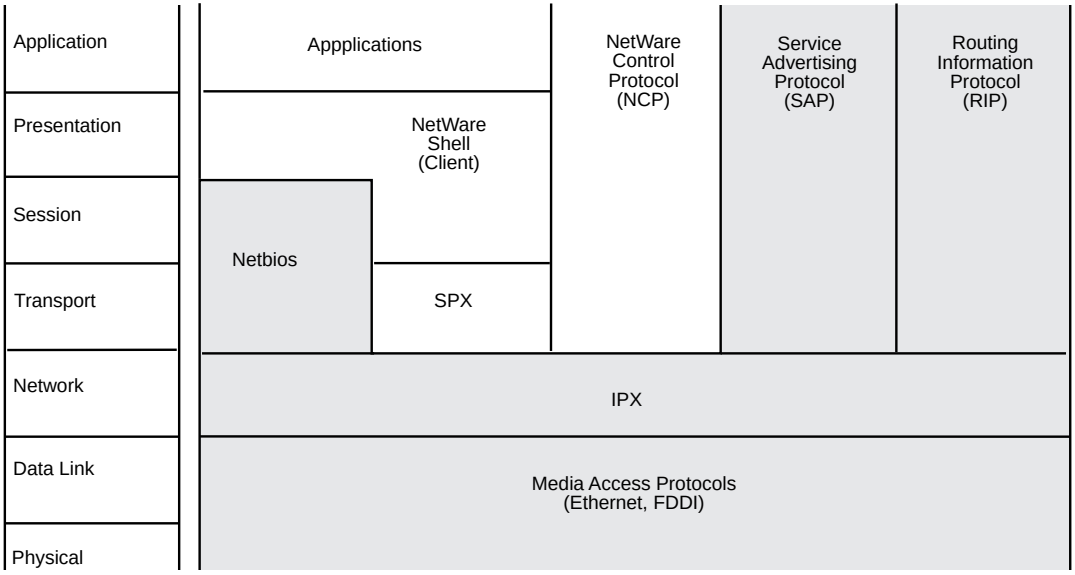


Figure 5-1 NetWare Protocols and the OSI Reference Model

The LANplex system uses the following protocols for routing in a Netware environment:

- Internet Packet Exchange (IPX)
- Routing Information Protocol (RIP)
- Service Advertisement Protocol (SAP)

Internet Packet Exchange (IPX)

IPX is the primary protocol used for routing in a netware environment. This datagram, connectionless protocol does not require an acknowledgment for each packet sent. Any packet acknowledgment, or connection control, must be provided by protocols above IPX.

IPX defines internetwork and intranode addressing schemes. IPX internetwork addressing is based on network numbers that are assigned to each interface in an IPX network. IPX intranode addressing is in the form of socket numbers. Since several processes are normally operating within a node, socket numbers provide a type of mail slot so that each process can distinguish itself to IPX.

Routing Information Protocol (RIP)

RIP allows the exchange of routing information on a NetWare network. IPX routers use RIP to dynamically create and maintain their routing tables.

RIP allows a router to exchange routing information with a neighboring router. As a router becomes aware of any changes in the network layout, it broadcasts this information to any neighboring routers. IPX routers also send periodic RIP broadcast packets containing all routing information known to the router. These broadcasts synchronize all routers on the network and age those networks that might become inaccessible due to a router going down abnormally.

Service Advertising Protocol (SAP)

SAP provides routers and servers (that contain SAP agents) with a means of exchanging network service information.

Through SAP, servers advertise their services and addresses. Routers gather this information and share it with other routers. This allows routers to dynamically create and maintain a database (server table) of network service information. Clients on the network can determine what services are available and obtain the network address of the nodes (servers) where they can access those services. Clients require this information to initiate a session with a file server.

SAP allows a router to exchange information with a neighboring SAP agent. As a router's SAP agent becomes aware of any change in the network server layout, it immediately broadcasts this information to any neighboring SAP agents. SAP broadcast packets containing all server information known to the SAP agent are also periodically sent. These broadcasts synchronize all servers on the network and age those servers that might become inaccessible due to any abnormal shut down of the router or server.

How IPX Routing Works

A router operates at the network layer of the OSI Reference Model. This means that it receives its instructions to route packets from one segment to another from a network layer protocol. IPX, with the help of RIP and SAP, performs these network layer tasks. These tasks include addressing, routing, and switching information packets to move single packets from one location to another. This section describes the information included in an IPX packet that helps it get delivered and the IPX packet delivery process.

IPX Packet Format

The IPX packet format consists of two parts: a 30-byte header and a data portion. The network, node, and socket address for both the destination and source are held within the packet's IPX header.

Figure 5-2 shows the IPX packet format.

Checksum	(2 Bytes)
Packet Length	(2 Bytes)
Transport Control (1 Byte)	Packet Type (1 Byte)
Destination Network	(4 Bytes)
Destination Node	(6 Bytes)
Destination Socket	(2 Bytes)
Source Network	(4 Bytes)
Source Node	(6 Bytes)
Source Socket	(2 Bytes)
Upper-layer Data	

Figure 5-2 IPX Packet Format

The packet format consists of the following elements:

- **Checksum** — The IPX packet begins with a 16-bit checksum field that is set to ones.
- **Packet Length** — This 16-bit field contains the length, in bytes, of the complete network packet. This includes both the IPX header and the data. The IPX length must be at least 30 bytes.
- **Transport Control** — This 1-byte field indicates how many routers a packet has passed through on its way to its destination. Packets are discarded when this value reaches 16. Sending nodes always set this field to zero when building an IPX packet.
- **Packet Type** — This 1-byte field specifies the upper-layer protocol to receive the packet's information.
- **Destination Network** — This 4-byte field provides the destination node's network number. When a sending node sets this field to zero, the destination node is assumed to be on the same local segment as the sending node.
- **Destination Node** — This 6-byte field contains the physical address of the destination node.
- **Destination Socket** — This 2-byte field contains the socket address of the packet's destination process.
- **Source Network** — This 4-byte field provides the source node's network number. If a sending node sets this field to zero, it means the source's local network is unknown.
- **Source Node** — This 6-byte field contains the physical address of the source node. Broadcast addresses are not allowed.
- **Source Socket** — This 2-byte field contains the socket address of the process that transmitted the packet.
- **Upper-Layer Data** — The data field follows the destination and source fields. It contains information for the upper-layer processes.

IPX Packet Delivery

On a NetWare network, the successful delivery of a packet depends on the proper addressing of the packet and the internetwork configuration. Packet addressing is handled in its Media Access Control (MAC) protocol header and IPX header address fields.

To send a packet to another node, the sending node must know the complete internetwork address (network, node, and socket) of the node it wishes to send to. Once the sending node has the destination node's address, it can proceed with addressing the packet. However, the way the MAC header of that packet is addressed depends on whether the sending and destination nodes are separated by a router. See Figure 5-3.

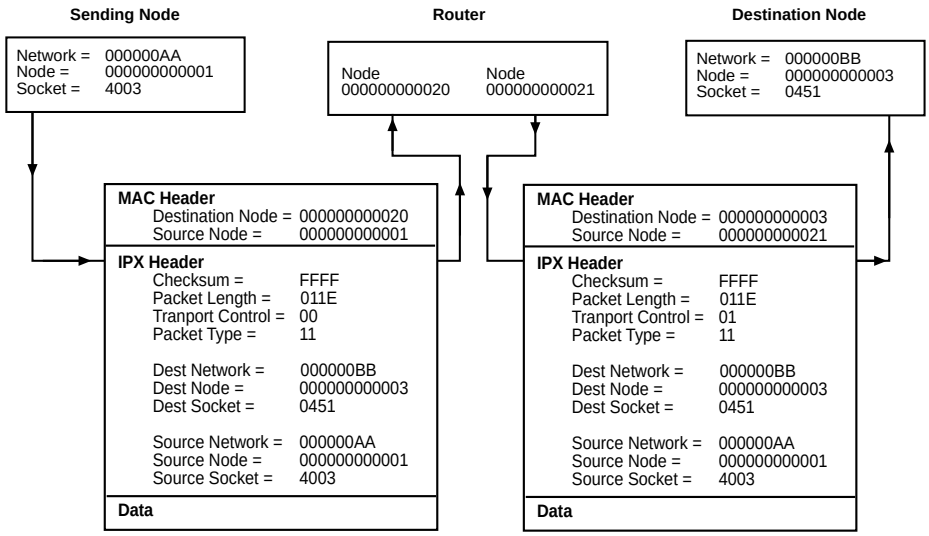


Figure 5-3 IPX Packet Routing

Sending Node's Responsibility

When a node wants to send information to another node with the same network number, the sending node can simply address and send packets directly to the destination node. However, if the two nodes have different network numbers, the sending node must find a router on its own segment that can forward packets to the destination node's network segment.

To find this router, the sending node broadcasts a RIP packet requesting the best route to the destination node's network number. The router residing on the sending nodes segment with the shortest path to the destination

segment responds to the RIP request. The router's response includes its network and node address in the IPX header. If the sending node is a router rather than a workstation, the router can get this information from its internal routing tables and need not send a RIP request.

Once the sending node knows the router's node address, it can send packets to the destination node.

Router's Responsibility

When a router receives an IPX packet, it handles the packet in one of the following methods:

- If the packet is destined for a network number that the router is directly connected to, the router performs the following:
  - **a** It places the destination node address from the IPX header in the destination address field of the MAC header.
  - **b** It places its own node address in the source address field of the MAC header.
  - **c** It increments the Transport Control field of the IPX header and transmits the packet on the destination node segment.
- If the router is not directly connected to the segment that the final destination node resides on, it sends the packet to the next router in the path to the destination node as follows:
  - **a** The router places the node address of the next router in the destination address field of the MAC header. This information is obtained from the Routing Information Table.
  - **b** It places its own node address in the source address field of the MAC header.
  - **c** It increments the Transport Control field in the IPX header and sends the packet to the next router.

The Elements of IPX Routing

IPX routers use the following elements to transmit packets over an intranetwork:

- Router interfaces
- Routing tables
- SAP

Router Interfaces

A router interface is the connection between the router and the network number (address). In traditional routing models, the interface would be the same as the port, since only one interface could exist per port.

In the LANplex system's IPX routing, more than one port can be connected to the network number. Therefore, the router interface is the relationship between the ports and the network number (address) in your IPX network.

Each router interface has a network address. This address defines the network number that the router interface is attached to. A router interface's IPX address serves two functions:

- It is used when sending IPX packets to or from the router itself.
- It defines the network number of the segment connected to that interface.

Routing Tables

A routing table holds information about all the network segments. It allows a router to send a packet toward its ultimate destination using the best possible route. The routing information table contains an entry for every network number that the router currently knows exists. A router uses the routing information table when the destination network number of the packet it is sending is not on a network to which it is directly connected. The routing information table provides the immediate address of a forwarding router that *can* forward the packet toward its destination.

The routing table consists of the following elements:

- **Interface** — Identifies the router's interface that will be used to reach the specific network segment.
- **Address** — Identifies the addresses for segments that the router currently knows exists.

- **Hops to Network** — Provides the number of routers that must be crossed to reach the network segment.
- **Ticks to Network** — Provides an estimate of the time necessary to reach the destination segment.
- **Node** — The node address of the router that can forward packets to each segment. When set to all zeroes, the route is directly connected.
- **Aging Timer** — The time since the network's last update.

Figure 5-4 shows a typical example of a routing information table.

Routing Table

Interface	Address	Hops	Ticks	Node	Age
1	1	1	1	00-00-00-00-00-00	0
2	45469f30	1	1	00-00-00-00-00-00	0
2	45469f33	2	3	08-00-17-04-33-45	40

Figure 5-4 Routing Table Example

Generating Routes

The routing information table is generated and updated as follows:

- **Statically** — You manually enter routes, which do not change until you change them (they do not time out).
- **Dynamically** — The router uses RIP to exchange information. Routes are recalculated at regular intervals.

Static Routes. A static route is one you manually configure in the routing table. Static routes are useful in environments where no routing protocol is used, or where you want to override some of the routes generated with a routing protocol. Because static routes do not automatically change in response to network topology changes, you should only manually configure a small number of reasonably stable routes.

Dynamic Routes Using RIP. Automated methods of learning routes help you keep up with a changing network environment, allowing routes to be reconfigured quickly and reliably. Interior Gateway Protocols (IGP), protocols that operate within intranetworks, provide this automated method. The

LANplex system uses RIP (one of the most widely used IGPs), to dynamically build its routing tables.

RIP operates in terms of active and passive devices. The active devices, usually routers, broadcast their RIP messages to all devices in a network; they update their own routing tables when they receive a RIP message. The passive devices, usually hosts, listen for RIP messages and update their routing tables; they do not send RIP messages.

An active router sends a RIP message every 60 seconds. This message contains both the network number and the number of hops for each destination. In RIP, each router that a packet must travel through to reach a destination equals one hop.

Selecting the Best Route

On large networks, there may be multiple routes to a single network. The criteria that should be used by the routers in selecting the "best route" to a network when choosing between alternate routes are listed below:

- Select the route that requires the lowest number of ticks
- If multiple routes exist with the number of ticks equal, select the route that also has the lowest number of hops
- If multiple routes exist with both ticks and hops equal, the router is free to choose any of the routes as the "best" route

Service Advertising Protocol (SAP)

The Service Advertising Protocol allows servers (for example, file servers, print servers, and gateway servers) to advertise their addresses and services. Through the use of SAP, adding and removing services on an internetwork becomes dynamic. As servers are booted up, they advertise their services using SAP. When they are brought down, they use SAP to indicate that their services are no longer available.

Internetwork Service Information

Using SAP, routers create and maintain a database of internetwork service information. This allows clients on the network to determine what services are available on the network and to obtain the internetwork address of the nodes (servers) where they can access those services.

Workstations cannot initiate a session with a file server without first knowing the server's address.

SAP Packet Structure

SAP uses IPX and the medium-access protocols for its transport. The packet structure allows for the following functions:

- A workstation request for the name and address of the nearest server of a certain type
- A router request for the names and addresses of either all the servers or all the servers of a certain type on the internetwork
- A response to either a workstation or router request
- Periodic broadcasts by servers and routers
- Changed server information broadcasts

Figure 5-5 provides an overview of the SAP packet structure. Notice that the packet structure is encapsulated within the data area of IPX.

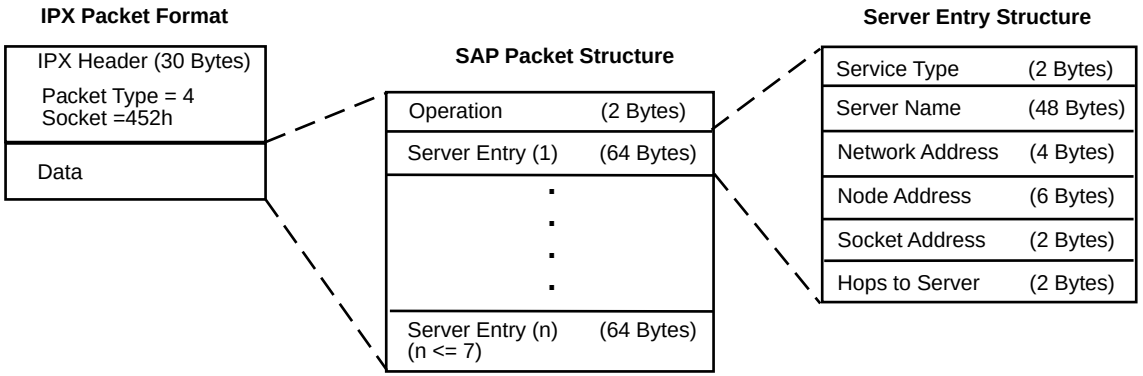


Figure 5-5 SAP Packet Structure

A SAP packet consists of the following fields:

- **Operation** — This field indicates the type of operation the SAP packet performs and can be set to one of the following values:

  1=Request

  2=Response

  3=Get Nearest Server Request

4=Get Nearest Server Response

- **Server Entry** — Each server entry includes information regarding a particular server and consists of the following fields:
  - **Service Type** — This field identifies the type of service the server provides.



Although IPX routers use SAP, routers typically do not act as servers and require no Server Type assignment.

  - **Server Name** — This field contains the 48 byte character string name that is assigned to a server. The Server Name, in combination with the Service Type, uniquely identifies a server on an internetwork.
  - **Network Address** — This field contains the server's network address.
  - **Node Address** — This field contains the server's node address.
  - **Socket Address** — This field contains the socket number the server uses to receive service requests.
  - **Hops to Server** — This field indicates the number of intermediate networks that must be passed through to reach the server associated with this field entry. Each time the packet passes through an intermediate network, the field is incremented by one.

By using SAP, servers can advertise their services and addresses. The information that these servers broadcast is not directly used by clients but is collected by a SAP agent within each router on the server's segment. The SAP agents store this information in a server information table. If the agents reside within a server, the information is also stored in their server's bindery. The clients can then contact the nearest router or file server SAP agent for server information.

The SAP broadcasts that servers and routers send are local broadcasts and, therefore, only received by SAP agents on their connected segments. However, SAP agents periodically broadcast their server information so that all SAP agents on the internetwork have information about all servers that are active on the internetwork.

Server Information Table

A server information table holds information about all the servers on the internetwork. It is this table that SAP agents use to store information received in SAP broadcasts. Figure 5-6 shows an example of a typical server information table.

Server Table

Interface	Name	Type	Network	Node	Socket	Hops	Age
1	LPX1102	4	45469f33	00-00-00-00-00-01	451	2	102
1	LPX1103	4	45469f44	00-00-00-00-00-01	451	5	65
2	LPX2001	4	45470001	00-00-00-00-00-01	451	4	33

Figure 5-6 Server Information Table

The server information table provides the following information:

- **Interface** — indicates which interface the information was received from
- **Server Name** — the name of the server
- **Server Type** — indicates the type of service provided
- **Network Address** — the address of the network on which the server resides
- **Node Address** — the node of the server
- **Socket Address** — the socket number on which the server will receive service requests
- **Hops to Server** — the number of intermediate networks that must be passed through to reach the server associated with this entry
- **Age of Server** — the time since the last update for that server

The server information table is either statically or dynamically generated and updated.

Static Servers. A static server is one you manually configure in the server information table. Static servers are useful in environments where no routing protocol is used, or where you want to override some of the servers generated with a routing/servers protocol. Because static servers do not

automatically change in response to network topology changes, you should only manually configure a small number of servers.

Dynamic Routes Using SAP. The automated method of adding and removing services help you keep up with a changing network environment, allowing servers to advertise their services and addresses quickly and reliably. SAP provides this automated method.

As servers are booted up, they advertise their services using SAP. When servers are brought down, they use SAP to indicate that their services are no longer available.

The information that these servers broadcast is not directly used by clients but instead is collected by a SAP agent within each router on the server's segment. The SAP agents store this information in the server information table. Clients can then contact the nearest router or file server SAP agent for server information.

Server Information Maintenance

When a router's SAP agent receives a SAP broadcast response indicating a change in the internetwork server configuration (for example, a server has gone down, been brought up, or is accessible through a better route), the agent must update its server information table and inform other SAP agents of these changes.

To relay this information to the rest of the internetwork, the SAP agent immediately sends a broadcast to all of its directly connected segments except the segment from which the information was received. This broadcast packet contains information regarding the server change that occurred. This information is also reflected in all future periodic broadcasts.

SAP Aging. Router SAP agents implement an aging mechanism to handle those conditions (for example, hardware failure, power glitch, power outage) that cause a SAP agent to go down suddenly without sending a DOWN broadcast. SAP agents maintain a timer for each entry in their server information table that keeps track of how much time has elapsed since information was received concerning a particular table entry. Since this information is either new or changed, the SAP agents that receive this information immediately pass it on and the change is quickly permeated throughout the internetwork.

SAP Request Handling. When a SAP agent receives a general request, a SAP response packet containing information about all servers of any type known to the SAP agent is sent to the sending source. This response includes the same information sent out in a periodic broadcast. When the request is specific, the SAP agent sends a SAP response directly to the requesting node. This response contains information regarding all servers of the type asked for by the requesting source (as far as the router knows this information).

6 Routing with AppleTalk

This chapter provides an overview of AppleTalk routing, and specifically defines these topics:

- Appletalk Network Elements
- AppleTalk Protocols
- About AARP

About AppleTalk

AppleTalk is a protocol suite defined by Apple Computer, Inc., for connecting computers and peripherals, and other devices on a network. AppleTalk protocols support most of the functions offered by the Open Standards Interconnection (OSI) reference model.

The AppleTalk protocols work together to provide file and printer sharing, and different applications such as electronic mail and database access. All Macintosh computers have Appletalk connectivity options built into them, making it the de facto standard for Apple computer networks.

AppleTalk Network Elements

An AppleTalk network consists of different nodes in groups of networks in an AppleTalk internet. These nodes can include workstations, routers, and printers, or services for other computers, or clients.

This section describes the different elements of an AppleTalk internet, which are as follows:

- AppleTalk Networks
- AppleTalk Nodes
- AppleTalk Zones
- Seed Routers

AppleTalk Networks

A network in an AppleTalk internet is a cable segment attached to a router. Each network is identified by a network number or range of network numbers. The network administrator assigns these numbers from a range of valid network numbers.

Two AppleTalk network numbering systems are currently in use: nonextended (Phase 1) and extended (Phase 2). 3Com routers support extended network numbers. While the LANplex system will not translate Phase 1 packets to Phase 2 packets, it will route packets to a Phase 1 network. When a LANplex system router does this, it anticipates that a gateway exists between the two networks to translate the packets.

An extended network can span a range of logical networks. Network numbers in an extended network consist of a range of numbers, such as 15-20. This numbering scheme allows for as many as 16,580, 608 nodes, although the actual cables will not support this many nodes.

AppleTalk Nodes

A node in a AppleTalk network is any addressable device, including workstations, printers, and routers. Nodes are physically attached to a network. All AppleTalk nodes are identified by a unique AppleTalk address that each node selects at initialization time. The address consists of the node's network number and a unique node number.

Named Entities

When a device on the network provides a service for other users, the network administrator can give the device a name. The name appears on the Chooser of the Macintosh with an associated icon. For example, the Chooser of the Macintosh can include a printer icon. When you select the printer icon, several printer names can appear in a list, such as Laser1, Laser2, Laser3, etc. The Name Binding Protocol (NBP), described later in this chapter, translates device names into AppleTalk addresses.

AppleTalk Zones

An AppleTalk zone is a logical collection of nodes on an AppleTalk internet. A zone can include all nodes in a single network or a collection of nodes in different networks. You assign a unique name to each zone to identify it in the internet. Figure 6-1 illustrates the relationship between physical AppleTalk networks and logical AppleTalk zones.

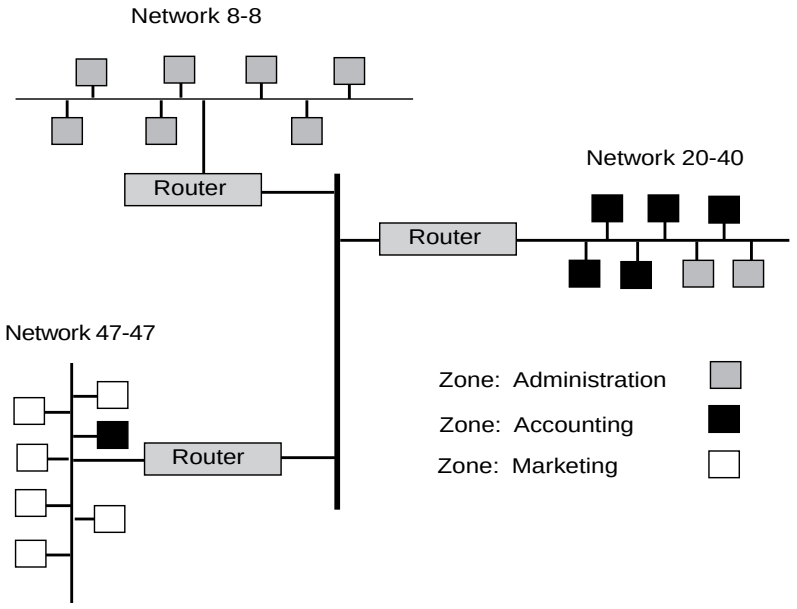


Figure 6-1 AppleTalk Networks and Zones

Figure 6-1 shows an AppleTalk internet with three networks: 47-47, 20-40, and 8-8. Three AppleTalk zones span the networks in this internet: administration, accounting, and marketing. Network 20-40 includes two nodes in the administration zone and five nodes in the accounting zone. Network 47-47 includes a node from the accounting zone as well as the marketing nodes. Network 8-8 consists of nodes in the administration zone only.

Creating zones within a network reduces the amount of searching a router has to do to find a resource on the network. For example, you may want to gain access to a printer on the network. Instead of searching the whole network for that printer, the router searches for it within a particular zone.

You will gain access to the printer more quickly within the zone because the zone includes fewer devices than the entire internet.

Seed Routers

A seed router initializes the internet with AppleTalk configuration information including network numbers and zone names. The seed router broadcasts this information so that nonseed routers can learn it. You can designate a seed router through the Administration Console.

Nonseed routers listen for a seed router and then take the configuration information from the first seed router they detect. After a nonseed router obtains the configuration information, it can participate in the network as if it were a seed router as well.

AppleTalk Protocols

AppleTalk protocols work together to ensure the seamless flow of information throughout the AppleTalk internet. Figure 6-2 shows a simplified view of AppleTalk protocols and their relationship to the OSI reference model. Together, these protocols provide the following services:

- Physical Connectivity
- End-to-End Services
- Reliable Data Delivery

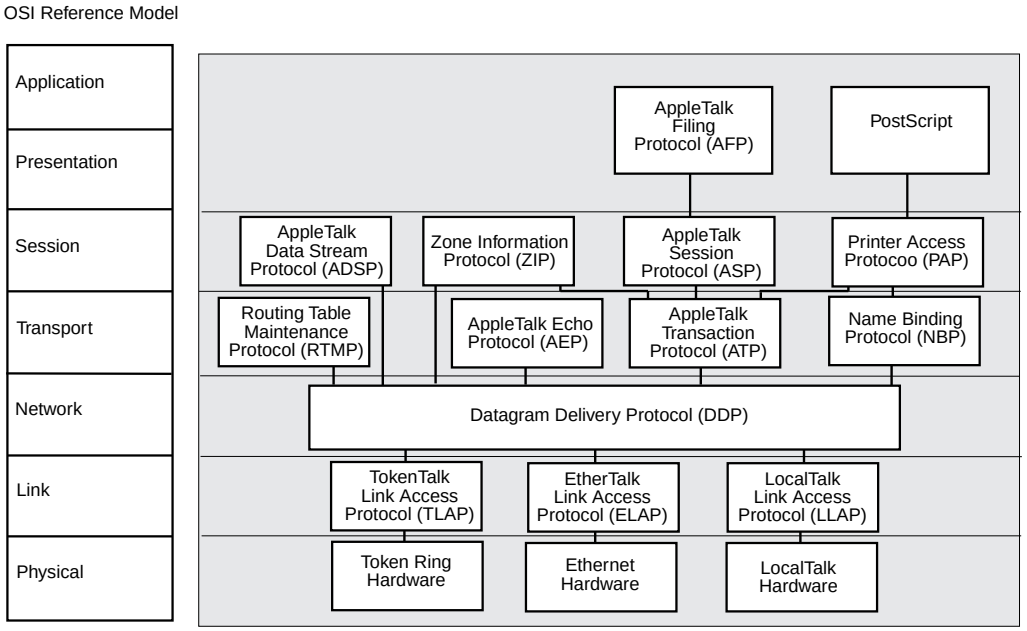


Figure 6-2 AppleTalk Protocols and the OSI Reference Model

The AppleTalk six-layer protocol suite is not fully compliant with the OSI seven-layer reference model. However, AppleTalk provides many of the functions and services provided by OSI. Note that AppleTalk has no specific protocols for the application layer, since the lower levels provide printer and file service.

Physical Connectivity

The physical layer of the protocol stack defines the network hardware. You can use standard network hardware, such as that defined for Ethernet and token ring networks with AppleTalk. Apple has also defined its own network hardware, called LocalTalk, which uses a synchronous RS-422A bus for communications.

The data link layer provides the interface between the network hardware and the upper layers of the protocol stack. The AppleTalk data link layer includes three link access protocols, or LAPs: TokenTalk LAP (TLAP), Ethernet LAP (ELAP), and LocalTalk Link Access Protocol (LLAP).

The AppleTalk Address Resolution Protocol (AARP), which translates hardware addresses to AppleTalk addresses, also exists at the datalink layer because it is closely related to the Ethernet and token ring LAPs. This protocol is usually included in the definition of each LAP, so it does not appear in the reference model. Refer to the section "About AARP" for more information about this protocol.

The Datagram Delivery Protocol (DDP)

The network layer accepts data from the layers above it and divides the data into packets that can be sent over the network through the layers below it. One protocol is present at the AppleTalk network layer: the Datagram Delivery Protocol (DDP).

The DDP transfers data in packets called datagrams. Datagram delivery is the basis for building other AppleTalk services, such as electronic mail. The DDP allows AppleTalk to run as a process-to-process, best-effort delivery system where the processes running in the nodes of interconnected networks can exchange packets with each other.

End-to-End Services

The transport layer and the session layer provide end-to-end services in the AppleTalk network. These services ensure that routers transmit data accurately between one another. Each layer includes four protocols that work together to support these services. This section describes all these protocols, and provides more detail for those that you can view using Administration Console.

Transport Layer Protocols

The four transport layer protocols are these:

- The Routing Table Maintenance Protocol (RTMP)
- The AppleTalk Echo Protocol (AEP)
- The Apple Talk Transaction Protocol (ATP)
- The Name Binding Protocol (NBP)

The Routing Table Maintenance Protocol (RTMP). The RTMP maintains information about AppleTalk addresses and connections between different networks. The RTMP specifies that each router learns about new routes from the other routers and deletes routes after a certain period if the local router no longer broadcasts the route to the network.

Each router builds a routing table that is the basis of dynamic routing operations in an AppleTalk internet. Every ten seconds, each router sends an RTMP data packet to the network. Routers use the information that they receive in the RTMP broadcasts to build their routing tables. Each entry in the routing table contains these items:

- The network range
- The distance in hops to the destination network
- The interface number of the destination network
- The state of each port (good, suspect, bad, really bad)

The router uses these items to determine the best path to forward a data packet to its destination on the network. The routing table contains an entry for each network that a datagram can reach within 15 hops of the router. The table is aged at set intervals as follows:

1. After a period of time, the RTMP changes the status of an entry from good to suspect.
2. After an additional period of time, the RTMP changes the status of an entry from suspect to bad.
3. After an additional period of time, the RTMP changes the status of an entry from bad to really bad.
4. Finally, the router will remove the entry of a nonresponding router with a really bad status from the table.

The data in the routing table is cross-referenced to the Zone Information Table (ZIT). This table maps networks into zones. The section on the session layer protocols includes information about the Zone Information Table.

Figure 6-3 represents a simple AppleTalk network and Table 6-1 shows the corresponding routing table.

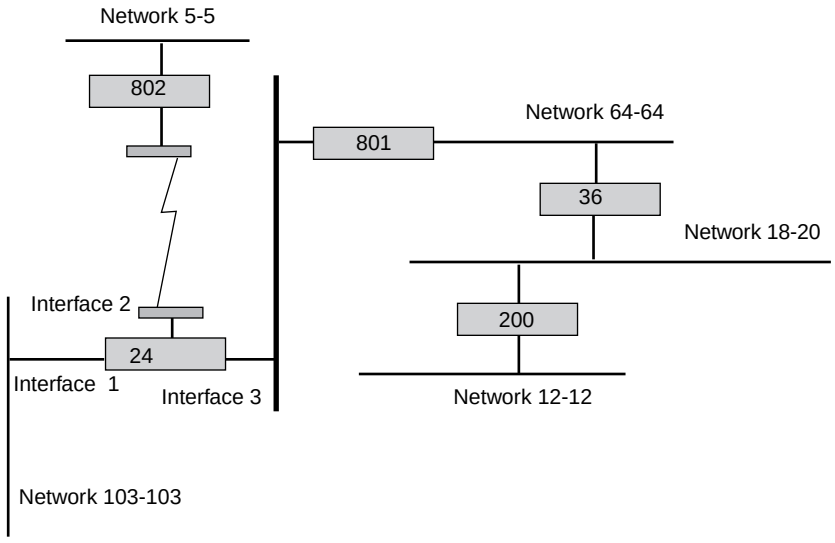


Figure 6-3 A Simple AppleTalk Network

Table 6-1 The Routing Table for Router 24

Network Range	Distance	Interface	State
5-5	1	2	Good
12-12	3	3	Good
18-20	2	3	Good
103-103	0	1	Good
64-64	1	3	Good

You can view the AppleTalk routing tables in your network through the Administration Console.

The AppleTalk Echo Protocol (AEP). AppleTalk nodes use the AEP to send datagrams to other nodes in the network. The AEP causes the destination node to return, or echo, the datagram to the sending node. This protocol can determine whether a node is accessible before any sessions are started, and can enable users to estimate the round-trip delay time between two nodes.

The AppleTalk Transaction Protocol (ATP). ATP, along with the AppleTalk Data Stream Protocol (ADSP), ensures that DDP packets are delivered to a destination without any losses or corruption.

The Name Binding Protocol (NBP). The NBP translates alphanumeric entity names to AppleTalk addresses. The NBP maintains a table that references the addresses of nodes and named entities that reside in that node. Because each node maintains its own list of named entities, the names directory within an AppleTalk network is not centralized. It is a distributed database of all nodes on the internet.

The Session Layer Protocols

The four session layer protocols are these:

- The Zone Information Protocol (ZIT)
- The AppleTalk Data Stream Protocol (ADSP)
- The AppleTalk Session Layer Protocol (ASP)
- The Printer Access Protocol (PAP)

The Zone Information Protocol (ZIP). ZIP works with RTMP to maintain a table that maps network numbers to network zones for the entire AppleTalk internet. Network zones are the logical groupings of AppleTalk networks. The table created by ZIP is called the Zone Information Table (ZIT). The Administration Console allows you to view the zone information table by network number or network zone.

ZIP creates a zone information table in each router. Each entry in the ZIT is a "tuple," or pair, that includes a network number and a network zone name. When an NBP packet arrives at the router, it includes the zone name which the router compares with entries in the zone table. The router then matches the network number from the matching ZIT tuple to that in the RTMP table to find the interface where it can route the packets.

The AppleTalk Data Stream Protocol (ADSP). The ADSP works with the ATP to ensure reliable data transmission. Unlike ATP, however, ADSP provides full-duplex byte-stream delivery. This means that two nodes can communicate simultaneously. ASDP also includes flow control, so that a fast sender does not overwhelm a slow receiver.

The AppleTalk Session Protocol (ASP). The ASP passes commands between a workstation and a server once a connection is made between the two. ASP ensures that the commands are delivered in the same order as they were sent, and returns the results of these commands to the workstation.

The Printer Access Protocol (PAP). The PAP maintains communications between a workstation and a printer, or print service. The PAP functions include setting up and maintaining a connection, transferring the data, and tearing down the connection on completion of the job. Like other protocols at the session layer, PAP relies on NBP to find the addresses of named entities. PAP also depends on ATP for sending data.

The Presentation Layer

The presentation layer maintains information about files, formats, and translations between formats. Two protocols are present at the presentation layer: the AppleTalk Filing Protocol (AFP) and PostScript. AFP provides remote access to files on the network. PostScript is a paged description language used by many printers.

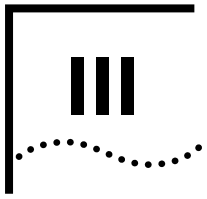
About AARP

The AppleTalk Address Resolution Protocol (AARP) maps the hardware address of an AppleTalk node to an AppleTalk protocol address. It does this for both extended and nonextended networks.

When a node on the network initializes, it randomly selects an AppleTalk address for itself. At the same time, it sends out 10 AARP probe packets. The probe packets determine whether any other nodes on the network are using the address it has chosen. If a node on the network is already using that address, the node randomly selects another address and sends out another probe packet.

The AARP maintains an Address Mapping Table (AMT) with the most recently used hardware addresses and their corresponding AARP addresses.

If an address is not in this table, AARP sends a request to the protocol address and adds the hardware address to the table when the destination node replies. You can view this table, called the AARP Cache, through the Administration Console.

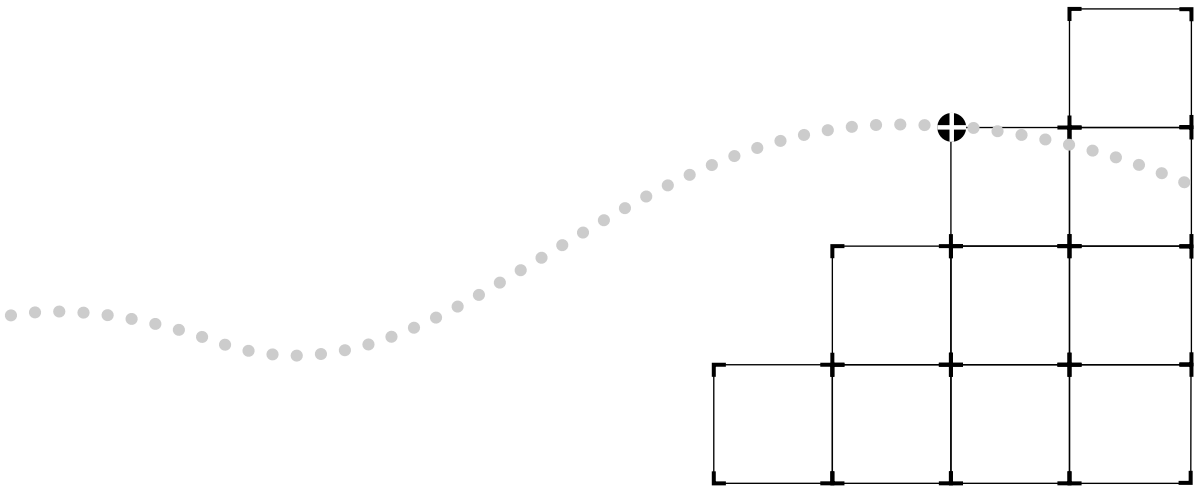


III Administering Routing Protocols

Chapter 7 Administering IP Routing

Chapter 8 Administering IPX Routing

Chapter 9 Administering AppleTalk Routing



7 Administering IP Routing

This chapter describes how to set up your LANplex system to route packets using IP. For more information about how IP works, see Part II of this Guide. You can display and/or configure the following:

- IP interfaces
- Routes
- Address Resolution Protocol (ARP) cache
- UDP Helper
- IP forwarding
- Routing Information Protocol (RIP)
- Ping
- IP statistics

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

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**

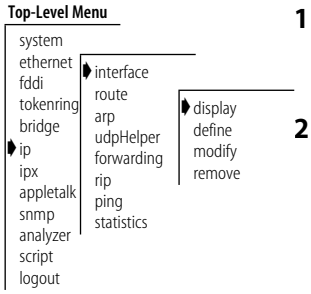
  A single interface may contain several bridge ports. All of the ports corresponding to one interface share the same IP address, subnet mask, broadcast address, and cost. An ESM contains nine ports: one FDDI and nine Ethernet. The port indices are always the following: 1 = FDDI and 2 – 9 = Ethernet. An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

  You do not have to include every switching module port in an interface (that is, some ports may remain unassigned). Packets will be bridged to and from unassigned ports. However, IP packets will not be forwarded to ports that are not assigned to any IP interface.

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 ports with commas (,). Enter a consecutive series of slots using a dash (-).

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

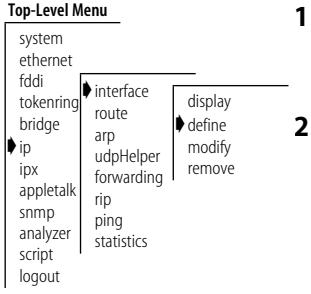
```
Slot 3 - IP forwarding is enabled, RIP is passive.

   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.

To define an IP interface:



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 nonconsecutive ports with commas (,). Enter a consecutive series of ports using a dash (-).

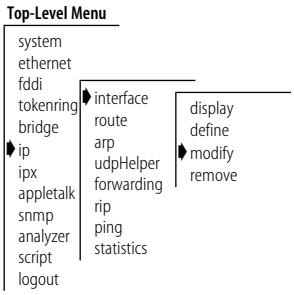
See the example below:

```
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
```

Modifying an Interface

You may want to change the configuration of an interface you have already defined.

To modify an IP interface:



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

```
ip interface modify
```

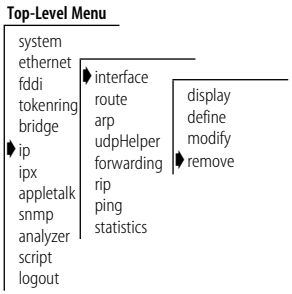
2 Enter the slot of the switching module for which you want to modify an interface.

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

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 route 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.

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 7-1.

Table 7-1 Route Status

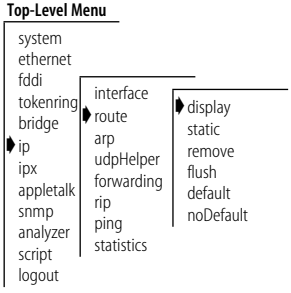
Status	Description
Direct	Route to a directly connected network
Static	Route was statically configured
Learned	Route was learned using indicated protocol
Timing out	Route was learned but is partially timed out
Timed out	Route has timed out and is no longer valid

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 a switching module's routing table to determine which routes are configured and if they are operational.

To display the contents of the routing table:



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.

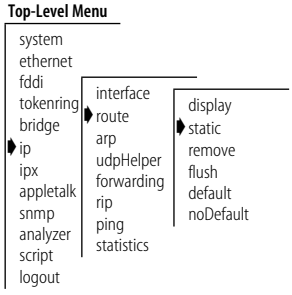
   Destination      Subnet mask      Metric  Gateway          Status
   Default Route     --                   2  158.101.112.250  Learned (RIP)
   10.0.0.0         255.0.0.0             8  158.101.112.254  Learned (RIP)
   129.213.0.0      255.255.0.0           7  158.101.112.254  Learned (RIP)
   137.39.0.0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

Prior to defining static routes on a given switching module, you must define at least one IP interface. 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.

Static routes are not included in periodic RIP updates sent by the switching module.

To define a static route:



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

   `ip route static`

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

2. Enter the slot of the switching module for which you want to define a static route.

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

3. Enter the destination IP address of the route.

4. Enter the subnet mask of the route.

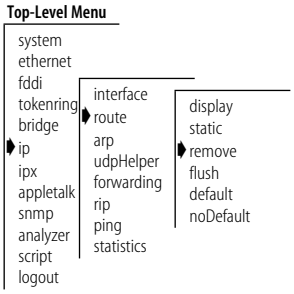
5. Enter the gateway IP address of the route.

   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.

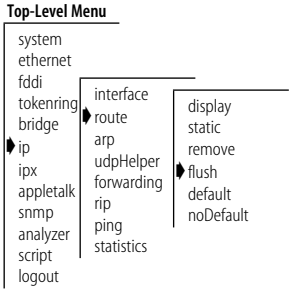
4. Enter the subnet mask of the route.

   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`

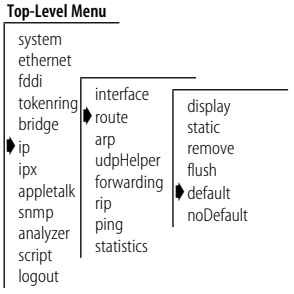
2. Enter the slot of the switching module for which you want to delete the learned routes.

   All learned routes are immediately deleted from the routing table.

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.

To statically configure the default route:



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

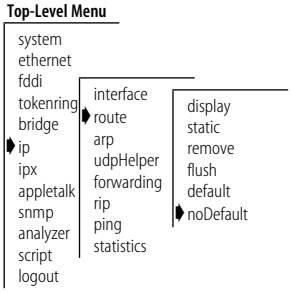
   `ip route default`

2. Enter the slot of the switching module for which you want to set a default route. Enter the gateway IP address of the route.

   The default route is immediately added to the routing table.

Removing the Default Route

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 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.

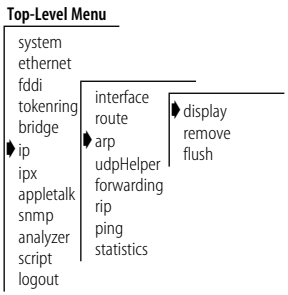
   The default route is immediately removed from the routing table.

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

You can display the contents of the ARP cache for each switching module in your LANplex system.



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 (-).

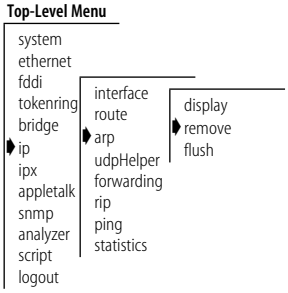
The contents of the ARP cache are displayed as shown in the example below.

```
Slot 3- IP forwarding is enabled,

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
Slot 3- IP forwarding is enabled
```

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`

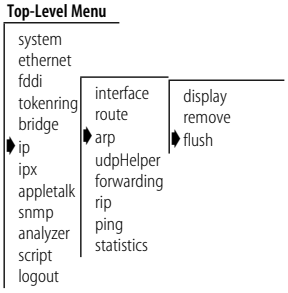
2. Enter the slot of the switching module for which you want to remove an ARP cache entry.

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.

Flushing the ARP Cache

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



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 all entries from the ARP cache.

   The ARP cache entries are immediately removed from the table.

Administering UDP Helper

UDP Helper allows you to send User Datagram Protocol (UDP) packets between routed networks. UDP Helper provides support for UDP services such as BOOTP or DHCP (Dynamic Host Configuration Protocol), that rely on the BOOTP relay agent. For example, by configuring the logical BOOTP port, you can boot hosts through the router. It also provides a relay agent for DHCP broadcasts. UDP packets that rely on the BOOTP relay agent are modified and then forwarded through the router.

The following are the ports for the UDP services mentioned in this section on UDP Helper:

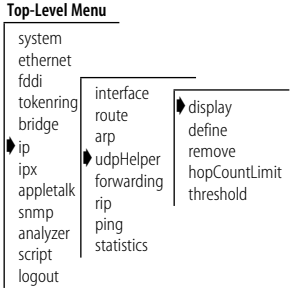
- BOOTP (including DHCP) = 67
- TIME = 37
- DNS = 53

UDP Helper allows you to configure the amount of time a UDP packet is forwarded between subnetworks. UDP packets are discarded based on the hop count and seconds value only for BOOTP and DHCP.

Displaying UDP Helper Information

You can display the Hop and Threshold configuration and list the ports with their IP forwarding addresses that are defined for each switching module in your LANplex system.

To display UDP Helper information:



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

```
ip udpHelper display
```

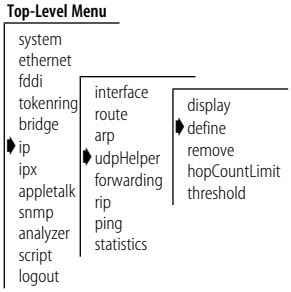
2 Enter the slot(s) of the switching module(s) for which you want to display the UDP Helper information. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

The contents of the UDP Helper are displayed as shown in the example below.

```
Slot 3- IP forwarding is enabled, BOOTP relay hopcount limit is 4,
BOOTP relay threshold is 0.

UDP port            forwarding address
67                  <158.101.1.112
```

Defining a Port and IP Forwarding Address



You can define port numbers and IP forwarding addresses for the UDP Helper. You can have up to 32 combinations of port numbers/IP forwarding addresses per router. You can also have multiple IP address entries for the same ports.

To define port numbers and IP forwarding addresses:

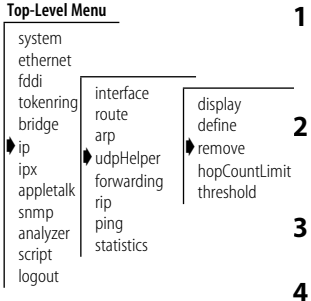
1. From the top level of the Administration Console, enter:

   ```
   ip udpHelper define
   ```

2. Enter the slot of the switching module for which you want to define port numbers and IP forwarding addresses.

3. Enter the port numbers and IP forwarding addresses you want to define.

Removing a Port and IP Forwarding Address



To remove a port number or IP forwarding address defined for UDP Helper:

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

   ```
   ip udpHelper remove
   ```

2. Enter the slot of the switching module for which you want to remove an port number/IP forwarding address.

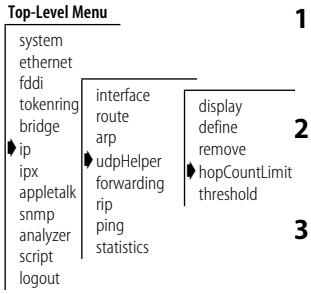
3. Enter the UDP port number that you want to remove.

4. Enter the IP forwarding address that you want to remove.

   The port numbers/IP forwarding addresses are immediately removed.

Setting the Hop Count Limit

You can set the maximum hop count that a packet is forwarded through the router. The range is 0 through 16. The default is 4.



To set the hop count limit:

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

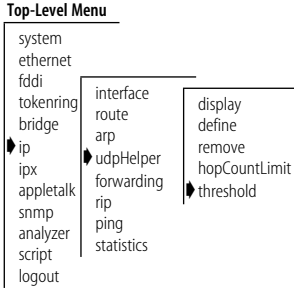
   ```
   ip udpHelper hopCountLimit
   ```

2. Enter the slot of the switching module for which you want to set the hop count limit.

3. Enter the BOOTP relay hop count limit.

Setting the BOOTP Relay Threshold

You can set the maximum time limit that a packet is forwarded through the router. If zero is used as the value, the seconds field is ignored by the router. If a non-zero value is used, the router uses that value along with the hop count value to determine whether to forward the UDP packet.



To set the BOOTP relay threshold:

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

   ```
   ip udpHelper threshold
   ```

2. Enter the slot of the switching module for which you want to set the BOOTP relay threshold.

3. Enter the BOOTP relay threshold.

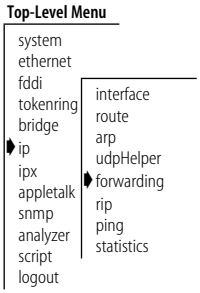
Enabling/ Disabling IP Forwarding

You can control whether a switching module forwards or discards IP packets addressed to other hosts. When you enable IP forwarding, the switching module acts as a normal IP router, forwarding IP packets from one subnet to another when required. When you disable IP forwarding, the switching module discards any IP packets not addressed directly to one of its defined IP interfaces.

IP forwarding default

By default, IP forwarding is enabled on all switching modules.

To enable or disable IP forwarding:



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

   `ip forwarding`

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

3. Enter the IP forwarding state (**`enable`** or **`disable`**).

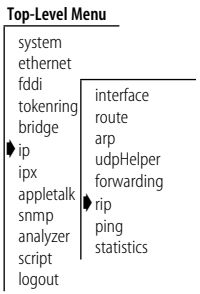
Setting the RIP Mode

You can select a RIP mode that is appropriate for your network. RIP can operate in any of three 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.
- **Active** — The switching module processes all incoming RIP packets, responds to explicit requests for routing information, and broadcasts periodic and 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`**, **`passive`**, or **`active`**). To use the value in brackets, press [Return] at the prompt.

See the example below:

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

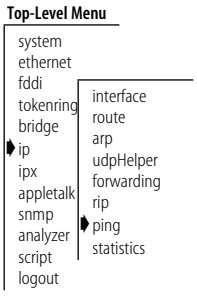
Pinging an IP Station

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
```

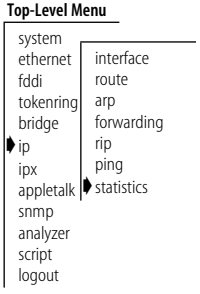
Displaying IP Statistics

The IP statistics you can view are described in Table 7-2.

Table 7-2 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, Slot 3- IP forwarding is
enabled,

      inReceives      forwDatagrams       inDelivers      outRequests
           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
```

8 Administering IPX Routing

This chapter describes how to set up your LANplex system to use the Internet Packet Exchange (IPX) protocol to route packets. For more information about how IPX works, see Part II of this Guide.

You can display and/or configure the following:

- IPX interfaces
- Routes
- Servers
- IPX forwarding
- Routing Information Protocol (RIP)
- Enhanced RIP mode
- Service Advertising Protocol (SAP)
- IPX statistics



Each Ethernet/FDDI Switching Module (EFSM) operates as a separate IPX router. This means that each module has its own interfaces, routing table, and statistics. IPX routing does not operate on the Ethernet Switching Module (ESM) or the Token Ring Switching Module (TRSM).

Administering Interfaces

You define interfaces to establish the relationship between the ports on your EFSMs and the network in your IPX network. You must define one interface for each group of ports that are connected to the same network. This means that every EFSM has one interface defined for each network to which it is directly connected.

An IPX interface has the following information associated with it:

- **IPX Network Address**

  This is a 4-byte address set by the network administrator. Each address within the network should be unique.

- **Cost**

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

- **Encapsulation Format**

  There are four Ethernet encapsulation formats and two FDDI encapsulation formats used in IPX routing. The Ethernet encapsulation formats are Ethernet Type II, Novell 802.3 raw, 802.2 LLC, and 802.3 SNAP. The FDDI encapsulation formats are FDDI 802.2 and FDDI SNAP. The two FDDI encapsulation formats correspond to the Ethernet 802.2 LLC and 802.3 SNAP encapsulation formats. If you select either of these Ethernet encapsulation formats, the corresponding FDDI encapsulation format is automatically selected for shared Ethernet and FDDI ports.

- **Ports**

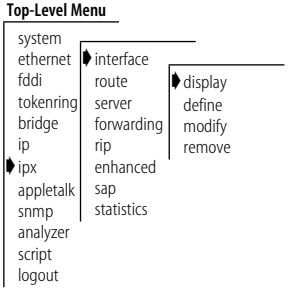
  A single interface may contain several bridge ports. All of the ports corresponding to one interface share the same IPX address, cost, and encapsulation format. An EFSM contains a maximum of eighteen ports: two FDDI and sixteen Ethernet. The port indices for the maximum configuration are the following: 1,2 = FDDI; 3 – 18 = Ethernet.

  You do not have to include every EFSM port in an interface (that is, some ports may remain unassigned). Packets will be bridged to and from unassigned ports. However, IPX packets will not be forwarded to ports that are not assigned to an IPX interface.

Displaying IPX Interfaces

You can display a table that shows all IPX interfaces and their parameter settings configured for each EFSM in the system.

To display IPX interface information:



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

```
ipx interface display
```

2 Enter the slot(s) of the EFSM(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 IPX forwarding and RIP and SAP information for that slot as well as IPX interface information.

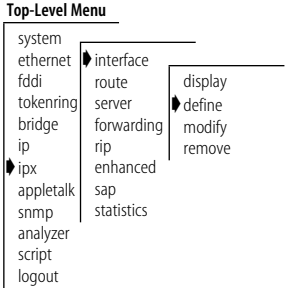
```
ıt 3 - IPX forwarding is enabled, RIP is active, SAP is active

Index   IPX Address     Cost  Format      Ports (1-2=FDDI, 3-18=Ethernet)
    1      45469f30        1  802.2       3-10
    2      5d41a110        1  802.2       11-18
```

Defining an Interface

When you define an interface, you define the interface's IPX address, cost, format, and the EFSM ports associated with the interface.

To define an IPX interface:



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

```
ipx interface define
```

2 Enter the slot of the EFSM 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 IPX network address of the interface.

4 Enter the cost of the interface.

5 Enter the format of the interface.

6 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:

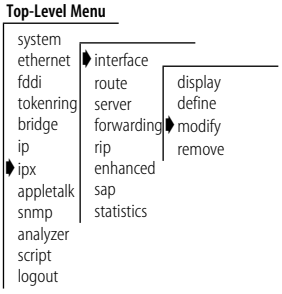
```
Select slot {3-4} [3-4]: 3

Enter IPX Address: 0x45469f30
Enter Cost: 1
Enter Frame Format (Ethernet II: 0, 802.2: 1, Raw 802.3: 2, SNAP: 3): 1
Enter ports(s) (1-2=FDDI, 3-18=Ethernet) (1-18|all): 3-10,13,16
```

Modifying an Interface

You may want to change the configuration of an interface you have already defined.

To modify an IPX interface:



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

```
ipx interface modify
```

2 Enter the slot of the EFSM for which you want to modify an interface.

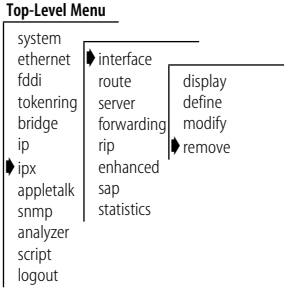
You are prompted for the interface parameters. Press [Return] at the prompts for which you do not want to modify the value.

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 perform routing on the ports associated with the interface.

To remove an IPX interface definition:



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

```
ipx interface remove
```

2 Enter the slot of the EFSM from which you want to remove an interface.

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

Administering Routes

Each EFSM maintains a table of routes to other IPX networks. You can either use the Routing Information Protocol (RIP) to exchange routing information automatically or make static entries in this table using the Administration Console.

Each routing table entry contains the following information:

- **Address**

  The 4-byte IPX network address of a segment currently known to the router.

- **Hops**

  The number of routers that must be crossed to reach the network segment. The maximum number of routers a packet can cross is fifteen. The maximum number of routers an IPX NetBIOS packet can cross is seven.

- **Tics**

  An estimate of how long it will take the packet to reach this segment. A tic is approximately 55 milliseconds.

- **Node**

  The 6-byte address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.

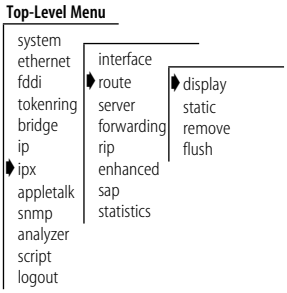
- **Age**

  This is the number of seconds that have elapsed since the last time the route was heard from.

Displaying the Routing Table

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

To display the contents of the routing table:



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

   `ipx route display`

2. Enter the slot(s) of the EFSM(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 are displayed. The configuration of IPX forwarding, RIP, and SAP is indicated in the display.

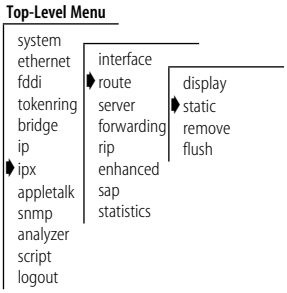
```
Slot 3 - IPX forwarding is enabled, RIP is active, SAP is active

   Interface     Address     Hops  Tics                Node   Age
           2    45469f02        5     6  08-00-02-04-80-b6    44
           2    c2c028ca        4    28  08-00-02-04-80-b6    85
           2    aaaaaaaa        6   671  08-00-02-04-80-b6    85
```

Defining a Static Route

Prior to defining static routes on a given EFSM, you must define at least one IPX interface (see the section "Defining an Interface" on page 8-3). 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. You can set up to 16 static routes.

To define a static route:



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

   `ipx route static`

2. Enter the slot of the EFSM for which you want to define a static route.

3. Enter the 4-byte IPX network address of the route.

4. Enter the cost of the route.

5. Enter the interface number of the route.

6 Enter the node address of the route.

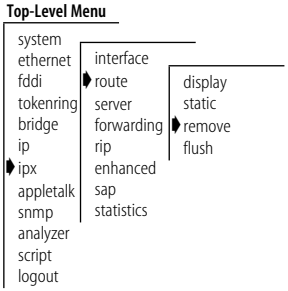
A static route is defined in the following example:

```
Select slot {3-4} [3-4]: 3

Enter IPX address: 0x45469f30
Enter Cost: 1
Enter Interface number: 1
Enter node address: 08-00-3e-22-15-78
```

Removing a Route

To remove a route:



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

```
ipx route remove
```

2 Enter the slot of the EFSM for which you want to remove any static or dynamic route.

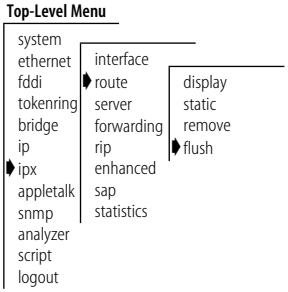
3 Enter the 4-byte IPX network address.

The route is immediately deleted from the routing table.

Flushing Routes

Flushing deletes all dynamically-learned routes from the routing table.

To flush all learned routes:



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

```
ipx route flush
```

2 Enter the slot of the EFSM for which you want to delete the learned routes.

All learned routes are immediately deleted from the routing table.

Administering Servers

Each EFSM maintains a table of servers on other IPX networks. You can either use the Service Advertising Protocol (SAP) to exchange server information automatically or make static entries in this table using the Administration Console.

Each server table contains the following information:

- **Name**

  The user-defined name of the server.

- **Type**

  The type of service provided by the server.

- **Node**

  The 6-byte address of the server that can forward packets to the segment.

- **Socket**

  The 2-byte socket address on which the server will receive service requests.

- **Hops**

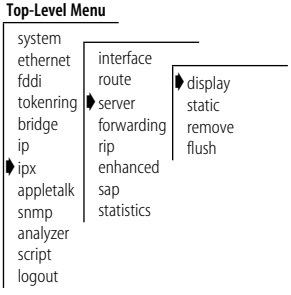
  The number of networks that must be crossed to reach the server. The maximum number is fifteen.

- **Age**

  This is the number of seconds that have elapsed since the last time a server in the table was heard from.

Displaying the Server Table

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



To display the contents of the server table:

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

   ```
   ipx server display
   ```

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

In the following example, servers known to an EFSM in slot 3 are displayed. The configuration of IPX forwarding, RIP, and SAP is indicated in the display.

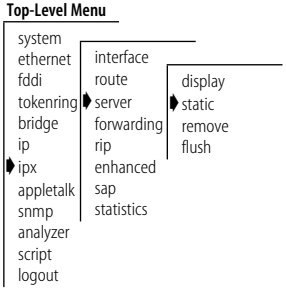
```
Slot 3 - IPX forwarding is enabled, RIP is active, SAP is active

   Interface        Name  Type  Network              Node Socket  Hops   Age
           2       GB201   39b 8c141bfe  08-00-02-04-80-b6   8059     4    73
           2     GB3COM2   39b af0bc60f  00-00-00-00-00-01   85fa     4    85
```

Defining a Static Server

Prior to defining static servers on a given EFSM, you must define at least one IPX interface (see the section "Defining an Interface" on page 8-3). Static servers remain in the table until you remove them, or until you remove the corresponding interface. Static servers take precedence over dynamically-learned servers to the same destination. You can have a maximum of eight static servers.

To define a static server:



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

   `ipx server static`

2. Enter the slot of the EFSM for which you want to define a static server.
3. Enter the interface number of the server.
4. Enter the service type of the server.
5. Enter the service name of the server.
6. Enter the IPX network address of the server.
7. Enter the socket value of the server.
8. Enter the node address of the server.
9. Enter the number of hops to the server.

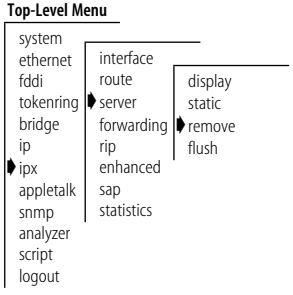
   A static server is defined in the following example:

```
Select slot {3-4} [3-4]: 3

Enter Interface number: 1
Enter service type: 4
Enter service name: gb201
Enter IPX address: 0x8c14a238
Enter socket: 0x8059
Enter node address: 00-00-2e-f3-56-01
Enter hops: 2
```

Removing a Server

To remove a server:



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

   `ipx server remove`

2. Enter the slot of the EFSM for which you want to remove any static or dynamic server.

3. Enter the service type of the server.

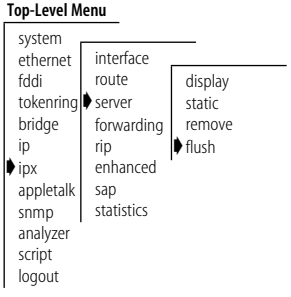
4. Enter the service name of the server.

   The server is immediately deleted from the server table.

Flushing Servers

Flushing deletes all dynamically-learned servers from the server table.

To flush all learned servers:



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

   `ipx server flush`

2. Enter the slot of the EFSM for which you want to delete the learned server.

   All learned servers are immediately deleted from the server table.

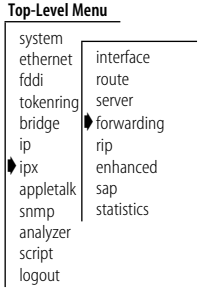
Setting IPX Forwarding

You can control whether an EFSM forwards or discards IPX packets addressed to other routers. When you enable IPX forwarding, the EFSM acts as a normal IPX router, forwarding IPX packets from one network to another when required. When you disable IPX forwarding, the EFSM discards any IPX packets not addressed directly to one of its defined IPX interfaces.

IPX forwarding default

By default, IPX forwarding is disabled on all EFSMs.

To enable or disable IPX forwarding:



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

   `ipx forwarding`

2. Enter the slot(s) of the EFSM(s) for which you want to enable or disable IPX forwarding. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

3. Enter the IPX forwarding state (**`enabled`** or **`disabled`**). To use the value in brackets, press [Return] at the prompt.

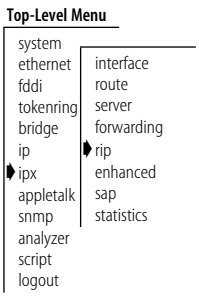
Setting the RIP Mode

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

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

RIP default mode By default, RIP is off.

To set the RIP operating mode:



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

   `ipx rip`

2. Enter the slot(s) of the EFSM(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 (-).

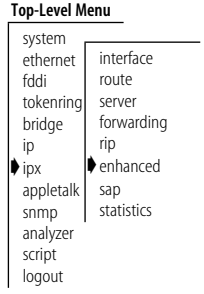
3. Enter the RIP mode (**`off`**, **`passive`**, or **`active`**). To use the value in brackets, press [Return] at the prompt.

Setting the Enhanced RIP Mode

Standard IPX RIP packets can include up to 50 route advertisements, but some routers allow up to 68. Enhanced RIP mode increases the number of entries in a RIP packet that the EFSM will accept, allowing the EFSM to have greater interoperability with routers that do not explicitly follow the IPX router implementation guidelines.

Enhanced RIP default By default, enhanced RIP is disabled on all EFSMs.

To enable or disable enhanced RIP mode:



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

   `ipx enhanced`

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

3. Enter the enhanced RIP state (**`enabled`** or **`disabled`**). To use the value in brackets, press [Return] at the prompt.

Setting the SAP Mode

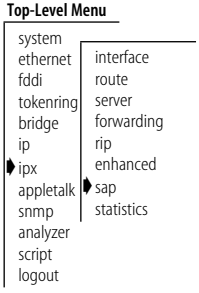
You can select a SAP mode that is appropriate for your network. SAP can operate in any of three modes:

- **Off** — The EFSM ignores all incoming SAP packets and does not generate any SAP packets of its own.
- **Passive** — The EFSM processes all incoming SAP packets, but does not broadcast periodic or triggered SAP updates, or respond to SAP requests.

  Active — The EFSM processes all incoming SAP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered SAP updates.

SAP default mode By default, SAP is off.

To set the SAP operating mode:



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

   `ipx sap`

2. Enter the slot(s) of the EFSM(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 SAP mode (**`off`**, **`passive`**, or **`active`**). To use the value in brackets, press [Return] at the prompt.

Displaying Statistics

The Administration Console allows you to display four types of IPX-related statistics:

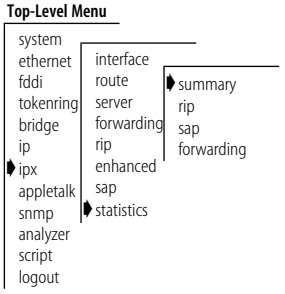
- IPX Summary statistics
- IPX RIP statistics
- IPX SAP statistics
- IPX Forwarding statistics

Displaying IPX Summary Statistics

To display IPX summary statistics:

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

```
ipx statistics summary
```



2 Enter the slot(s) of the EFSM(s) for which you want to view IPX statistics. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

Statistics appear as in the example below:

```
Slot 3 - IPX forwarding is enabled, RIP is active, SAP is active

    Received       Transmitted       Dropped       Msg Pool Empty
     1170878            565099             0                    0
```

The IPX summary statistics you can view are described in Table 8-1.

Table 8-1 IPX Summary Statistics

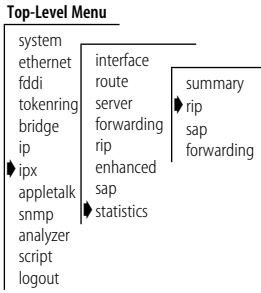
Field	Description
Received	Number of IPX packets received
Transmitted	Number of IPX packets transmitted
Dropped	Number of IPX packets dropped
Msg Pool Empty	Number of IPX RIP or IPX SAP messages delivered to the IPX application that are dropped due to resource limitations

Displaying IPX RIP Statistics

To display IPX RIP statistics:

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

```
ipx statistics rip
```



2 Enter the slot(s) of the EFSM(s) for which you want to view IPX RIP statistics. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

Statistics appear as in the example below:

```
Slot 3 - IPX forwarding is enabled, RIP is active, SAP is active

  RIP Received        RIP Transmitted         RIP dropped
        106195                   7929                   0

 RIP Responses          RIP Requests         RIP Entries
        100552                   5643                   2
```

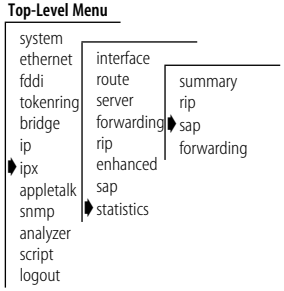
The IPX RIP statistics you can view are described in Table 8-2.

Table 8-2 IPX RIP Statistics

Field	Description
RIP Received	Number of IPX RIP packets received
RIP Transmitted	Number of IPX RIP packets transmitted
RIP Dropped	Number of IPX RIP packets dropped
RIP Responses	Number of IPX RIP Responses that have been processed
RIP Requests	Number of IPX RIP Requests that have been processed
RIP Entries	Number of routes in the routing table

Displaying IPX SAP Statistics

To display IPX SAP statistics



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

```
ipx statistics sap
```

2 Enter the slot(s) of the EFSM(s) for which you want to view IPX SAP 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 - IPX forwarding is enabled, RIP is active, SAP is active

        SAP Received        SAP Transmitted        SAP dropped
             1064015                  22493                  0

      SAP Responses           SAP Requests        SAP Entries
             1063532                     45                  0

  SAP GNS Responses       SAP GNS Requests
                  0                    438
```

The IPX SAP statistics you can view are described in Table 8-3.

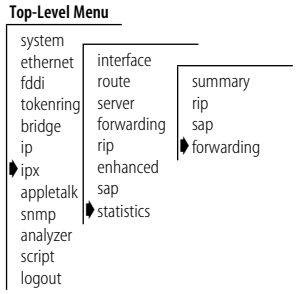
Table 8-3 IPX SAP Statistics

Field	Description
SAP Received	Number of IPX SAP packets received
SAP Transmitted	Number of IPX SAP packets transmitted
SAP Dropped	Number of IPX SAP packets dropped
SAP Responses	Number of IPX SAP Responses that have been processed
SAP Requests	Number of IPX SAP Requests that have been processed
SAP Entries	Number of servers in the server table
SAP GNS Responses	Number of IPX SAP Get Nearest Service Responses that have been received
SAP GNS Requests	Number of IPX SAP Get Nearest Service Requests processed

Displaying IPX Forwarding Statistics

To display IPX Forwarding statistics:

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




```
ipx statistics forwarding
```

2 Enter the slot(s) of the EFSM(s) for which you want to view IPX forwarding 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:

The IPX forwarding statistics you can view are described in Table 8-4.

```
Slot 3 - IPX forwarding is enabled, RIP is active, SAP is active

                  Received                    Transmitted                   Forwarded
                   1335653                         565105                           0

                Hdr Errors                Hop Count Errors                 Addr Errors
                     13758                             0                       13758

                No Routes                     Misc Errors
                         2                             411

                NetBIOS Rx                    NetBIOS Tx              NetBIOS Max Hops
                    150604                        125781                           0

                   Host Rx                       Host Tx
                   1171190                        565105
```

Table 8-4 IPX Forwarding Statistics

Field	Description
Received	Number of IPX Forwarding packets received
Transmitted	Number of IPX Forwarding packets transmitted
Forwarded	Number of IPX packets forwarded by the IPX router
Hdr Errors	Number of IPX packets dropped due to IPX Network layer header errors
Hop Count Errors	Number of IPX packets dropped due to exceeded maximum transport control
Addr Errors	Number of IPX packet dropped due to IPX Address errors in network layer header
No Routes	Number of IPX packets dropped because the IPX route is unknown
Misc Errors	Number of multicasts attempted to be forwarded
NetBios Rx	Number of IPX NetBIOS packets received
NetBios Tx	Number of IPX NetBIOS packets transmitted
NetBios Max Hops	Number of IPX NetBIOS packets that exceeded the Transport control maximum
Host Rx	Number of IPX packets delivered to the IPX host's RIP and SAP applications
Host Tx	Number of IPX packets transmitted from IPX host's RIP and SAP applications

9 Administering AppleTalk Routing

This chapter describes how to set up your LANplex system to use the AppleTalk protocol to route packets. For more information on how AppleTalk routing works, see Part II of this Guide.

You can display and/or configure the following:

- AppleTalk interfaces
- Routes
- AARP cache
- Zones
- AppleTalk Forwarding
- Checksum generation/verification
- AppleTalk statistics

Each Ethernet/FDDI Switching Module (EFSM) operates as a separate AppleTalk router. This means that each module has its own interfaces, routing table, and statistics. AppleTalk routing does not operate on the Ethernet Switching Module (ESM) or the Token Ring Switching Module (TRSM).

Administering Interfaces

You define interfaces to establish the relationship between the ports on your EFSMs and the subnets in your network. You must define one interface for each group of ports that are connected to the same subnet. This means that every EFSM has one interface defined for each network to which it is directly connected.

The maximum number of interfaces you can configure per router is 18.

An AppleTalk interface has the following information associated with it:

- **Seed Interface**

  You can configure the interface to be a seed interface or nonseed interface. Seed interfaces initialize the network with the configuration information the administrator enters (network range, address, zone name, and ports). Nonseed interfaces wait and listen for a seed interface and then take this configuration initialization information from the first seed interface they hear. After the nonseed interface obtains a network configuration, it begins to participate in the routing of the network.

- **Network Range**

  A range of numbers used to designate a network segment's identity. This allows the physical segment between two LANplex systems to be a range of multiple networks.

- **Address**

  The AARP address based on the network range and the network node (1-253).

- **Zone**

  The default zone name, as well as up to 15 additional defined zones.

- **Ports**

  A single interface may contain several bridge ports. All of the ports corresponding to one interface share the same AppleTalk address, cost, and format. 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; 3 – 18 = Ethernet.

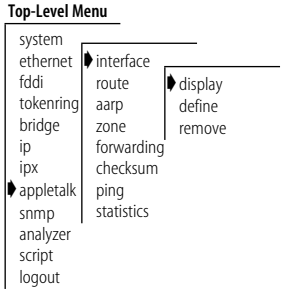
  You do not have to include every EFSM port in an interface (that is, some ports may remain unassigned). Packets will be bridged to and from

unassigned ports. However, AppleTalk packets will not be forwarded to ports that are not assigned to an AppleTalk interface.

Displaying AppleTalk Interfaces

You can display a table that shows all AppleTalk interfaces and their parameter settings configured for each EFSM in the system.

To display the AppleTalk interfaces defined on the router:



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

```
appletalk interface display
```

2 Enter the slot(s) of the EFSM(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 (-).

An example of interfaces defined is shown below:

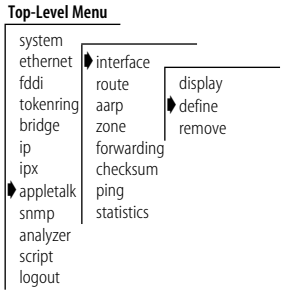
```
lot 2 - DDP forwarding is enabled.

  Index  Network Range    Address      State  Ports (1-2=FDDI, 3-18=Ethernet)
      1    20112-20112   20112.27    enabled      3
      2    20124-20124    20124.1    enabled      4-10
      3    20125-20125    20125.1    enabled      11-18
```

Defining an Interface

When you define an interface, you define the interface's network range, zone name, and the EFSM ports associated with the interface.

To define an AppleTalk interface:



1 At the Administration Console's top-level menu, enter:

```
appletalk interface define
```

2 Enter the slot of the EFSM 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.

The following message appears:

```
Configure seed interface? (n,y) [y]:
```

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

4 Enter the start of the network range associated with the interface.

5 Enter the end of the network range associated with the interface.

6 Enter the default zone name.

The default zone name is used by clients that have not been configured to use a particular zone.

7 Enter the zone name.

You can enter up to 16 zone names per interface.

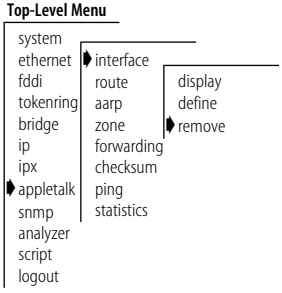
8 Type **q** after entering all the zone names.

9 Enter the ports you want to include on the interface. Separate non-consecutive ports with commas (,). Enter a consecutive series of ports using a dash (-).

Removing an Interface

You may want to remove an interface if you no longer perform routing on the ports associated with the interface.

To remove an AppleTalk interface:



1 At the Administration Console's top-level menu, enter:

```
appletalk interface remove
```

2 Enter the slot of the EFSM from which you want to remove an interface.

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

The interface is no longer defined on the router.

Administering Routes

Each EFSM maintains a table of routes to other AppleTalk networks. The routing table is generated automatically by the Routing Table Maintenance Protocol (RTMP). RTMP defines the rules for exchanging information between routers so that the routers can maintain their routing tables, as well as the rules for the information contained within each routing table.

Each routing table entry contains the following information:

- **Network Range**

  A range of numbers used to designate a network segment's identity

- **Distance**

  The distance in hops to the destination network

- **Interface**

  The defined interface number

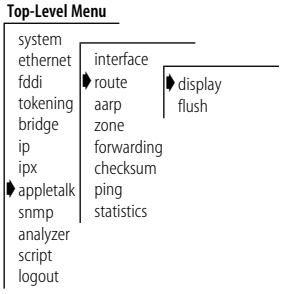
- **State**

  The status (good, suspect, bad, or really bad) of each route

Displaying the Routing Table

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

To display the contents of the routing table:



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

```
appletalk route display
```

2 Enter the slot(s) of the EFSM(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 (-).

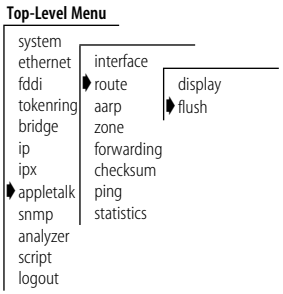
An example of a routing table display is shown below:

```
Slot 2 - DDP forwarding is enabled.

   Network Range   Distance   Interface        State
             1-1         10           1         good
               3          4           1         good
           10-14          6           1         good
           15-19          7           1         good
              61          6           1         good
         100-100         10           1         good
         201-300          7           1         good
       2010-2015          2           1         good
     10009-10009          5           1         good
     10010-10010          7           1         good
     10060-10060          8           1         good
     10110-10113          5           1         good
     10116-10117          5           1         good
     10118-10118          6           1         good
     10119-10119          4           1         good
     10120-10120          7           1         good
     10122-10122          9           1         good
     10310-10329         10           1         good
     10410-10410          8           1         good
     11010-11019          9           1         good
```

Flushing all Routes

Flushing deletes all dynamically learned routes from the routing table.



To flush all learned routes:

1. At the Administration Console's top-level menu, enter:

   ```
   appletalk route flush
   ```

2. Enter the slot(s) of the EFSM(S) for which you want to flush all learned routes.

Administering the AARP Cache

AARP allows hardware addresses to be mapped to an AppleTalk protocol address. AppleTalk uses dynamically assigned 24-bit addresses, unlike the statically-assigned 48-bit addresses used by Ethernet and token ring.

To make the address mapping process easier, AARP uses an Address Mapping Table (AMT). The most recently used addresses are maintained in the AMT. If an address is not in the AMT, AARP sends a request to the desired protocol address and the hardware address is added to the table when the destination node replies.

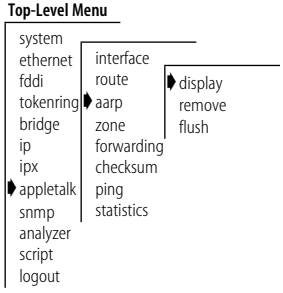
AARP is also responsible for registering a node's dynamically assigned address on the network. This process is described below:

- AARP randomly assigns an address.
- AARP broadcasts AARP probe packets to this address to determine if another node is already using the address.
- If there is no reply, the address becomes that node's address.
- If there is a reply, AARP repeats this process until an available address is discovered.

In the Administration Console, you can:

- Display the cache
- Remove entries
- Flush the cache

Displaying the AARP Cache



You can display the AARP cache for the EFSMs in a system to determine which routes are configured and if they are operational.

To display the contents of the AARP cache:

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

```
appletalk aarp display
```

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

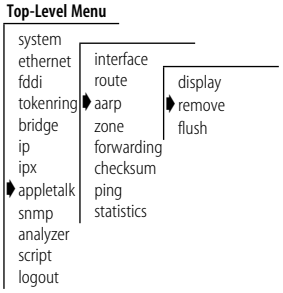
An example of an AARP cache display is shown below:

```
Slot 2 - DDP forwarding is enabled.

   AARP Address         MAC Address   Interface   Age (secs)
      20112.125   00-20-af-0b-e1-7c           1          211
      20112.177   00-00-89-01-91-f0           1           20
      20112.192   00-00-89-01-91-f3           1            6
      20112.150   00-00-89-01-8b-51           1           18
        20112.1   08-00-02-04-80-b6           1           31
      20125.193   08-00-07-d7-69-1f           3          388
       20125.76   08-00-07-66-62-9d           3          862
       20125.67   08-00-07-ee-10-a2           3          851
       20124.41   08-00-07-7c-c3-d8           2          864
      20112.225   00-20-af-0b-d8-f1           1          270
      20112.135   00-20-af-9e-68-62           1          174
      20112.147   00-00-94-41-de-79           1           26
      20112.132   08-00-09-7f-98-c5           1           24
      20112.112   08-00-07-7c-20-61           1          121
      20112.148   08-00-07-ac-56-4b           1         1098
      20112.244   00-20-af-0b-ff-72           1           35
       20112.21   08-00-07-dc-e5-c4           1         8932
      20112.131   08-00-07-54-88-b1           1          397
       20124.35   08-00-07-57-ec-58           2          368
       20112.97   08-00-07-9e-09-86           1         1925
        20112.4   08-00-07-ec-98-3d           1          121
      20112.180   08-00-07-f7-cf-de           1          110
      20112.108   08-00-07-4f-74-7e           1         5833
       20112.56   08-00-07-bc-10-fc           1          120
      20112.110   00-40-10-56-1a-b5           1          110
      20112.155   08-00-07-6c-88-77           1         5536
      20112.243   08-00-07-66-72-c7           1         4940
      20112.253   08-00-20-12-75-bf           1           70
      20125.104   08-00-07-66-2b-c2           3          848
      20112.236   00-80-3e-02-81-66           1         3841
```

Removing an Entry in the Cache

To remove an AARP cache entry:



1. At the Administration Console's top-level menu, enter:

   `appletalk aarp remove`

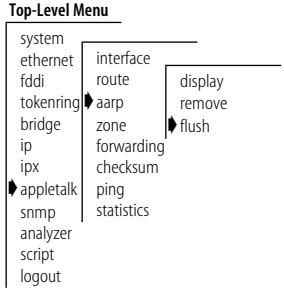
2. Enter the slot(s) of the EFSM(s) for which you want to remove the AARP cache entry.

3. Enter the AARP address at the prompt.

   The entry is removed.

Flushing All Cache Entries

To flush all AARP cache entries:



1. At the Administration Console's top-level menu, enter:

   `appletalk aarp flush`

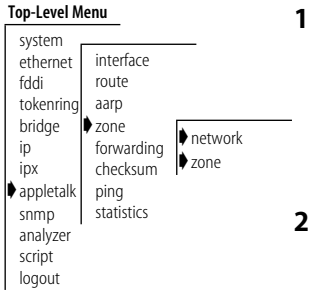
2. Enter the slot(s) of the EFSM(s) for which you want to flush all AARP cache entries.

Displaying the Zone Table

AppleTalk allows for the logical grouping of nodes into zones to make navigation through the network easier. This is done with the Zone Information Protocol (ZIP). ZIP helps routers maintain a mapping of network numbers to zones in the entire network. To do this, ZIP creates and maintains a Zone Information Table (ZIT) in each router. The entries in this table match the network numbers with the zone names.

In the Administration Console, you can display the zone table either by network numbers or by zones.

To display the zone table:



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

`appletalk zone display network`

OR

`appletalk zone display zone`

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

Depending on the command entered, the zone table is displayed by network or zone. An example of each type of display is shown below:

Zone Table by Network Numbers

```
Slot 2 - DDP forwarding is enabled.

   Network 1-1 has 1 known zone
      Munich GmbH

   Network 3 has 1 known zone
      Ethernet A5D85800

   Network 10-14 has 1 known zone
      Freds_Ethernet

   Network 15-19 has 1 known zone
      Freds_Token

   Network 61 has 1 known zone
      DevMacNet

   Network 100-100 has 1 known zone
      France Les Ulis

   Network 201-300 has 1 known zone
      Fred_Wilma

   Network 2010-2015 has 1 known zone
      NY

   Network 10009-10009 has 2 known zones
      Hemel NSOPS
      3Com Arpeggio

   Network 10010-10010 has 1 known zone
      Marlow EUR
```

Zone Table by Zones

```
Slot 2 - DDP forwarding is enabled.

   Zone Holmdel is assigned to 2 networks
      21105-21105
      21010-21010

   Zone NY is assigned to 2 networks
      63535-63535
      2010-2015

   Zone Manchester UK is assigned to 1 network
      10310-10329

   Zone DC8 is assigned to 1 network
      30110-30129

   Zone Chicago is assigned to 1 network
      22030-22030

   Zone Startek-Enet1 is assigned to 1 network
      20033-20033

   Zone Startek-TR1 is assigned to 1 network
      20037-20037

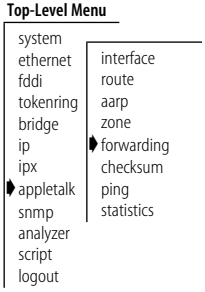
   Zone Test GmbH is assigned to 1 network
      12010-12012

   Zone Madrid3Com is assigned to 1 network
      14010-14029

   Zone NSDEng is assigned to 1 network
      32910-32910
```

Configuring Forwarding

You can control whether the router forwards or discards AppleTalk packets addressed to other hosts. When you enable forwarding, the router processes packets as usual, forwarding AppleTalk packets from one subnet to another when required. When you disable IP forwarding, the router discards any AppleTalk packets not addressed directly to one of its defined interfaces.



1. At the Administration Console's top-level menu, enter:

   `appletalk forwarding`

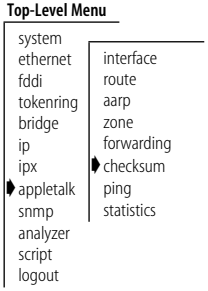
2. Enter the slot(s) of the EFSM(s) for which you want to enable AppleTalk forwarding.

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

Configuring Checksum

Checksum is a simple method used for detecting errors in the transmission of data. Checksum generation totals the bytes comprising the data and adds this sum to the end of the data packet. Checksum verification allows you to verify the integrity of the data that is routed. You can enable or disable checksum generation and verification states.

To enable or disable checksum generation/verification:



1. At the Administration Console's top-level menu, enter:

   `appletalk checksum`

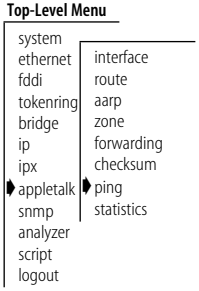
2. Enter **`enable`** or **`disable`** at the checksum generation prompt.

3. Enter **`enable`** or **`disable`** at the checksum verification prompt.

Pinging an AppleTalk Node

The AppleTalk Echo Protocol (AEP) sends a datagram (an Echo Request) from one node to another, which causes the destination node to return or *echo*, the datagram (an Echo Reply) to the sender. This allows you to determine whether a node is accessible before any sessions are started.

To ping an AppleTalk node:



1 At the Administration Console's top-level menu, enter:

```
appletalk ping
```

You are prompted for a node address.

2 Enter the address of the node you want to ping.

If the node is accessible, you receive a response.

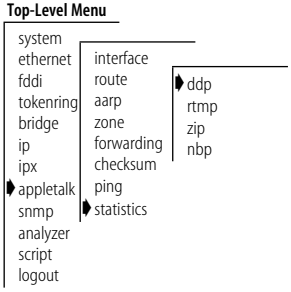
Viewing Appletalk Statistics

You can view statistics specific to the following AppleTalk protocols:

- Datagram Delivery Protocol (DDP)
- Routing Table Maintenance Protocol (RTMP)
- Zone Information Protocol (ZIP)
- Name Binding Protocol (NBP)

Displaying DDP Statistics

To display DDP statistics:



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

```
appletalk statistics ddp
```

2 Enter the slot(s) of the EFSM(s) for which you want to view AppleTalk statistics. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

An example of summary statistics is shown below:

```
Slot 2 - DDP forwarding is enabled.

      inReceives     inForwards       inLocals     inNoRoutes
          131131         113171          17906             22

     inNoClients    inTooShorts     inTooLongs    inShortDdps
               0              0              0              0

    inCsumErrors  inBcastErrors      inTooFars     inDiscards
               0              0              0             54

       outLocals
           15600
```

The AppleTalk DDP statistics you can view are described in Table 9-1:

Table 9-1 AppleTalk Statistics

Field	Description
inReceives	Total number of packets received, including those with errors
inForwards	Total number of packets forwarded, including those with errors
inLocals	Number of DDP datagrams for which this entity was their final DDP destination
inNoRoutes	Number of DDP datagrams dropped because a route could not be found
inNoClients	Number of DDP datagrams dropped because of an unknown DDP type
inTooShorts	Number of input DDP datagrams dropped because the received data length was less than the data length specified in the DDP header or the received data length was less than the length of the expected DDP header
inTooLongs	Number of input DDP datagrams dropped because they exceeded the maximum DDP datagram size
inShortDdps	Number of input DDP datagrams dropped because this entity was not their final destination and their type was short DDP
inCsumErrors	Number of DDP datagrams which were dropped because of a checksum error
inBcastErrors	Number of DDP datagrams for which this DDP entity was their final destination, and which were dropped because of a broadcast error

(continued)

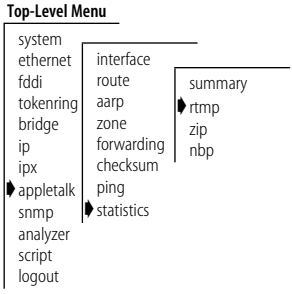
Table 9-1 AppleTalk Statistics (continued)

Field	Description
inTooFars	Number of input datagrams dropped because this entity was not their final destination and their hop count would exceed 15
inDiscards	Number of DDP Datagrams thrown out during the routing process
outLocals	Number of host-generated DDP datagrams

Displaying RTMP Information

To display RTMP statistics:

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




```
appletalk statistics rtmp
```

2 Enter the slot(s) of the EFSM(s) for which you want to view RTMP statistics. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

An example of summary statistics is shown below:

```
Slot 2 - DDP forwarding is enabled.

          inDatas     inRequests       outDatas    outRequests
             7204              0           4865              6

      routeEqChgs  routeLessChgs   routeDeletes routeOverflows
                0              0              0              0

    inVersionErrs    inOtherErrs
                0            119
```

The RTMP statistics you can view are described in Table 9-2:

Table 9-2 RTMP Statistics

Field	Description
inDatas	Number of good RTMP data packets received
inRequests	Number of good RTMP request packets received
outDatas	Number of good RTMP data packets sent
outRequests	Number of RTMP request packets sent

(continued)

Table 9-2 RTMP Statistics (continued)

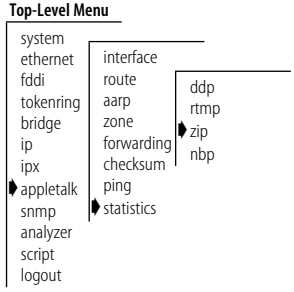
Field	Description
routeEqChgs	Number of times RTMP changes the Next Internet Router in a routing entry because the hop count advertised in a routing table was equal to the current hop count for a particular network
routeLessChgs	Number of times RTMP changes the Next Internet Router in a routing entry because the hop count advertised in a routing table was less than the current hop count for a particular network
routeDeletes	Number of times RTMP deletes a route because it was aged out of the table
routeOverflows	Number of times RTMP attempted to add a route to the RTMP table but failed due to lack of space
inVersionErrs	Number of RTMP packets received that were rejected due to a version mismatch
inOtherErrs	Number of RTMP packets received that were rejected for an error other than due to a version mismatch

Displaying ZIP Information

To display ZIP statistics:

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

```
appletalk statistics zip
```



2 Enter the slot(s) of the EFSM(s) for which you want to view ZIP statistics. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

An example of summary statistics is shown below:

```
Slot 2 - DDP forwarding is enabled.

        inQueries      inReplies    inExReplies  inGniRequests
              248             14              0            182

     inGniReplies   inLocalZones    inZoneLists
               22             30              0

      inObsoletes     inZoneCons     inZoneInvs       inErrors
                0              0             22              0

       outQueries     outReplies   outExReplies outGniRequests
                9              0            277             13

    outGniReplies  outLocalZones   outZoneLists
              182              0             30

      outZoneInvs    outAddrInvs
```

The ZIP statistics you can view are described in Table 9-3:

Table 9-3 ZIP Statistics

Field	Description
inQueries	Number of ZIP queries received
inReplies	Number of ZIP replies received
inExReplies	Number of ZIP extended replies received
inGniRequests	Number of ZIP GetNetInfo request packets received
inGniReplies	Number of ZIP GetNetInfo reply packets received
inLocalZones	Number of Zip GetLocalZones requests packets received
inZoneLists	Number of Zip GetZoneLists requests packets received
inObsoletes	Number of ZIP Takedown or ZIP Bringup packets received
inZoneCons	Number of times a conflict has been detected between this entity's zone information and another entity's zone information
inZoneInvs	Number of times this entity has received a ZIP GetNetInfo reply with the zone invalid bit set because the corresponding GetNetInfo request had an invalid zone name
inErrors	Number of ZIP packets received that were rejected for any error
outQueries	Number of ZIP queries sent

(continued)

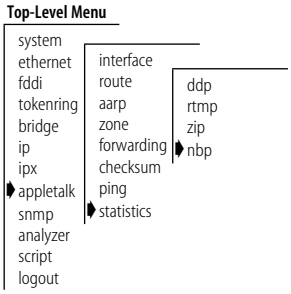
Table 9-3 ZIP Statistics (continued)

Field	Description
outReplies	Number of ZIP replies sent
outExReplies	Number of ZIP extended replies sent
outGniRequests	Number of ZIP GetNetInfo packets sent
outGniReplies	Number of ZIP GetNetInfo reply packets sent out of this port
outzoneInvs	Number of times this entity has sent a ZIP GetNetInfo reply with the zone invalid bit set in response to a GetNetInfo request with an invalid zone name
outAddrInvs	Number of times this entity had to broadcast a ZIP GetNetInfo reply because the GetNetInfo request had an invalid address

Displaying NBP Information

The NBP handles the translations between the numeric internet address and the alphanumeric entity names used by AppleTalk.

To display NBP statistics:



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

   `appletalk statistics nbp`

2. Enter the slot(s) of the EFSM(s) for which you want to view NBP statistics. Separate non-consecutive slots with commas (,). Enter a consecutive series of slots using a dash (-).

   An example of summary statistics is shown below:

```
Slot 2 - DDP forwarding is enabled.

       inLkupReqs    inBcastReqs      inFwdReqs  inLkupReplies
             3093            611           5951              0

         inErrors
                0
```

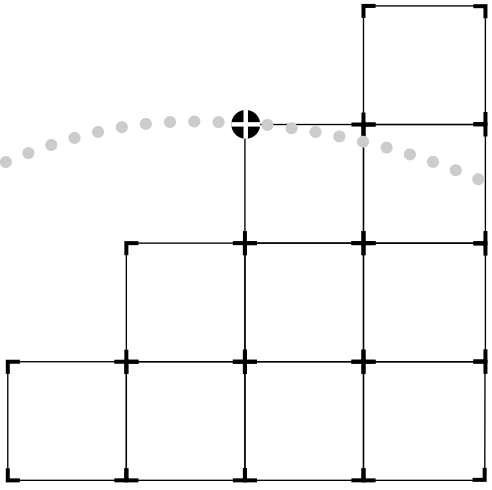
The NBP statistics you can view are described in Table 9-4:

Table 9-4 NBP Statistics

Field	Description
inLkupReqs	Number of NBP Lookup Requests received
inBcastsReqs	Number of NBP Broadcast Requests received
inFwdReqs	Number of NBP Forward Requests received
inLkupReplies	Number of NBP Lookup Replies received
inErrors	Number of NBP packets received that were rejected for any error

IV Appendix

Appendix A Technical Support



A 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®
- 3ComFacts℠ 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

Country	Data Rate	Telephone Number
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

Country	Telephone Number
U.K.	(44) (1442) 278279
U.S.	(1) (408) 727 7021

Local access numbers are available within the following countries:

Country	Telephone Number
Australia	800 123853
Belgium	0800 71279
Denmark	800 17319
Finland	98 001 4444
France	05 90 81 58
Germany	0130 8180 63
Italy	1678 99085

Country	Telephone Number
Netherlands	06 0228049
Norway	800 11062
Portugal	0505 442607
Russia (Moscow only)	956 0815
Spain	900 964445
Sweden	020 792954
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
Australia (Sydney)	(61) (2) 959 3020
(Melbourne)	(61) (3) 653 9515
Belgium*	0800 71429
Brazil	(55) (11) 546 0869
Canada	(905) 882 9964
Denmark*	800 17309
Finland*	0800 113153
France*	05 917959
Germany*	0130 821502
Hong Kong	(852) 868 9111
Ireland*	1 800 553117
Italy*	1678 79489

Country	Telephone Number
Japan	(81) (3) 3345 7251
Mexico	(525) 531 0591
Netherlands*	06 0227788
Norway*	800 13376
Singapore	(65) 538 9368
South Africa	(27) (11) 803 7404
Spain*	900 983125
Sweden*	120 795482
Taiwan	(886) (2) 577 4352
United Arab Emirates	(971) (4) 349049
U.K.*	0800 966197
U.S.	(1) (408) 492 1790

* These numbers are toll-free.

Returning Products for Repair

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) A-1
3Com sales offices A-4
3ComFacts A-2

A

AARP 6-10
AARP cache
  administering 9-6
  displaying 9-7
  removing an entry from 9-8
address
  classes 4-3
  IP 7-1
  IP to MAC, translating 7-9
  MAC 3-3
  network 3-3
Address Resolution Protocol. *See* ARP
Administration Console
  menu descriptions 1-2
  top-level menu 1-2
ADSP 6-10
AEP 6-8
AppleTalk
  address resolution protocol (AARP) 6-10
  checksum 9-10
  configuring forwarding 9-10
  data stream protocol (ADSP) 6-10
  echo protocol (AEP) 6-8
  interface, displaying 9-3
  main menu 1-5
  name binding protocol (NBP) 6-9
  network layer 6-6
  nodes 6-2
  physical layer 6-5
  printer access protocol (PAP) 6-10
  protocols, about 6-1
  protocols, and OSI levels 6-4
  routing table maintenance protocol (RTMP) 6-6
  routing tables 6-8
  session layer protocol (ASP) 6-10
  statistics, viewing 9-11
  transaction protocol (ATP) 6-9
  zone information protocol (ZIP) 6-9
  zones 6-3
AppleTalk networks 6-2
  extended 6-2
  nonextended 6-2
AppleTalk node
  pinging an 9-11
AppleTalk routing 6-1
ARP
  cache 4-7
  defined 4-7, 7-9
  location in OSI reference model 4-1
  reply 4-8
  request 4-8
  See also ARP cache 7-9
ARP cache 4-7, 7-9
  displaying 7-9
  displaying contents 7-9
  flushing 7-10
  removing an entry from 7-10, 7-12, 7-13
ASP 6-10
ATP 6-9

B

BOOTP relay threshhold 7-13
bridge
  menus 1-3
bridging/routing
  LANplex model 3-4
  traditional model 3-4
broadcast address 7-2
bulletin board service A-1

C

checksum
  configuring AppleTalk 9-10
chooser, Macintosh 6-2
CompuServe A-2

conventions
notice icons 2
text 2
cost
of IP interface 7-2
See also metric

D

datagram delivery protocol 6-6
datagrams, statistics 7-16
data-link layer 4-1
DDP statistics 9-11
default route, IP
defined 4-7, 7-6
removing 7-9
setting 7-8
direct, route status 7-6
documentation
for the LANplex system 3
DOS
copying software to 2-3
software media 2-1
dynamic routes 4-6, 5-14
See also RIP
See also SAP
dynamic routes, IPX 5-9

E

extended network numbers 6-2
extended switching, overview 1-1

F

fax service. *See* 3ComFacts
flushing
ARP cache 7-10
learned routes, AppleTalk 9-6
learned routes, IP 7-8
learned routes, IPX 8-7
for 8-8
forwarding
configuring AppleTalk 9-10
ftp
IP address 7-1
server in software load 2-4

G

gateway
IP address 7-5
routing table, and the 4-5
See also router

H

hard disk
copying software to 2-1

I

ICMP
defined 4-9
echo (request and reply) 7-15
Echo Reply 4-9
Echo Request 4-9
ping and 7-15
Redirect 4-9
Time Exceeded 4-9
installing software 2-1
interface
defining an IP 7-3
interface, AppleTalk
defining an 9-3
displaying an 9-3
removing an 9-4
interface, IP
displaying an 7-3
parts of 7-1
parts of an 7-2
removing definition 7-5
interface, IPX
defining an 8-3
displaying an 8-3
modifying an 8-4
removing an 8-4
Interior Gateway Protocols (IGP) 4-6, 5-9
Internet address. *See* IP address
Internet Control Message Protocol. *See* ICMP
Internet Protocol. *See* IP
intranetwork routing
diagram 3-2
IP
address translation 7-9
ARP cache 7-9
enabling forwarding 7-13
interface 7-1
main menu 1-3
pinging a station 7-15

RIP mode 7-14
routes 7-5
statistics, displaying 7-16
IP address
address classes 4-3
configuring 7-3
defined 4-2
derived from 4-2
division of network and host 4-2
example 4-4
for IP interface 7-1
network layer and the 4-1
RIP, and 4-6
routing table, and the 4-5
software installation, and 2-4
subnet mask, and the 4-3
subnet part 4-3
IP forwarding
configuring 7-13
IP interface
address 7-1
broadcast address 7-2
cost 7-2
defining 7-3
displaying 7-3
removing definition 7-5
subnet mask 7-2
IP route
default 7-6, 7-8
defining static 7-7
displaying table 7-6
gateway IP address 7-5
metric 7-5
removing from table 7-8
status 7-6
IP router
transmission process 4-2
IP routing
address classes 4-3
basic elements 4-2
ICMP 4-9
OSI reference model 4-1
references 4-10
router interface 4-4
routing table 4-5
transmission errors 4-9
IPX
forwarding statistics, displaying 8-17
main menu 1-4
RIP statistics, displaying 8-15
SAP statistics 8-16
IPX routing
and RIP 5-10
packet format 5-5
router interface 5-8
routing table 5-8
SAP, and 5-10
server table 5-13

L

LANplex
bridging/routing model 3-6
documentation 3
intranetwork router, as an 3-2
subnetting with 3-2
learned routes
flushing AppleTalk 9-6
flushing IP 7-8
flushing IPX 8-7
learned, IP route status 7-6

M

MAC address 3-3
ARP and 7-9
bridging in switching modules, and 3-6
compared to IP address 4-2
in ARP Request 4-8
located with ARP 4-7
use in IP routing 4-8
Macintosh, chooser 6-2
management
IP interface 7-1
media types 2-1
menu
AppleTalk main 1-5
bridge 1-3
IP main 1-3
IPX main 1-4
metric
defined 4-5
in IP routing table 7-5

N

name binding protocol 6-9
named entities 6-2
NBP 6-9
NetWare
defined 5-1
OSI reference model, and the 5-2
protocols 5-1 to 5-3
network address 3-3
network layer, and IP address 4-1

network layer, AppleTalk 6-6
network numbers
  extended 6-2
  nonextended 6-2
network supplier support A-3
nodes
  AppleTalk 6-2
nonextended network numbers 6-2

O

on-line technical services A-1
OSI 6-5
OSI Reference Model
  AppleTalk routing and 6-5
  IP routing and 4-1
  IPX routing and 5-2

P

PAP 6-10
physical layer, AppleTalk 6-5
pinging
  AppleTalk node 9-11
  IP station 7-15
port
  including in IP interface 7-2
printer access protocol 6-10
protocol
  AppleTalk routing table maintenance 6-6

R

references
  Comer 4-10
  Perlman 4-10
  routing RFCs 4-10
returning products for repair A-4
RIP
  active mode 7-14
  broadcast address, and 7-2
  default mode 7-14
  defined 4-6, 5-10
  displaying state 7-3
  off mode 7-14
  passive mode 7-14
  route configuration, and 4-6, 5-9
  setting mode 7-14
  using for dynamic routes 5-9
RIP statistics
  IPX RIP 8-15
route, IP
  default 7-6
  defining static 7-7
  gateway address 7-5
  metric 7-5
  removing default 7-9
  removing from table 7-8
  status 7-6
  subnet mask 7-5
route, IPX
  defining a static 8-6
  removing a 8-7
router interface, IP
  described 4-4
  diagram 4-5
  routing table, and the 4-5
router interface, IPX
  described 5-8
routers, seed 6-4
routing
  and bridging in switching modules 3-4
  and bridging, traditional model 3-4
  implementation in LANplex 3-4
  LANplex system, and the 3-1 to 3-7
  See also IP routing, IPX routing, and AppleTalk routing
Routing Information Protocol. *See* RIP
routing table, IP
  contents 4-5, 7-5
  default route 4-7
  default route, setting 7-8
  described 4-5
  display routes 7-6
  dynamic routes 4-6
  example 4-6
  flushing learned routes 7-8
  metric 4-5
  removing default route 7-9
  removing route 7-8
  static routes 4-6
routing table, IPX
  contents 5-8
  described 5-8
  displaying 8-6
  dynamic routes 5-9
  example 5-9
  flushing learned routes 8-7
  removing a route 8-7
  static routes 5-9
routing table, AppleTalk 6-8
RTMP
  description of 6-6

S

SAP
aging mechanism 5-14
packet structure 5-11
request handling 5-15
using for dynamic routes 5-14
SAP mode
setting 8-13
SAP statistics, displaying 8-16
seed routers 6-4
segmentation, increasing 3-3
server
defining a static IPX 8-8
server table
contents 5-13
described 5-13
displaying 8-8
Service Advertisement Protocol. *See* SAP
session layer protocols
AppleTalk 6-9
software
copying to hard disk 2-1
corrupted on install 2-5
installation 1-1, 2-1
loading time 2-4
static route, IP 4-6
defining 7-7
status of 7-6
static route, IPX 5-9
defining 8-6
static server, IPX
defining a 8-8
statistics
AppleTalk, viewing 9-11
IP 7-16
IPX forwarding 8-17
IPX SAP 8-16
ZIP, displaying 9-14
subnet mask
defined 4-3
diagram 4-4
example 4-4
for IP address 7-2
in IP routing table 7-5
in routing table 4-5
subnetting
defined 4-3
Ethernet switching and 3-2
subnet mask, and the 4-3
with the LANplex 3-2

T

technical support A-1
ThreeComForum A-2
timing out, IP route status 7-6
transmission errors
ICMP Redirect 4-9
reasons for 4-9

U

UNIX
copying software to 2-2
software media 2-1

Z

ZIP 6-9
statistics, displaying 9-14
zone information protocol (ZIP) 6-9
zone information table (ZIT) 6-9
displaying the 9-8
zone, AppleTalk
default 9-4
example of 6-3
naming 9-4

LIMITED WARRANTY

HARDWARE: 3Com warrants its hardware products to be free from defects in workmanship and materials, under normal use and service, for the following lengths of time from the date of purchase from 3Com or its Authorized Reseller:

Internetworking products	*One year*
Network adapters	*Lifetime*
Ethernet stackable hubs and Unmanaged Ethernet fixed port repeaters	*Lifetime One year if not registered*
Power supply and fans in these stackable hubs and unmanaged repeaters	*One year*
Other hardware products	*One year*
Spare parts and spares kits	*0 days*

If a product does not operate as warranted during the applicable warranty period, 3Com shall, at its option and expense, repair the defective product or part, deliver to Customer an equivalent product or part to replace the defective item, or refund to Customer the purchase price paid for the defective product. All products that are replaced will become the property of 3Com. Replacement products may be new or reconditioned. Any replaced or repaired product or part has a ninety (90) day warranty or the remainder of the initial warranty period, whichever is longer.

3Com shall not be responsible for any software, firmware, information, or memory data of Customer contained in, stored on, or integrated with any products returned to 3Com pursuant to any warranty.

SOFTWARE: 3Com warrants that the software programs licensed from it will perform in substantial conformance to the program specifications therefor for a period of ninety (90) days from the date of purchase from 3Com or its Authorized Reseller. 3Com warrants the magnetic media containing software against failure during the warranty period. No updates are provided. 3Com's sole obligation hereunder shall be (at 3Com's discretion) to refund the purchase price paid by Customer for any defective software products, or to replace any defective media with software which substantially conforms to 3Com's applicable published specifications. Customer assumes responsibility for the selection of the appropriate applications program and associated reference materials. 3Com makes no warranty that its software products will work in combination with any hardware or applications software products provided by third parties, that the operation of the software products will be uninterrupted or error free, or that all defects in the software products will be corrected. For any third party products listed in the 3Com software product documentation or specifications as being compatible, 3Com will make reasonable efforts to provide compatibility, except where the non-compatibility is caused by a "bug" or defect in the third party's product.

STANDARD WARRANTY SERVICE: Standard warranty service for hardware products may be obtained by delivering the defective product, accompanied by a copy of the dated proof of purchase, to 3Com's Corporate Service Center or to an Authorized 3Com Service Center during the applicable warranty period. Standard warranty service for software products may be obtained by telephoning 3Com's Corporate Service Center or an Authorized 3Com Service Center, within the warranty period. Products returned to 3Com's Corporate Service Center must be pre-authorized by 3Com with a Return Material Authorization (RMA) number marked on the outside of the package, and sent prepaid, insured, and packaged appropriately for safe shipment. The repaired or replaced item will be shipped to Customer, at 3Com's expense, not later than thirty (30) days after receipt by 3Com.

WARRANTIES EXCLUSIVE: IF A 3COM PRODUCT DOES NOT OPERATE AS WARRANTED ABOVE, CUSTOMER'S SOLE REMEDY SHALL BE REPAIR, REPLACEMENT, OR REFUND OF THE PURCHASE PRICE PAID, AT 3COM'S OPTION. THE FOREGOING WARRANTIES AND REMEDIES ARE EXCLUSIVE AND ARE IN LIEU OF ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, EITHER IN FACT OR BY OPERATION OF LAW, STATUTORY OR OTHERWISE, INCLUDING WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. 3COM NEITHER ASSUMES NOR AUTHORIZES ANY OTHER PERSON TO ASSUME FOR IT ANY OTHER LIABILITY IN CONNECTION WITH THE SALE, INSTALLATION, MAINTENANCE OR USE OF ITS PRODUCTS.

3COM SHALL NOT BE LIABLE UNDER THIS WARRANTY IF ITS TESTING AND EXAMINATION DISCLOSE THAT THE ALLEGED DEFECT IN THE PRODUCT DOES NOT EXIST OR WAS CAUSED BY CUSTOMER'S OR ANY THIRD PERSON'S MISUSE, NEGLECT, IMPROPER INSTALLATION OR TESTING, UNAUTHORIZED ATTEMPTS TO REPAIR, OR ANY OTHER CAUSE BEYOND THE RANGE OF THE INTENDED USE, OR BY ACCIDENT, FIRE, LIGHTNING, OR OTHER HAZARD.

LIMITATION OF LIABILITY: IN NO EVENT, WHETHER BASED IN CONTRACT OR TORT (INCLUDING NEGLIGENCE) SHALL 3COM BE LIABLE FOR INCIDENTAL, CONSEQUENTIAL, INDIRECT, SPECIAL, OR PUNITIVE DAMAGES OF ANY KIND, OR FOR LOSS OF REVENUE, LOSS OF BUSINESS, OR OTHER FINANCIAL LOSS ARISING OUT OF OR IN CONNECTION WITH THE SALE, INSTALLATION, MAINTENANCE, USE, PERFORMANCE, FAILURE, OR INTERRUPTION OF ITS PRODUCTS, EVEN IF 3COM OR ITS AUTHORIZED RESELLER HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Some states do not allow the exclusion of implied warranties or the limitation of incidental or consequential damages for consumer products, so the above limitations and exclusions may not apply to you. This warranty gives you specific legal rights which may vary from state to state.

GOVERNING LAW: This Limited Warranty shall be governed by the laws of the state of California.

3Com Corporation
5400 Bayfront Plaza
Santa Clara, CA 95052-8145
(408) 764-5000
1/1/94

Artisan Technology Group is an independent supplier of quality pre-owned equipment

Gold-standard solutions

Extend the life of your critical industrial, commercial, and military systems with our superior service and support.

We buy equipment

Planning to upgrade your current equipment? Have surplus equipment taking up shelf space? We'll give it a new home.

Learn more!

Visit us at **artisantg.com** for more info on price quotes, drivers, technical specifications, manuals, and documentation.

Artisan Scientific Corporation dba Artisan Technology Group is not an affiliate, representative, or authorized distributor for any manufacturer listed herein.

We're here to make your life easier. How can we help you today?

(217) 352-9330 | sales@artisantg.com | artisantg.com

ARTISAN® TECHNOLOGY GROUP