




CoreBuilder™ 2500 Extended Switching User Guide

Extended Switching software
Revision 8.3.1



<http://www.3com.com/>

Part No. 10005361
Published May 1998



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

Copyright © 1998, 3Com Corporation. 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 written 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, term, or condition of any kind, either implied or expressed, including, but not limited to, the implied warranties, terms or conditions of merchantability, satisfactory quality, 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.

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

UNITED STATES GOVERNMENT LEGEND

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

All technical data and computer software are commercial in nature and developed solely at private expense. Software is delivered as "Commercial Computer Software" as defined in DFARS 252.227-7014 (June 1995) or as a "commercial item" as defined in FAR 2.101(a) and as such is provided with only such rights as are provided in 3Com's standard commercial license for the Software. Technical data is provided with limited rights only as provided in DFAR 252.227-7015 (Nov 1995) or FAR 52.227-14 (June 1987), whichever is applicable. You agree not to remove or deface any portion of any legend provided on any licensed program or documentation contained in, or delivered to you in conjunction with, this User Guide.

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

3Com, the 3Com logo, Net Age, and Transcend are registered trademarks of 3Com Corporation. CoreBuilder is a trademark of 3Com Corporation. 3ComFacts is a service mark of 3Com Corporation.

Apple, AppleTalk, EtherTalk, Macintosh, and TokenTalk are registered trademarks of Apple Computer, Inc. Banyan and VINES are registered trademarks of Banyan Systems. DECnet is a trademark of Digital Equipment Corporation. NetBIOS is a trademark of Micro Computer Systems, Inc. Netware and Novell are registered trademarks of Novell Inc. Sun and Solaris are registered trademarks of Sun Microsystems, Inc. Xerox is a registered trademark of Xerox Corporation.

All other company and product names may be trademarks of the respective companies with which they are associated.

Guide written by Colleen McCaffrey, Dave Powell, and Beth Britt. Edited by Bonnie Jo Collins and Ben Mann. Illustrated and produced by Corrinne Hamilton.

CONTENTS

ABOUT THIS GUIDE

Introduction	1
How to Use This Guide	2
Conventions	2
Year 2000 Compliance	3
CoreBuilder 2500 Documentation	4
Paper Documents	4
Documents on CD	5
Documents on the Web	5
Documentation Comments	5

PART I GETTING STARTED

1 COREBUILDER 2500 EXTENDED SWITCHING FEATURES

About CoreBuilder 2500 Extended Switching	1-1
Using Menus	1-2
Extended Switching Features	1-3
Bridge Menus	1-3
IP Menus	1-4
IPX Menus	1-7
AppleTalk Menus	1-8

PART II VIRTUAL LAN TECHNOLOGY

2 VLANS ON THE COREBUILDER 2500 SYSTEM

About VLANs	2-1
Types of VLANs	2-1
Port Group VLANs	2-2
MAC Address Group VLANs	2-2
Application-Oriented VLANs	2-2
Protocol-Sensitive VLANs	2-3
CoreBuilder 2500 Protocol-Sensitive VLAN Configuration	2-3
Protocol Suite	2-3
Switch Ports	2-4
Layer 3 Addressing Information	2-4
Default VLAN	2-5
Modifying the Default VLAN	2-5
Flooding Decisions in the CoreBuilder 2500	2-5
VLAN Exception Flooding	2-6
Overlapped IP VLANs	2-7
Routing Between VLANs	2-8
VLAN Configuration Examples	2-9

PART III ABOUT ROUTING PROTOCOLS

3 BRIDGING AND ROUTING IN THE COREBUILDER 2500 SYSTEM

What Is Routing?	3-1
CoreBuilder 2500 in a Subnetworked Environment	3-2
Integrating Bridging and Routing	3-3
Bridging and Routing Models	3-4
Traditional Bridging and Routing	3-4
CoreBuilder 2500 Bridging and Routing	3-6

4 ROUTING WITH IP TECHNOLOGY

IP Routing and the OSI Reference Model	4-1
Elements of IP Routing	4-2
IP Addresses	4-2
Network Part	4-3
Subnetwork Part	4-3
Router Interfaces	4-4
Routing Table	4-5
Default Route	4-6
Address Resolution Protocol (ARP)	4-7
IP Routing Transmission Errors	4-8
Routing with Classical IP over ATM	4-9
About Logical IP Subnets (LISs)	4-9
ATM ARP Servers	4-10
Forwarding to Nodes Within a LIS	4-10
IP Routing References	4-11

5 ROUTING WITH IP MULTICAST

About IP Multicast Routing	5-1
IGMP	5-2
DVMRP	5-2
MBONE	5-2
Multicast Routing Algorithms	5-3
Flooding	5-3
Spanning Tree	5-3
Reverse Path Forwarding	5-4
Pruning	5-4
Multicast Interfaces	5-5
DVMRP Metric Value	5-5
Time-To-Live (TTL) Threshold	5-5
Rate Limit	5-5
Multicast Tunnels	5-5

6 ROUTING WITH IPX

IPX Routing in the NetWare Environment	6-1
Internet Packet Exchange (IPX)	6-3
Routing Information Protocol (RIP)	6-3
Service Advertising Protocol (SAP)	6-3
How IPX Routing Works	6-4
IPX Packet Format	6-4
IPX Packet Delivery	6-5
Sending Node's Responsibility	6-6
Router's Responsibility	6-7
Elements of IPX Routing	6-7
Router Interfaces	6-7
Routing Tables	6-8
Generating Routing Table Information	6-9
Selecting the Best Route	6-10
Service Advertising Protocol in IPX	6-10
Internetwork Service Information	6-10
SAP Packet Structure	6-10
Server Information Tables	6-12
Server Information Maintenance	6-14

7 ROUTING WITH OSPF

Elements of OSPF Routing	7-3
Autonomous Systems	7-3
Areas	7-3
Area Border Routers	7-4
Routing Databases	7-5
Neighbors	7-6
Protocol Packets	7-7
Router Types	7-7
Router IDs	7-8
Interface Characteristics	7-9
Mode	7-9
Priority	7-9
Area ID	7-10

Cost	7-10
Delay	7-10
Hello Timer	7-10
Retransmit Timer	7-10
Dead Interval	7-11
Password	7-11
Stub Default Metrics	7-11
Virtual Links	7-11
How OSPF Routing Works	7-12
Starting Up	7-12
Finding Neighbors	7-12
Establishing Adjacencies	7-12
Electing the Backup Designated Router	7-12
Electing the Designated Router	7-13
Calculating Shortest Path Trees	7-13
Routing Packets	7-14

8 ROUTING WITH APPLE TALK TECHNOLOGY

About AppleTalk	8-1
AppleTalk Network Elements	8-1
AppleTalk Networks	8-2
AppleTalk Nodes	8-2
Named Entities	8-2
AppleTalk Zones	8-2
Seed Routers	8-4
AppleTalk Protocols	8-4
Physical Layer Protocols	8-5
Link Layer Protocols	8-5
Network Layer Protocols	8-5
Transport Layer Protocols	8-5
Routing Table Maintenance Protocol (RTMP)	8-6
AppleTalk Echo Protocol (AEP)	8-8
AppleTalk Transaction Protocol (ATP)	8-8
Name Binding Protocol (NBP)	8-8

Session Layer Protocols	8-8
AppleTalk Data Stream Protocol (ADSP)	8-8
Zone Information Protocol (ZIP)	8-8
AppleTalk Session Protocol (ASP)	8-9
Printer Access Protocol (PAP)	8-9
Presentation Layer Protocols	8-9
About AARP	8-10

PART IV ADMINISTERING EXTENDED SWITCHING FEATURES

9 ADMINISTERING VLANS

Displaying VLAN Information	9-1
Defining VLAN Information	9-3
Modifying VLAN Information	9-4
Removing a VLAN Definition	9-5

10 ADMINISTERING IP ROUTING

Administering Interfaces	10-2
IP VLAN Interfaces	10-2
IP LIS Interfaces	10-2
Interface Characteristics	10-2
Displaying Interfaces	10-3
Defining an IP VLAN Interface	10-4
Defining an IP LIS Interface	10-5
Modifying an Interface	10-6
Removing an Interface	10-6
Adding a Permanent Virtual Circuit (PVC)	10-7
Removing a Permanent Virtual Circuit (PVC)	10-7
Administering Routes	10-8
Displaying the Routing Table	10-9
Defining a Static Route	10-9
Removing a Route	10-10
Flushing a Route	10-10
Setting the Default Route	10-10
Removing the Default Route	10-11

Administering the ARP Cache	10-11
Displaying the ARP Cache	10-11
Removing an ARP Cache Entry	10-12
Flushing the ARP Cache	10-12
Administering ATM ARP Servers	10-13
Displaying ATM ARP Servers	10-13
Defining an ATM ARP Server	10-13
Removing an ATM ARP Server	10-14
Displaying the ATM ARP Cache	10-14
Removing an ATM ARP Cache Entry	10-15
Flushing the ATM ARP Cache	10-15
Administering UDP Helper	10-15
Displaying UDP Helper Information	10-16
Defining a Port and IP Routing Address	10-16
Removing a Port Number and IP Routing Address	10-16
Setting the BOOTP Hop Count Limit	10-17
Setting the BOOTP Relay Threshold	10-17
Configuring Overlapped Interfaces	10-18
Enabling and Disabling IP Routing	10-18
Enabling and Disabling ICMP Router Discovery	10-19
Configuring RIP	10-19
Displaying the RIP Interface Configuration	10-19
Setting the RIP Mode	10-20
Setting the RIP Interface Cost	10-20
Setting the Poison Reverse Mode	10-21
Adding an RIP Advertisement Address	10-21
Removing an RIP Advertisement Address	10-22
Displaying RIP General Statistics	10-22
Pinging an IP Station	10-23
Displaying IP Statistics	10-24

11 ADMINISTERING IP MULTICAST ROUTING

- Enabling and Disabling DVMRP 11-1
- Enabling and Disabling IGMP 11-2
- Administering IP Multicast Interfaces 11-3
 - Multicast Interface Characteristics 11-3
 - Displaying Multicast Interfaces 11-3
 - Enabling Multicast Interfaces 11-4
 - Disabling Multicast Interfaces 11-4
- Administering Multicast Tunnels 11-5
 - Displaying Multicast Tunnels 11-5
 - Defining a Multicast Tunnel 11-6
 - Removing a Multicast Tunnel 11-6
- Displaying Routes 11-7
- Displaying the Multicast Cache 11-8

12 ADMINISTERING IPX ROUTING

- Administering Interfaces 12-2
 - Displaying IPX Interfaces 12-3
 - Defining an Interface 12-3
 - Modifying an Interface 12-4
 - Removing an Interface 12-4
- Administering Routes 12-4
 - Displaying the Routing Table 12-5
 - Defining a Static Route 12-5
 - Removing a Route 12-6
 - Flushing Routes 12-6
- Administering Servers 12-6
 - Displaying the Server Table 12-7
 - Defining a Static Server 12-7
 - Removing a Server 12-8
 - Flushing Servers 12-8
- Setting IPX Forwarding 12-8
- Setting the RIP Mode 12-9
- Setting the Enhanced RIP Mode 12-9
- Setting RIP Triggered Updates 12-10
- Setting the SAP Mode 12-10
- Setting SAP Triggered Updates 12-11

Displaying Statistics	12-11
Displaying IPX Summary Statistics	12-11
Displaying IPX RIP Statistics	12-12
Displaying IPX SAP Statistics	12-13
Displaying IPX Forwarding Statistics	12-14

13 ADMINISTERING OSPF ROUTING

Administering Areas	13-1
Displaying Areas	13-2
Defining Areas	13-2
Modifying Areas	13-3
Removing Areas	13-3
Adding Network Ranges	13-4
Modifying Network Ranges	13-4
Removing Network Ranges	13-5
Setting the Default Route Metric	13-5
Displaying Default Route Metrics	13-5
Defining Default Route Metrics	13-6
Removing Default Route Metrics	13-6
Configuring OSPF Interfaces	13-6
Displaying OSPF Interface Information	13-7
Displaying OSPF Interface Statistics	13-8
Setting Modes	13-10
Setting Priorities	13-11
Setting Area IDs	13-11
Setting Costs	13-12
Setting Transmit Delays	13-12
Setting Hello Timers	13-13
Setting Retransmit Timers	13-13
Setting Dead Intervals	13-14
Setting Passwords	13-14

Displaying the Link State Database	13-15
Displaying a Database Summary	13-15
Displaying Router LSAs	13-16
Displaying Network LSAs	13-18
Displaying Summary Network LSAs	13-19
Displaying External Network LSAs	13-20
Administering Neighbors	13-21
Displaying Neighbors	13-21
Adding Neighbors	13-22
Removing Neighbors	13-22
Setting OSPF Router IDs	13-23
Administering Memory Partitions	13-24
Displaying Memory Partitions	13-24
Modifying Memory Partitions	13-24
Administering Stub Default Metrics	13-25
Displaying Stub Default Metrics	13-25
Defining Stub Default Metrics	13-25
Removing Stub Default Metrics	13-25
Administering Virtual Links	13-26
Displaying Virtual Links	13-26
Defining Virtual Links	13-26
Modifying Virtual Links	13-27
Removing Virtual Links	13-27
Displaying OSPF Statistics	13-28

14 ADMINISTERING APPLE TALK ROUTING

Administering Interfaces	14-1
Displaying AppleTalk Interfaces	14-2
Defining an Interface	14-3
Removing an Interface	14-4
Administering Routes	14-4
Displaying the Routing Table	14-4
Flushing all Routes	14-5
Administering the AARP Cache	14-6
Displaying the AARP Cache	14-6
Removing an Address from the Cache	14-7
Flushing All Cache Entries	14-8

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

PART V RMON AND THE COREBUILDER 2500 SYSTEM

15 REMOTE MONITORING (RMON) TECHNOLOGY

Overview	15-1
RMON Benefits	15-2
RMON in the CoreBuilder 2500	15-2
RMON Groups	15-3
Statistics and axFDDIStatistics Groups	15-4
History and axFDDIHistory Groups	15-5
Alarm Group	15-5
Setting Alarm Thresholds	15-6
RMON Hysteresis Mechanism	15-7
Host Group	15-8
HostTopN Group	15-8
Matrix Group	15-8
Event Group	15-8
3Com Transcend RMON Agents	15-9
RMON Management Information Base (MIB)	15-10
MIB Objects	15-10

PART VI APPENDIX

A TECHNICAL SUPPORT

Online Technical Services	A-1
World Wide Web Site	A-1
3Com FTP Site	A-2
3Com Bulletin Board Service	A-2
Access by Analog Modem	A-2
Access by Digital Modem	A-3
3ComFacts Automated Fax Service	A-3
Support from Your Network Supplier	A-3
Support from 3Com	A-3
Returning Products for Repair	A-5

INDEX

ABOUT THIS GUIDE

Introduction

This *Extended Switching User Guide* explains the features of the CoreBuilder™ 2500 Extended Switching software. These features include IP, IP multicast, classical IP over ATM, IPX, AppleTalk routing, virtual LAN (VLAN) configuration, and remote monitoring (RMON).

Use this guide with the *CoreBuilder 2500 Administration Console User Guide* when you configure your system.



See the CoreBuilder 2500 Extended Switching Software Installation and Release Notes for information about how to install Extended Switching software on your CoreBuilder 2500 system.

This guide is intended for the system or network administrator who is responsible for configuring, using, and managing the CoreBuilder 2500 system. It assumes that you have 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 that accompany your product differs from the information in this guide, follow the instructions in 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
Virtual LANs (VLANs)	Chapter 2
General bridging and routing models	Chapter 3
IP routing	Chapter 4
IP multicast routing	Chapter 5
IPX routing	Chapter 6
OSPF routing	Chapter 7
AppleTalk routing	Chapter 8
Administering VLANs	Chapter 9
Administering IP routing	Chapter 10
Administering IP multicast routing	Chapter 11
Administering IPX routing	Chapter 12
Administering OSPF routing	Chapter 13
Administering AppleTalk routing	Chapter 14
Remote monitoring (RMON)	Chapter 15
3Com Technical Support	Appendix A

Conventions

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

Table 1 Notice Icons





Icon	Notice Type	Description
	Information note	Information that describes important features or instructions
	Caution	Information that alerts you to potential loss of data or potential damage to an application, system, or device
	Warning	Information that alerts you to potential personal injury

Table 2 Text Conventions

Convention	Description
Screen displays	This typeface represents information as it appears on the screen.
Commands	<p>The word "command" means that you must enter the command exactly as shown and then press Return or Enter. Commands appear in bold. Example:</p> <p>To display the default route metric, from the top level of the Administration Console, enter:</p> <p>ip ospf defaultRouteMetric display</p> <p> This guide always gives the full form of a command in uppercase and lowercase letters. However, you can abbreviate commands by entering just enough characters to distinguish one command from another similar command. Commands are not case sensitive.</p>
The words "enter" and "type"	When you see the word "enter" in this guide, you must type something, and then press Return or Enter. Do not press Return or Enter when an instruction simply says "type."
Keyboard key names	<p>If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example:</p> <p>Press Ctrl+Alt+Del.</p>
Words in <i>italics</i>	<p>Italics are used to:</p> <ul style="list-style-type: none"> ■ Emphasize a point. ■ Denote a new term at the place where it is defined in the text. ■ Identify menu names, menu commands, and software button names. Examples: <p>From the <i>Help</i> menu, select <i>Contents</i>.</p> <p>Click <i>OK</i>.</p>

Year 2000 Compliance

For information on Year 2000 compliance and 3Com products, visit the 3Com Year 2000 Web page:

<http://www.3com.com/products/yr2000.html>

CoreBuilder 2500 Documentation

This section describes the CoreBuilder 2500 system documentation set. Paper documents are shipped with your system. Additional documents are included on the 3Com compact disc. To order a paper copy of a document that you see on the compact disc, or to order additional compact discs, call your 3Com sales representative.

Paper Documents

These documents are shipped with your CoreBuilder 2500 system:

- *CoreBuilder 2500 Unpacking Instructions*

Instructions about how to unpack your CoreBuilder 2500 system; also an inventory list of the items that are shipped with your system.

- *CoreBuilder 2500 Software Installation and Release Notes*

Information about the software release, including new features, bug fixes, and any changes to the documentation.

- *CoreBuilder 2500 Getting Started Guide*

All the procedures necessary for installing, cabling, powering up, configuring management access to, and troubleshooting your CoreBuilder 2500 system.

- *CoreBuilder 2500 Intelligent Switching Administration Console Command Quick Reference card*

A summary of the Intelligent Switching commands for the CoreBuilder 2500 system Administration Console.

In addition, the package for each module contains a guide:

- *Module Installation Guides*

An overview, installation instructions, LED status information, and pin-out information for each module.

Documents on CD

In addition to the paper documents that are shipped with your product, the compact disc that comes with your system contains these books:

- Online versions of the paper guides that are shipped with your system
- *CoreBuilder 2500 Operation Guide*
Information to help you understand system management and administration, bridging, Fast Ethernet, and FDDI technology. Also, how these concepts are implemented in the CoreBuilder 2500 system.
- *CoreBuilder 2500 Administration Console User Guide*
Instructions about how to use the Administration Console to configure and manage your CoreBuilder 2500 system.
- *CoreBuilder 2500 Extended Switching User Guide* (this book)
Information about how the CoreBuilder 2500 system implements routing protocols, VLAN technology, and RMON technology. Also, information about using the Administration Console to configure these features.

Documents on the Web

You can view most 3Com documentation on the World Wide Web at our Web site:

<http://www.3Com.com>

Documentation Comments

Your suggestions are very important to us. They help make our documentation more useful to you. Please send e-mail comments about this document to 3Com at:


sdtechpubs_comments@3Com.com

Please include the following information when commenting:

- Document title
- Document part number (on the title page)
- Page number (if appropriate)


Example:


- *CoreBuilder 2500 Extended Switching User Guide*
- Part No. 10005361
- Page 2-5 (chapter 2, page 5)



GETTING STARTED

Chapter 1 CoreBuilder 2500 Extended Switching Features





COREBUILDER 2500 EXTENDED SWITCHING FEATURES

This chapter gives an overview of the Revision 8.3.1 Extended Switching software and describes the enhanced Administration Console menus.

About CoreBuilder 2500 Extended Switching

The CoreBuilder™ 2500 Extended Switching software replaces your existing CoreBuilder 2500 software and adds new functionality to your system. Extended Switching software contains all of the features of CoreBuilder 2500 Intelligent Switching software and adds these features:

- Virtual LANs (VLANs)
- Internet Protocol (IP) routing, an enhanced version that includes:
 - UDP Helper support for overlapped IP interfaces
 - Enhanced IP Routing Information Protocol (RIP) administration
- IP multicast routing
- Classical IP routing over Asynchronous Transfer Mode (ATM)
- Internet Packet Exchange (IPX) routing
- Open Shortest Path First (OSPF) routing
- AppleTalk routing
- Remote Monitoring (RMON)



The RMON feature is available only through an SNMP connection, not through the Administration Console. See Chapter 15 for details.



For information about how to gain access to online help, use scripts, and use the Administration Console, see the CoreBuilder 2500 Administration Console User Guide.



for information about how to install Extended Switching software on your CoreBuilder 2500 system, see the CoreBuilder 2500 Extended Switching Software Installation and Release Notes.

Using Menus

When you gain access to the Administration Console, the top-level menu appears. The Extended Switching software contains top-level menus and additions to the Intelligent Switching's `ethernet`, `fddi`, `bridge`, and `ip` menu options, as shown in Figure 1-1.

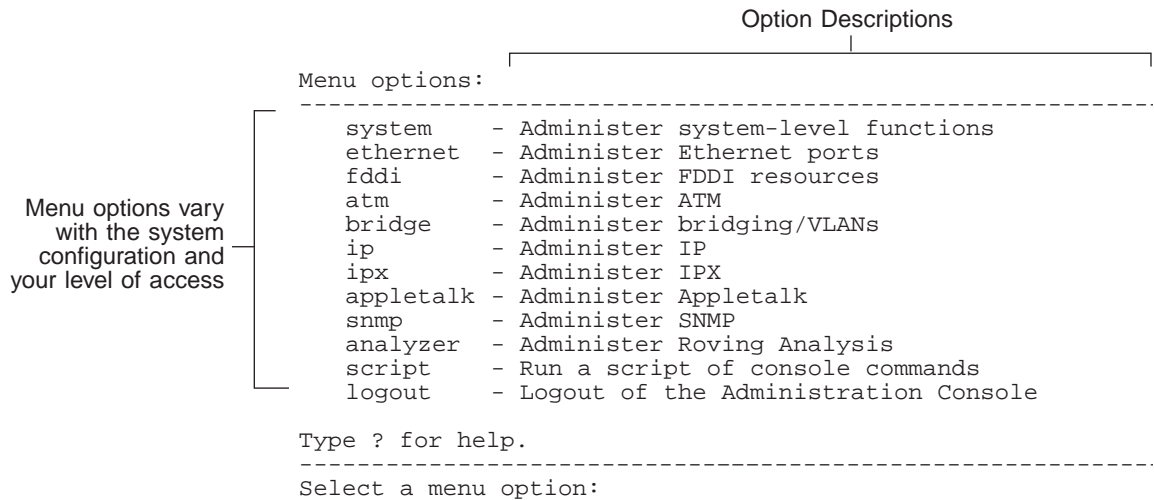


Figure 1-1 Top-level Menu Options for Extended Switching Software

Extended Switching Features

The rest of this chapter reviews the features in the CoreBuilder 2500 Extended Switching software. Subsequent chapters describe these features in greater detail.

Bridge Menus

From the `bridge` menu, you can configure and manage virtual LANs (VLANs). Figure 1-2 shows the `bridge` menu hierarchy. For example, to create a new VLAN, from the top level of the Administration Console, enter: **bridge vlan define**

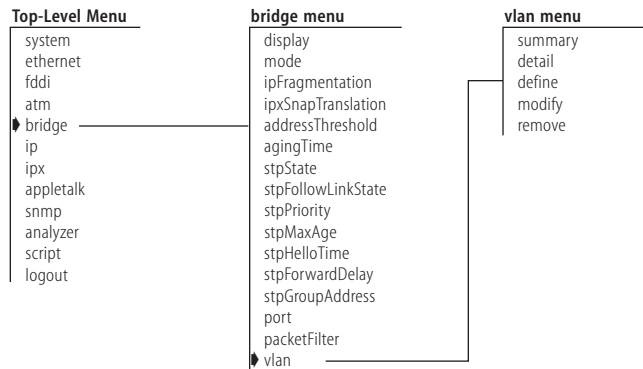
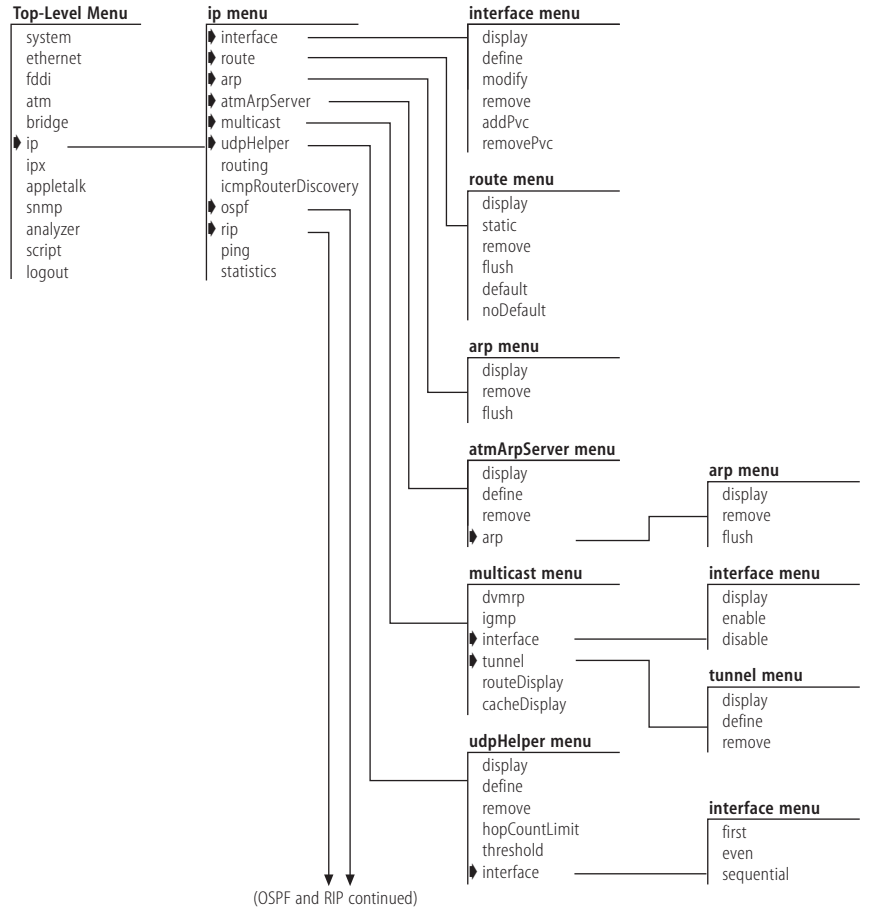


Figure 1-2 bridge Menus

For more information about VLANs, see Chapter 2 and Chapter 9.

- IP Menus** From the `ip` menu, you can configure and manage:
- Internet Protocol (IP) interfaces and routes (see Chapter 4)
 - Address Resolution Protocol (ARP) cache (see Chapter 4 and Chapter 10)
 - Multicast routing (see Chapter 5 and Chapter 11)
 - UDP Helper (see Chapter 10)
 - IP routing (see Chapter 4 and Chapter 10)
 - Open Shortest Path First (OSPF) routing (see Chapter 7 and Chapter 13)
 - Routing Information Protocol (RIP) (see Chapter 10 and Chapter 12)
 - Pinging (see Chapter 10)
 - Statistics displays (see Chapter 10)
 - ATM ARP servers, if you are running classical IP over ATM (see Chapter 4 and Chapter 10)

Figure 1-3 (on this page and the next) shows the `ip` menu hierarchy. For example, to define a new IP interface, from the top level of the Administration Console, enter: **ip interface define**



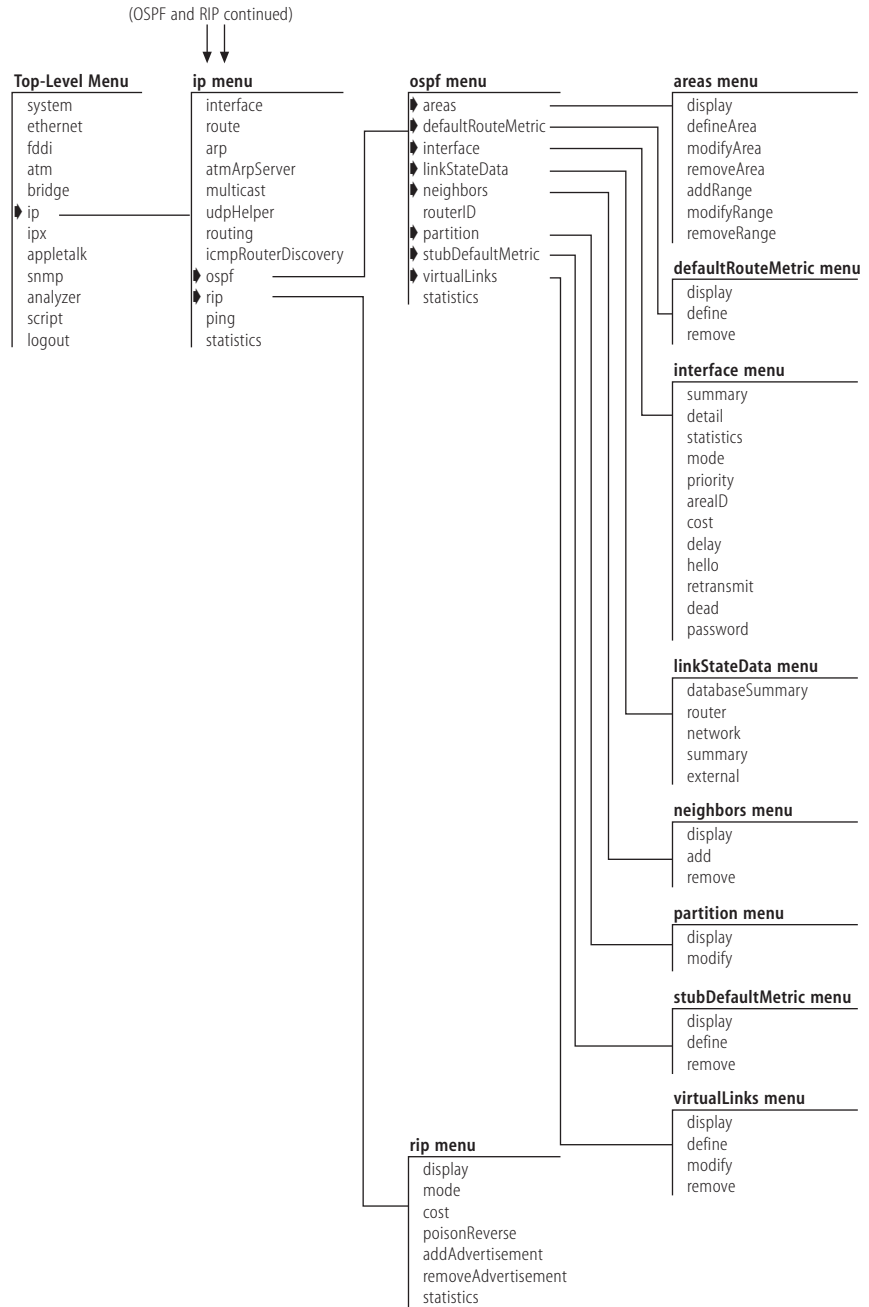


Figure 1-3 ip Menus

- IPX Menus** From the `ipx` menu, you can configure and manage:
- Internet Packet Exchange (IPX) interfaces, routes, and servers
 - Forwarding
 - Routing Information Protocol (RIP) and enhanced RIP
 - Service Advertising Protocol (SAP)
 - Statistics displays

Figure 1-4 shows the `ipx` menu hierarchy. For example, to define a new IPX interface, from the top level of the Administration Console, enter:

`ipx interface define`

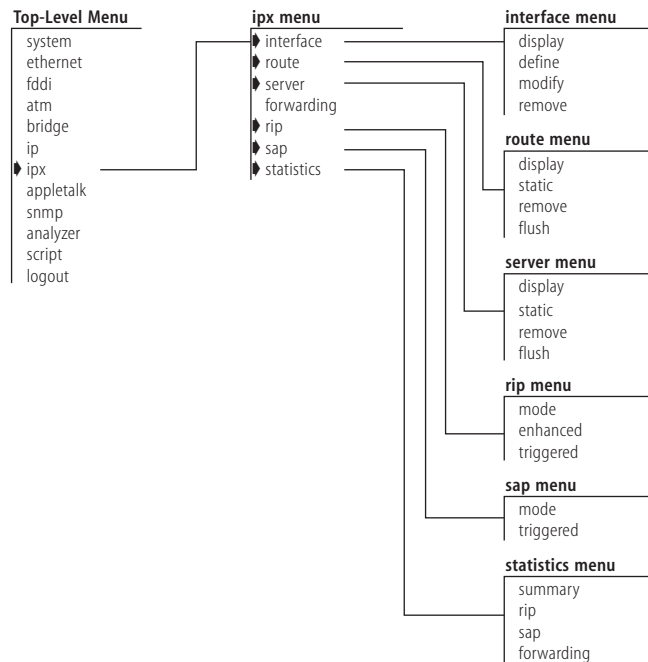


Figure 1-4 `ipx` Menus

For more information about IPX routing, see Chapter 6 and Chapter 12.

AppleTalk Menus From the `appletalk` menu, you can configure and manage:

- Interfaces, routes, and zones
- Address Resolution Protocol (ARP) cache
- Forwarding
- Checksum generation and verification
- Pinging
- Statistics reports

Figure 1-5 shows the `appletalk` menu hierarchy. For example, to define a new AppleTalk interface, from the top level of the Administration Console, enter: **appletalk interface define**

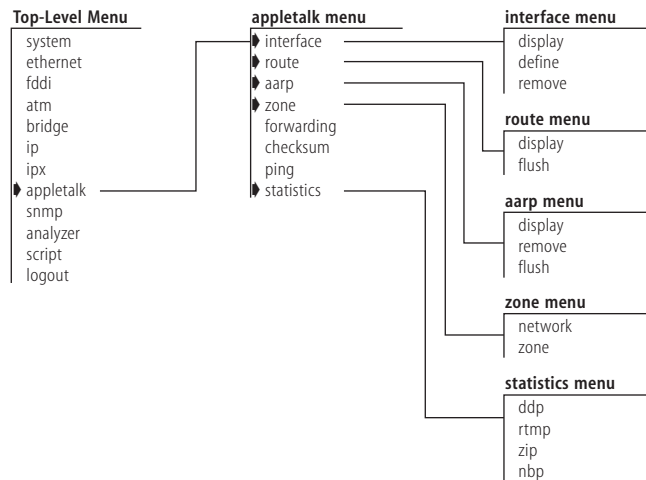




Figure 1-5 `appletalk` Menus

For more information about AppleTalk routing, see Chapter 8 and Chapter 14.



VIRTUAL LAN TECHNOLOGY

Chapter 2 VLANs on the CoreBuilder 2500 System



2

VLANS ON THE COREBUILDER 2500 SYSTEM

In this chapter:

- About VLANs: How VLANs operate in the CoreBuilder™ 2500 system
- VLAN Configuration Examples

About VLANs

Virtual LANs (VLANs) are logical subnetworks composed of selected CoreBuilder 2500 network interfaces. VLANs help minimize broadcast and multicast traffic across networks. VLANs also make it easier for you to move, add, and change end stations.

In the CoreBuilder 2500 system, VLANs allow you to:

- Create independent broadcast domains to optimize network performance and create firewalls
- Form flexible user groups independent of the users' physical network locations

Types of VLANs

You can use several types of VLANs to group users:

- Port group VLANs
- MAC address group VLANs
- Application-oriented VLANs
- Protocol-sensitive VLANs

The packet filtering capabilities in the CoreBuilder 2500 system provide support for port group, MAC address group, and application-oriented VLANs. For information about how to filter port groups and MAC address groups, see the Operation Guide and the Administration Console User Guide. For information about how to configure and manage protocol-sensitive VLANs, see the description of the Administration Console menus in Chapter 9.

Port Group VLANs

Port group VLANs, which connect one or more switch ports, require little configuration. Port groups are useful when traffic patterns are directly associated with particular ports. Port groups can benefit the network by restricting traffic based on a set of simple rules.

A port group VLAN groups all frames received on a port and keeps the frames within that port group, regardless of the data contained in the frames.

MAC Address Group VLANs

In a MAC address group VLAN, a switch filters by MAC addresses. This type of VLAN is very secure because you can configure these MAC address groups so that stations in the group can communicate only with each other or with specific network resources.

This type of VLAN is also easy to maintain because a VLAN association moves when a station moves. However, MAC address group VLANs may require complex configuration in comparison to other types of VLANs.

Application-Oriented VLANs

Using the CoreBuilder 2500 filtering capability, a switch can filter application-specific traffic such as telnet traffic or FTP traffic based on higher-layer information. To create this application-oriented VLAN, you configure packet filters that specify data and offsets of the data within received packets. For example, to use a filter on a particular port for all telnet traffic, create a filter that discards all TCP traffic that is received on the telnet port.

You can also use IP multicast routing and autocast VLANs to group IP multicast traffic for specific applications.

Protocol-Sensitive VLANs

The CoreBuilder 2500 system forwards, to all ports, any data that has a broadcast, multicast, or unknown destination address. This process is referred to as *bridge flooding*.

With protocol-sensitive VLANs, you can restrict flood traffic for routable and nonroutable protocols. Protocol-sensitive VLANs have a relatively simple configuration, grouping one or more switch ports together for a specified network Layer 3 protocol, such as IP or AppleTalk. These VLANs determine flooding based on the network layer protocol of the frame. In addition, for IP VLANs, you can make flooding decisions based on Layer 3 subnetwork address information.

In a multiprotocol environment, protocol-sensitive VLANs can effectively control broadcast and multicast flooding. They operate independently of each other. In addition, the same switch port can belong to multiple VLANs. For example, you can assign port 1 on a CoreBuilder 2500 system to several IP subnetwork VLANs: one IPX VLAN, one AppleTalk VLAN, and one NetBIOS VLAN.



Although two or more types of VLANs can coexist within the CoreBuilder 2500 system, when a switch evaluates received data in a multiple VLAN configuration, port group VLANs, MAC address group VLANs, and application-oriented VLANs always take precedence over protocol-sensitive VLANs.

CoreBuilder 2500 Protocol-Sensitive VLAN Configuration

The CoreBuilder 2500 protocol-sensitive VLAN configuration includes three elements: the protocol suite, the switch ports, and the Layer 3 addressing information for IP VLANs.

Protocol Suite

The protocol suite describes which protocol entities can comprise a protocol-sensitive VLAN. For example, CoreBuilder 2500 VLANs support the IP protocol suite, which is made up of the IP, ARP, and RARP protocols.

Table 2-1 lists the protocol suites that the CoreBuilder 2500 supports, as well as the protocol types included in each protocol suite.

Table 2-1 Supported Protocols for VLAN Configuration

Protocol Suite	Protocol Types
IP	IP, ARP, RARP (Ethernet II)
Novell IPX	IPX (Ethernet II, 802.2, 802.3, 802.3 SNAP)
AppleTalk	DDP, AARP (Ethernet II, 802.3 SNAP)
Xerox XNS	XNS IDP, XNS Address Translation, XNS Compatibility (Ethernet II, 802.3 SNAP)
DECnet	DEC MOP, DEC Phase IV, DEC LAT, DEC LAVC (Ethernet II, 802.3 SNAP)
SNA	SNA Services over Ethernet (Ethernet II, 802.2, 802.3 SNAP)
Banyan VINES	Banyan (Ethernet II, 802.3 SNAP)
X25	X.25 Layer 3 (Ethernet II, 802.3 SNAP)
NetBIOS	NetBIOS (802.2)
Default	Default (all protocol types)

Switch Ports

A group of switch ports is any combination of ports on a CoreBuilder 2500 system bridge. Included are switch ports created as ATM LAN Emulation Clients (ATM LECs). VLANs support only media implementations that run over switch (bridge) ports, for example, ATM Logical IP Subnets (ATM LISs).

Layer 3 Addressing Information

For IP VLANs only, the CoreBuilder 2500 system optionally supports configuring of individual IP VLANs with network layer subnetwork addresses. With this additional Layer 3 information, you can create independent IP VLANs that share the same switch ports for multiple IP VLANs. To distinguish among multiple IP VLANs on the same switch port, the CoreBuilder 2500 system floods data according to both the protocol (IP) and the Layer 3 information in the IP header. This configuration is discussed on page 2-7 in "Overlapped IP VLANs."

Default VLAN

When you start the CoreBuilder 2500 Extended Switching software, the system creates a default VLAN. Initially, the default VLAN includes all the system's switch ports. The CoreBuilder 2500 default VLAN defines:

- The flood domain for protocols that are not supported by any VLAN in the system
- The flood domain for protocols that are supported by a VLAN in the system but received on nonmember ports

Both cases represent exception flooding conditions that are described in the following sections.

Modifying the Default VLAN

If you insert a LAN card or create an ATM LEC, new switch ports can dynamically appear. When a new switch port that is not part of a default VLAN appears in the system at initialization, the system software adds that switch port to the first default VLAN defined in the system.



With CoreBuilder 2500 VLANs you can modify the initial default VLAN to form two or more subsets of switch ports. If you remove the default VLAN and no other VLANs are defined for the system, no flooding of traffic can occur.

Flooding Decisions in the CoreBuilder 2500

Protocol-sensitive VLANs directly affect how the CoreBuilder 2500 system performs flooding. Without protocol-sensitive VLANs, the flooding process forwards data to all switch ports in the system. With protocol-sensitive VLANs, the flooding process follows this model:

- When the system receives a frame that needs to be flooded, the system decodes the frame's protocol.
- If a VLAN for the frame's protocol exists in the system and the frame's source port is a member of the VLAN, then the system floods the frame to the ports that are assigned to that VLAN.
- If a VLAN for the frame's protocol exists in the system but the frame's source port is not a member of the VLAN, then the system floods the frame to the default VLAN that is assigned to that port.
- If there is no VLAN for the frame's protocol, then the system floods the frame to the default VLAN for the port that received it.

The following example shows how flooding occurs according to VLANs set up by protocol. The example assumes an 18-port switch.

VLAN Index	VLAN Protocol	VLAN Ports
1	Default	1 — 18
2	IP	1 — 12
3	IPX	11 — 16

Data received on this port	Is flooded on this VLAN	Because
IP — port 1	VLAN 2	The received IP data matches the IP VLAN on the source port.
IPX — port 11	VLAN 3	The received IPX data matches the IPX VLAN on the source port.
XNS — port 1	VLAN 1	The received XNS data matches no protocol VLAN, so the default VLAN is used.

VLAN Exception Flooding

Data for a protocol may arrive on a switch port that has no defined VLAN for that protocol. In such cases, called *VLAN exception flooding*, the default VLAN defines the flooding domain for the data, even if a VLAN for the protocol exists elsewhere in the system.

The following example shows how VLAN exception flooding occurs. The example assumes an 18-port switch.

VLAN Index	VLAN Protocol	VLAN Ports
1	Default	1 — 18
2	IP	1 — 10

Data received on this port	Is flooded on this VLAN	Because
XNS — port 1	VLAN 1	The received XNS data does not match any defined VLAN in the system.
IP — port 2	VLAN 2	The received IP data matches IP VLAN 2 for source ports 1 through 10.
IP — port 12	VLAN 1	The received IP data on source port 12 does not match any defined source port for IP VLAN, so the default VLAN is used.

Overlapped IP VLANs

You can assign network layer information to IP VLANs so you can manage your VLANs by subnetwork. The CoreBuilder 2500 system makes flooding decisions by first matching the incoming frame using the protocol (IP) and then matching the frame with Layer 3 subnetwork information. If the received data is IP but does not match any defined IP subnetwork VLAN, the data is flooded within all IP VLANs using the relevant switch port.

For example, you can configure two overlapping IP VLANs for ports 1 through 10 as follows:

- **IP VLAN 1** — Subnetwork 158.101.112.0, ports 1 through 10, with subnet mask 255.255.255.0
- **IP VLAN 2** — Subnetwork 158.101.113.0, ports 1 through 10, with subnet mask 255.255.255.0

The following example shows how flooding decisions are made using overlapping IP VLANs. The example assumes a 12-port switch.

VLAN Index	VLAN Protocol	Network Address/Mask	VLAN Ports
1	Default	none	1 – 12
2	IP	158.103.122.0/ 255.255.255.0	1 – 6
3	IP	158.103.123.0/ 255.255.255.0	6 – 12

Data received on this port	Is flooded on this VLAN	Because
IP subnetwork 158.103.122.2 on port 6	VLAN 2	The IP network layer matches the Layer 3 address for VLAN 2.
IP subnetwork 158.103.123.2 on port 6	VLAN 3	The IP network layer matches the Layer 3 address for VLAN 3.
IP subnetwork 158.103.124.2 on port 6	VLAN 2 and VLAN 3	The IP network layer does not match any Layer 3 address for IP VLANs.
IPX on port 6	VLAN 1	The IPX frame does not match any defined VLAN.

When the subnetwork address of an IP packet does not match any subnetwork address of any defined IP VLAN in the system, the system floods the data to all of the IP VLANs that share the source switch port. In this example, the shared source port is port 6.

Routing Between VLANs

Stations in two different VLANs communicate only by routing between them. The CoreBuilder 2500 system supports internal routing among IP, IPX, and AppleTalk VLANs. If VLANs are configured for other routable network layer protocols, they communicate only through an external router.

You configure routing protocol interfaces based on a VLAN defined for that protocol. To assign a routing interface, you first create a VLAN for that protocol and then associate it with that interface.

For example, to create an IP interface that routes through a VLAN:

- 1 Create an IP VLAN for a group of switch ports.

This IP VLAN does not need to contain Layer 3 information unless you want to further restrict flooding according to the Layer 3 subnetwork address.

- 2 Configure an IP interface with a network address, subnet mask, broadcast address, cost, and type (`vlan`). Select an IP VLAN to bind to that IP interface.



If Layer 3 information is provided in the IP VLAN for which you are configuring an IP interface, the subnetwork portion of both addresses must be the same.

For example:

IP VLAN subnetwork 157.103.54.0 with subnet mask of 255.255.255.0

IP host interface address 157.103.54.254 with subnet mask of 255.255.255.0



The group of ports within an IP VLAN or router interface can communicate at the Layer 2 (bridging) level. IP data uses the IP routing interface to reach a different IP subnetwork, even if the destination subnetwork is on a shared port.

VLAN Configuration Examples

In Figure 2-1, three protocol-sensitive VLANs (two IP and one IPX) interconnect over a high-speed FDDI link. The end stations and servers are on 10 Mbps ports with traffic that is segregated by protocol. Traffic aggregates only over the FDDI link.

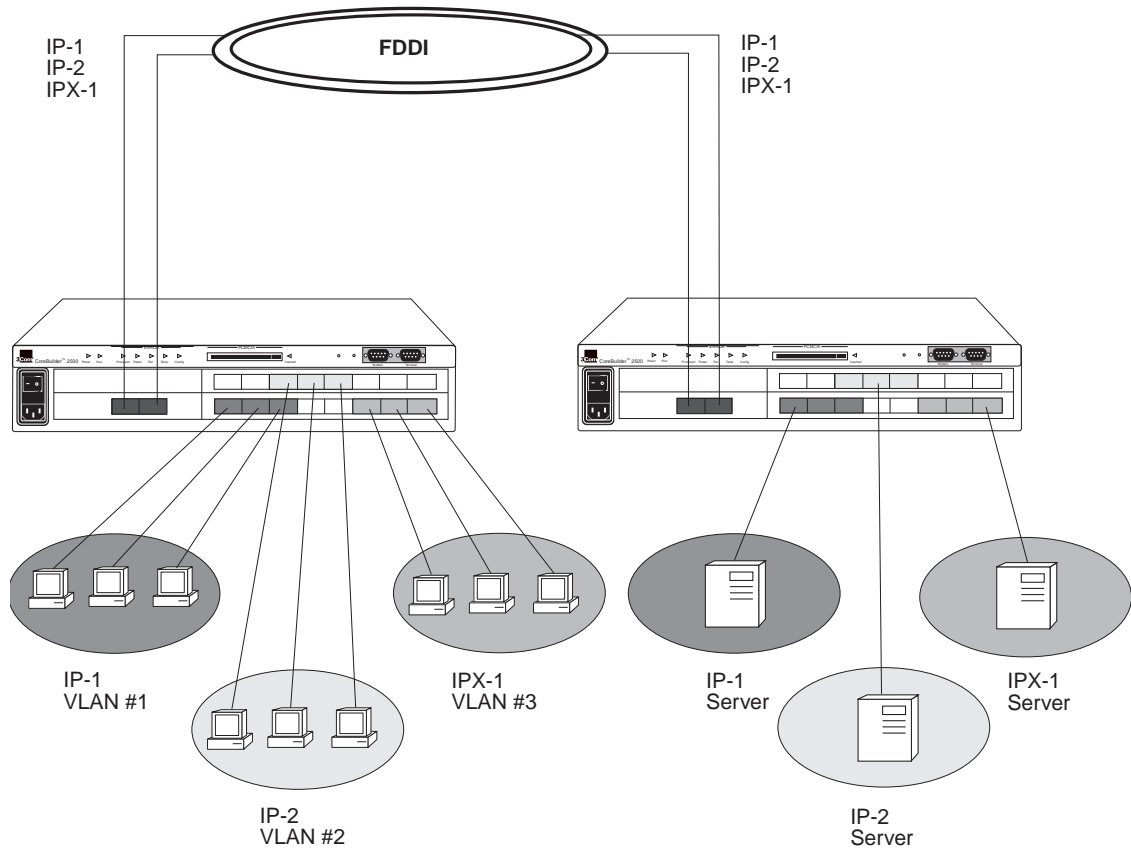


Figure 2-1 Example of a Protocol-Sensitive VLAN Configuration

In Figure 2-2, two overlapping protocol-sensitive VLANs (IP and IPX) are connected to servers on separate, high-speed 100BASE-T ports. The client end stations share the same switch ports, yet the IP traffic and IPX traffic remain separate.




Figure 2-2 Overlapping VLAN Configuration with Servers on Separate Ports



ABOUT ROUTING PROTOCOLS

Chapter 3 Bridging and Routing in the CoreBuilder 2500 System


Chapter 4 Routing with IP Technology

Chapter 5 Routing with IP Multicast

Chapter 6 Routing with IPX

Chapter 7 Routing with OSPF

Chapter 8 Routing with AppleTalk



3

BRIDGING AND ROUTING IN THE COREBUILDER 2500 SYSTEM

This chapter describes how the CoreBuilder™ 2500 system operates in a subnetted routing environment and how the system's bridging and routing model compares with traditional models. The chapter contains these topics:

- What Is Routing?
- Bridging and Routing Models

What Is Routing?

Routing distributes packets over potentially dissimilar networks. A router is the device that accomplishes this task. Routers typically:

- Connect enterprise networks
- Connect subnetworks (or client/server networks) to the main enterprise network

Figure 3-1 shows where routers are typically used in a network. CoreBuilder 2500 system routing connects subnetworks to the enterprise network, providing connectivity between devices within a workgroup, department, or building.

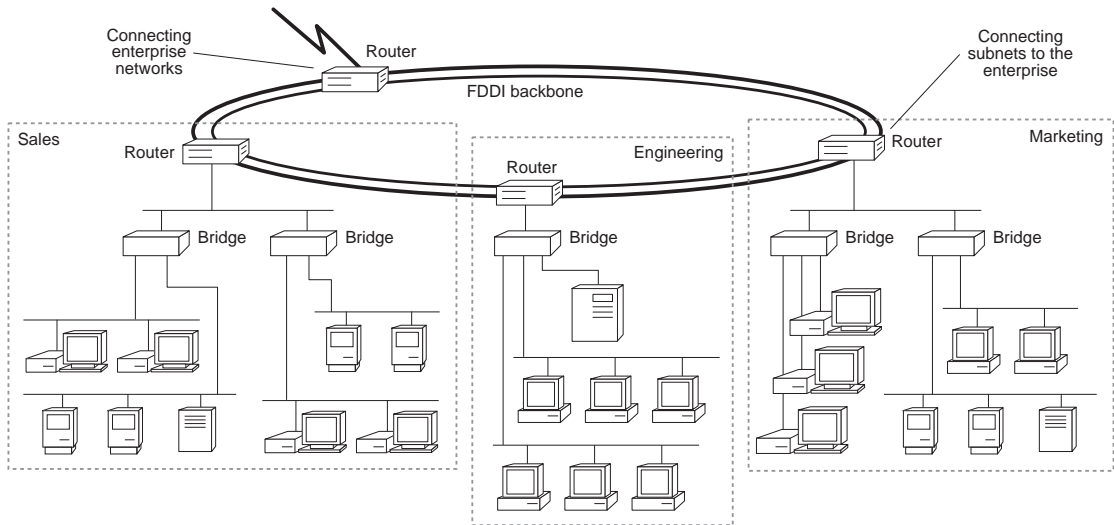


Figure 3-1 Traditional Routing Architecture

CoreBuilder 2500 in a Subnetworked Environment

With the CoreBuilder 2500 system, you fit Ethernet switching capability into subnetworked environments. When you put the CoreBuilder 2500 system into such a network, the system streamlines your network architecture by *routing* traffic between subnetworks and *switching* within subnetworks. See Figure 3-2.




Figure 3-2 CoreBuilder 2500 Subnetwork Routing Architecture

Integrating Bridging and Routing

The CoreBuilder 2500 system integrates bridging and routing. You can assign multiple switch ports to each subnetwork. See Figure 3-3.

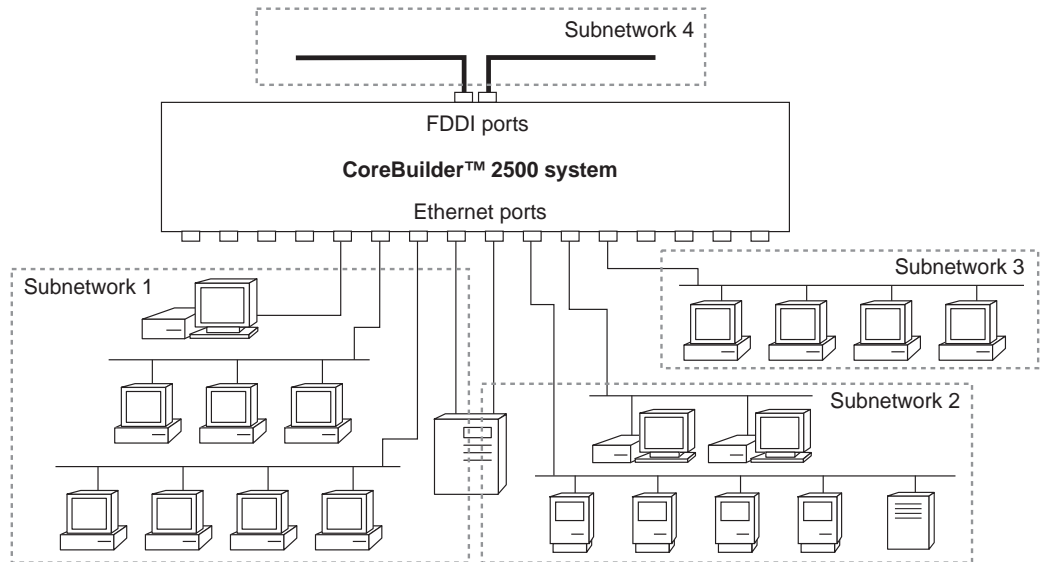


Figure 3-3 Multiple Ethernet Ports per Subnetwork

Transparent bridging or Express switching (described in the *CoreBuilder 2500 Operation Guide*) switches traffic between ports that are assigned to the same subnetwork. Traffic traveling to different subnetworks is routed using one of the supported routing protocols.



*In the following descriptions of bridging and routing on the CoreBuilder 2500 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.*

Because the CoreBuilder 2500 model of bridging and routing allows several segments to be connected to the same subnetwork, you can increase the level of segmentation in your network without creating new subnetworks or assigning new network addresses. Instead, you can use additional Ethernet ports to expand existing subnetworks.

This model differs from traditional bridging and routing, in which at most one port is connected to any subnetwork. Traditionally, to increase the level of segmentation in your network, you must create additional subnetworks and assign new network addresses to existing hosts.

Bridging and Routing Models

The CoreBuilder 2500 system implements routing differently from the way bridging and routing usually coexist. Traditionally, network systems first try to route packets that belong to recognized protocols; all other packets are bridged. In the CoreBuilder 2500 model, the system first tries to bridge a packet. Then, if a packet's destination network address is not on the same subnetwork, the system routes the packet.

The next sections describe these approaches.

Traditional Bridging and Routing

In traditional routing, the bridge or router determines what to do with a packet based on the packet's protocol. If the packet belongs to a recognized protocol, the packet is routed. Otherwise, the packet is bridged. Figure 3-4 illustrates traditional bridging:

- 1 The packet enters the bridge or router.
- 2 The bridge or router determines that the packet does *not* belong to a recognized routing protocol, so the packet is passed to the bridge.
- 3 The bridge examines the destination MAC address and forwards the packet to the port where that address was learned.

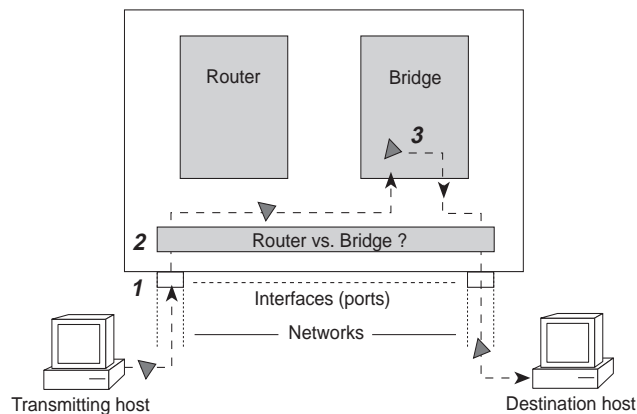


Figure 3-4 Traditional Bridging

Figure 3-5 illustrates traditional routing:

- 1 The packet enters the bridge or router.
- 2 The bridge or router determines that the packet belongs to a recognized routing 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) that is connected to the destination subnetwork.

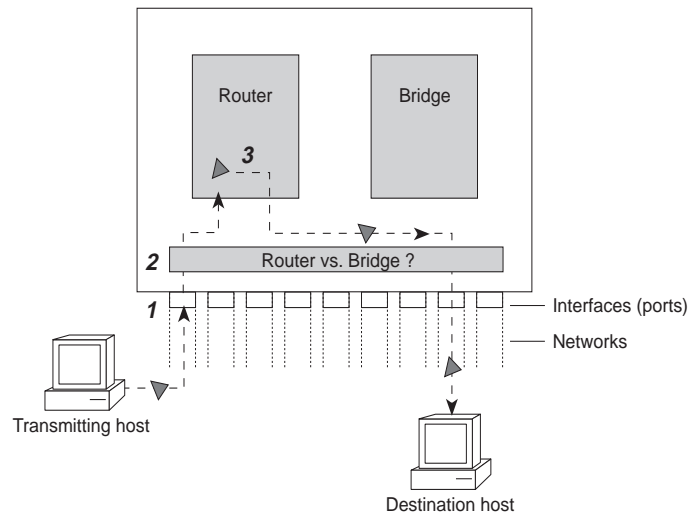


Figure 3-5 Traditional Routing

CoreBuilder 2500 Bridging and Routing

The destination MAC address determines whether the CoreBuilder 2500 system bridges or routes a packet. Before a host system sends a packet to another host, the host system compares its own network address to the network address of the other host as follows:

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

Figure 3-6 illustrates CoreBuilder 2500 bridging:

- 1 The packet enters the CoreBuilder 2500 system.
- 2 The bridging layer examines the packet's destination MAC address. The destination MAC address does *not* correspond to the MAC address of one of the system ports configured for routing.
- 3 The bridging layer selects a segment (port) based on the destination MAC address and forwards the packet to that segment.




Figure 3-6 CoreBuilder 2500 Bridging

Figure 3-7 illustrates CoreBuilder 2500 routing:

- 1 The packet enters the CoreBuilder system.
- 2 The bridging layer examines the packet's destination address. The destination address corresponds to the address of one of the system ports configured for routing (as opposed to a learned end station address).
- 3 The packet is passed to the router interface that is associated with the port where 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.

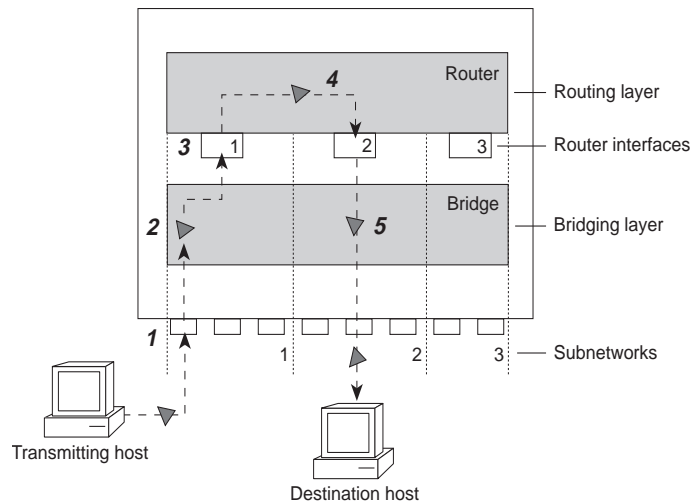


Figure 3-7 CoreBuilder 2500 Routing

4

ROUTING WITH IP TECHNOLOGY

This chapter reviews IP routing technology in these sections:

- IP Routing and the OSI Reference Model
- Elements of IP Routing
- IP Routing Transmission Errors
- Routing with Classical IP over ATM
- IP Routing References

IP Routing and the OSI Reference Model

An IP router, unlike a bridge, operates at the network layer of the Open Systems Interconnection (OSI) Reference Model. An IP router routes packets by examining the network layer address (IP address). Bridges use data link layer MAC addresses to perform forwarding. See Figure 4-1.

OSI Reference Model

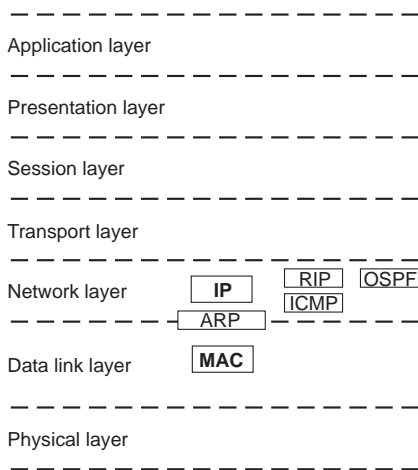


Figure 4-1 OSI Reference Model and IP Routing

When an IP router sends a packet, it does not know the complete path to a destination — only the next hop. Each hop involves three steps:

- 1 The IP routing algorithm computes the *next hop* IP address and the next router interface, using routing table entries.
- 2 The Address Resolution Protocol (ARP) translates the next hop IP address into a physical MAC address.
- 3 The router sends the packet over the network across the next hop.

Elements of IP Routing

IP routers use the following elements to transmit packets:

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

IP Addresses

IP addresses are 32-bit addresses composed of a *network part* (the address of the network where the host is located) and a *host part* (the address of the host on that network). See Figure 4-2.




Figure 4-2 IP Address: Network Part and Host Part

IP addresses differ from Ethernet and FDDI MAC addresses, which are unique hardware-configured 48-bit addresses. A central agency assigns the network part of the IP address, and you assign the host part. All devices connected to the same network share the same network part (also called the *prefix*).

Network Part

The location of the boundary between the network part and the host part depends on the class that the central agency assigns to your network. The three primary classes of IP addresses are A, B, and C:

- **Class A address** — Uses 8 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 address** — Uses 16 bits for the network part and 16 bits for the host part.
- **Class C address** — Uses 24 bits for the network part and 8 bits for the host part. Each Class C network can contain only 254 hosts, but many such networks can be created.

The high-order bits of the network part of the address designate the IP network class.

Subnetwork Part

In some environments, the IP address contains a *subnetwork part*, at the beginning of the host part of the IP address. Thus, you can divide a single Class A, B, or C network internally, allowing the network to appear as a single network to other networks. The subnetwork part of the IP address is visible only to hosts and gateways on the subnetwork.

When an IP address contains a subnetwork part, a *subnet mask* identifies the bits that constitute the subnetwork address and the bits that constitute the host address. A subnet mask is a 32-bit number in the IP address format. The *1* bits in the subnet mask indicate the network and subnetwork part of the address. The *0* bits in the subnet mask indicate the host part of the IP address. See Figure 4-3.

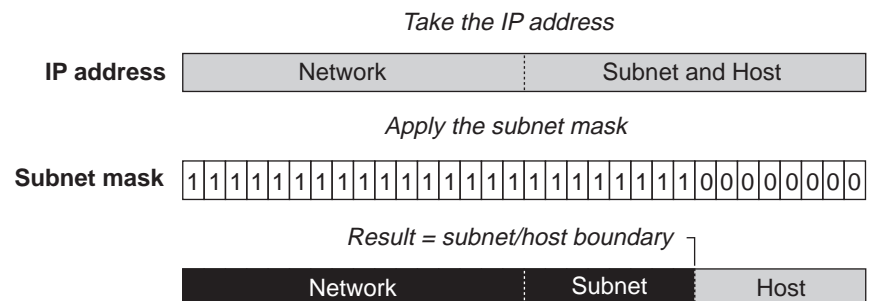


Figure 4-3 How a Subnet Mask Is Applied to the IP Address

An example of an IP address that includes network, subnetwork, 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 subnetwork part
- `52` is the host part



As shown in this example, the 32 bits of an IP address and subnet mask are usually written using an integer shorthand. This notation translates four consecutive 8-bit groups into four integers ranging from 0 through 255. The subnet mask in the example is written as `255.255.255.0`.

Router Interfaces

A router interface connects the router to a subnetwork. In traditional routing models, the interface is the same as the port because only one interface can exist per port. In the IP routing model for the CoreBuilder™ 2500 system, more than one port can connect to the same subnetwork.

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

- For sending IP packets to or from the router.
- For defining the network and subnetwork numbers of the segment connected to that interface. See Figure 4-4.

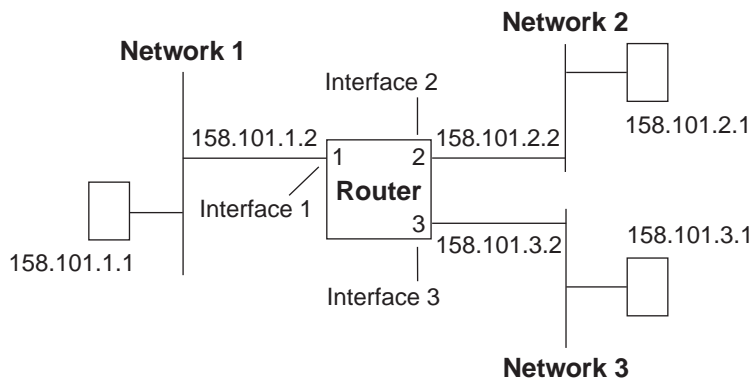


Figure 4-4 Router Interfaces in the CoreBuilder 2500 System

Routing Table

With a routing table, a router or host determines how to send a packet toward its ultimate destination. The routing table contains an entry for every network, subnetwork, and host to which the router or host can forward packets. A router or host uses the routing table when the packet's destination IP address is not on a network or subnetwork 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 elements:

- **Destination IP address** — The destination network, subnetwork, or host.
- **Subnet mask** — The subnet mask for 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 through routers.
- **Gateway** — The IP address of the router interface through which the packet travels on its next hop.
- **Status** — Information that the routing protocol has about the interface.

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

Routing table				
Destination IP address	Subnet mask	Metric	Gateway	Status
158.101.1.1	255.255.255.0	1	158.101.1.2	learned - RIP
158.101.2.1	255.255.255.0	1	158.101.2.2	learned - OSPF - INTRA
158.101.3.1	255.255.255.0	1	158.101.3.2	learned - OSPF - INTRA
default route	255.255.255.0	1	158.101.1.2	learned - OSPF - INTRA

Figure 4-5 Sample CoreBuilder 2500 Routing Table

Routing table data is updated statically or dynamically:

- **Statically** — You manually enter static routes 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 that are generated with a routing protocol. Because static routes do not automatically change in response to network topology changes, manually configure only a small number of reasonably stable routes. Static routes do not time out.
- **Dynamically** — Routers use a protocol such as the Routing Information Protocol (RIP) to automatically exchange routing data. Routes are recalculated at regular intervals. This process helps you to keep up with network changes and allows the system to reconfigure routes quickly and reliably. Interior Gateway Protocols (IGPs), which operate within networks, provide this automated method. The CoreBuilder 2500 system uses RIP and Open Shortest Path First (OSPF) Protocol to configure its routing tables dynamically.

RIP operates using both active and passive devices. *Active* devices, usually routers, broadcast RIP messages to all devices in a network or subnetwork and update their internal routing tables when they receive a RIP message. *Passive* devices, usually hosts, listen for RIP messages and update their internal routing tables, but do not send RIP messages.

An active router sends an RIP message every 30 seconds. This message contains the IP address and a metric (distance) from the router to each destination in the router's internal table. In RIP, each router through which a packet must travel to reach a destination counts as one network *hop*.

OSPF routes packets within and between predefined autonomous systems and areas based on the cost of network links. The OSPF protocol handles network topology changes with a minimum of administrator involvement and routing traffic.

Default Route

In addition to the routes to specific destinations, a routing table can contain a *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 that all have the same gateway IP address and interface number. The default route can be configured statically, or it can be learned dynamically.

Address Resolution Protocol (ARP)

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

When the host or router knows the IP address of the *next* hop toward a packet's destination, the host or router translates that IP address into a MAC address before sending the packet. To perform this translation, the host or router first searches its *ARP cache*, which is 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, the host or router broadcasts an *ARP request* packet to all the devices on the network. The ARP request contains information about the target and source addresses for both the hardware (MAC addresses) and the protocol (IP addresses). See Figure 4-7.

ARP request packet

00802322b00ad	Source hardware address
158.101.2.1	Source protocol address
?	Target hardware address
158.101.2.15	Target protocol address

Figure 4-7 Example of an ARP Request Packet

When devices on the network receive this packet, they examine it. 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 matches the target protocol address, the receiving 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 this *ARP reply*, it places the new MAC address in its ARP cache next to the corresponding IP address. See Figure 4-8.

ARP cache	
IP address	MAC address
158.101.1.1	00308e3d0042
158.101.2.1	0080232b00ab
158.101.3.1	0134650f3000

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

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

IP Routing Transmission Errors

Because a router knows only about the next network hop, it is not aware of problems that may be closer to the destination. Destinations may be unreachable if:

- Hardware is temporarily out of service.
- You specified a nonexistent destination address.
- The routers do not have a route to the destination network.

To help routers and hosts discover problems in packet transmission, a mechanism called Internet Control Message Protocol (ICMP) reports errors back to the source when routing problems occur. With ICMP, you can determine whether a delivery failure resulted from a local or a remote problem.

ICMP performs these tasks:

- **Tests whether nodes can be reached (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 to the sender. This process tests whether the destination is reachable and responding and verifies that the network's transport hardware and software are working. The `ping` command is frequently used to invoke this process.
- **Creates more efficient routing (ICMP Redirect)** — Often the host route configuration specifies the minimum possible routing data that is 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 that a host is not using the best route. The router sends an ICMP redirect to this host, requesting that the host use a different gateway when it sends packets to a destination. The host then sends packets to that destination using the new route.
- **Informs sources that a packet has exceeded its allocated time to exist within the network (ICMP Time Exceeded)**

Routing with Classical IP over ATM

CoreBuilder Extended Switching software supports classical IP routing over ATM ARP in an ATM network. The Classical IP over ATM model uses Logical IP Subnetworks (LISs) to forward packets within the network environment.



See the CoreBuilder 2500 Operation Guide for detailed information about the ATM protocol architecture. See the CoreBuilder 2500 Administration Console User Guide for information about how to configure ATM ports.

About Logical IP Subnets (LISs)

A LIS is a group of IP nodes that belong to the same subnetwork and are directly connected to a single ATM network. When you add a node to a LIS through the Administration Console IP interface menu, you define its IP address, subnet mask, and the address of an ATM ARP server that supports it.

ATM ARP Servers

An ATM ARP server maintains a table of IP addresses and their corresponding ATM addresses and circuit information. To forward IP packets over an ATM interface, the network node learns the ATM address for the corresponding IP address from the ATM ARP server.

Each ATM ARP server supports a single LIS. You can associate two or more LISs with the same ATM network, but each LIS operates independently of other LISs on the network.

Several types of network nodes can function as ATM ARP servers:

- Any CoreBuilder system with revision 8.1.0 or later of Extended Switching software
- An ATM switch
- A UNIX workstation

The following sequence describes how the ATM ARP server learns and stores information about the IP and ATM addresses of nodes in the network:

- 1 A node establishes a connection to the ATM ARP server.
- 2 The ATM ARP server sends an inverse ATM ARP request to the node, requesting its IP and ATM address.
- 3 When the node returns this information, the ATM ARP server stores, or *caches*, it in the ATM ARP server table.

Forwarding to Nodes Within a LIS

Nodes can forward packets directly to other nodes in the same LIS. To forward a packet within the same LIS, the sending node requests a translation from the destination IP address to the corresponding ATM address from the ATM ARP server.

- If the address is known to the server, the server returns a message with this address
- If the address is not known to the server, the server returns a message to advise the sending node that the packet is discarded.

When the server returns a destination address, the sending node uses this learned address to create a *virtual circuit* (VC) and to forward this and all subsequent packets to the destination address. The sending node adds this VC to its ATM ARP cache.

**IP Routing
References**

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

Perlman, Radia. *Interconnections: Bridges and Routers*. 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*.

RFC 1577. *Classical IP over ATM*.

RFC 1583. *OSPF Version 2*

5

ROUTING WITH IP MULTICAST

This chapter describes the IP multicast routing implementation on the CoreBuilder™ 2500 system. It includes these sections:

- About IP Multicast Routing
- IGMP
- DVMRP
- Multicast Routing Algorithms
- Multicast Interfaces
- Multicast Tunnels

About IP Multicast Routing

IP multicast routing is an extension of the Internet Protocol. With multicast routing, a router or switch sends packets to a specific group of hosts without using broadcasts or multiple unicast transmissions. Multicast destinations include:

- Hosts that reside on the local LAN
- Hosts that reside on different sites within a private network
- Hosts that are scattered throughout the Internet

Multicast routing operates without loops or excess transmissions.

The CoreBuilder 2500 system supports two IP multicast protocols:

- Internet Group Management Protocol (IGMP)
- Distance Vector Multicast Routing Protocol (DVMRP)

This chapter describes these protocols and the algorithms that the CoreBuilder 2500 system uses for multicast routing.

IGMP

The CoreBuilder 2500 system is capable of dynamic multicast filtering based on the Internet Group Management Protocol (IGMP). This protocol ensures that multicast packets are flooded only to appropriate ports in a routing interface.

IGMP tracks end station group membership within a multicast group. Membership in a group is dynamic, and hosts can be a member of more than one group at a time. The system avoids propagating multicast broadcasts to the entire subnetwork by confining them within the multicast group (a process called IGMP *snooping*).

DVMRP

The Distance Vector Multicast Routing Protocol (DVMRP) establishes multicast delivery paths over a series of routing devices. DVMRP is a distance-vector-routing protocol, similar to the IP Routing Information Protocol (RIP). Multicast routers exchange distance vector updates that contain lists of destinations and the distance in router hops to each destination. They maintain this information in a routing table.

The Internet Multicast Backbone (MBONE) uses DVMRP. Because of DVMRP, the CoreBuilder 2500 system can establish delivery paths without direct connections to multicast routers.

MBONE

The *MBONE* is the Internet's experimental multicast backbone network. Users can test multicast applications and technology on the MBONE without waiting for Internet multicast standards to be set. You can connect to the MBONE through any Internet service provider (ISP).

MBONE routers forward multicast packets over an interface or over a multicast tunnel only if the Time-To-Live (TTL) value in the packet is larger than the tunnel's threshold. See "Multicast Tunnels" on page 5-5 for more information about tunnels.



At software revisions earlier than 8.0, CoreBuilder 2500 systems that are connected to the MBONE network support up to 16 IP multicast tunnels or routing interfaces. CoreBuilder 2500 systems at revision 8.0 or later support up to 32 IP multicast tunnels or routing interfaces connected to the MBONE.

Multicast Routing Algorithms

The CoreBuilder 2500 system uses three multicast routing algorithms:

- Flooding
- Spanning Tree
- Reverse Path Forwarding

Flooding In most flooding algorithms, a network node receives a packet that was sent to a multicast destination. The node determines whether the packet is an original that it has not received before or a duplicate of a packet that it has received before. If the packet is original, the node forwards the packet on all interfaces except the incoming interface. If the packet is a duplicate, the node discards it.

This flooding algorithm is useful when network robustness is important. The algorithm does not depend on routing tables. Multicast destinations receive packets as long as at least one path to the destinations exists and no errors occur during transmission.

Spanning Tree The Spanning Tree algorithm detects loops and logically blocks redundant paths within the network. The paths form a loopless graph, or *tree*, spanning all the nodes in the network. A port in the Spanning Tree *blocking* state does not forward or receive data packets.

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

Figure 5-1 shows a simple network with six links.




Figure 5-1 Simple Network Implemented Without Using Spanning Tree

Figure 5-2 shows a spanning tree for the network in Figure 5-1. The spanning tree consists of links 1, 2, 3, and 4.




Figure 5-2 Spanning Tree Algorithm Implemented to Block Redundant Paths

Reverse Path Forwarding

The MBONE network uses the reverse path forwarding (RPF) multicast algorithm. RPF avoids duplicate paths on multiaccess links by using a routing table to compute a logical spanning tree for each network source. The algorithm consists of three steps:

- 1 When the system receives a multicast packet, the algorithm notes the packet's source network and the interface that received the packet.
- 2 If the receiving interface is on the shortest path toward the source network, the system forwards the packet to all interfaces except the interface where the packet was received.
- 3 If the receiving interface is not on the shortest path toward the source network, the system drops the packet.

Pruning

Pruning is a method used in the RPF algorithm to forward packets to a spanning tree only if group members exist in the tree. This method results in fewer spanning trees, but it requires dynamic updates to the routing table.

Nodes that are at the border of the network and have no point beyond them in the RPF spanning tree are called *leaf* nodes. Leaf nodes all receive the first multicast packet. If a group member is attached to a leaf node, the node continues to accept packets. If no group member is attached to the leaf node, the node sends back a *prune message* to the router that sent the packet. The message notifies the router not to send any further packets to this group. In the CoreBuilder 2500 system, the Administration Console IP multicast `cacheDisplay` includes information about when pruning occurs on the Spanning Tree.

Multicast Interfaces	This section describes the characteristics of the CoreBuilder 2500 multicast interface.
DVMRP Metric Value	The DVMRP metric value determines the <i>cost</i> of a multicast interface. A higher cost results in a slower link. The default value is <i>1</i> .
Time-To-Live (TTL) Threshold	The TTL threshold determines whether the interface forwards multicast packets to other switches and routers in the subnetwork. If the interface TTL is greater than the packet TTL, then the interface does not forward the packet. The default value of <i>1</i> means that the interface forwards most packets.
Rate Limit	The rate limit determines how many multicast packets can travel over the interface per second. The CoreBuilder 2500 system drops multicast traffic that travels faster than this rate. The default value of <i>0</i> means that no rate limit is in effect. A lower rate limit results in fewer multicast packets traveling over the interface.

Multicast Tunnels Multicast *tunnels* logically connect two multicast routers through one or more unicast routers. The multicast router at the local endpoint of the tunnel puts multicast packets in a format that unicast routers can interpret and forward. The multicast router at the receiving endpoint reformats the packets into their multicast format. Tunnels are virtual links through the unicast IP network.

Multicast tunnels have characteristics similar to those of a multicast interface: a DVMRP metric value, a TTL threshold, and a rate limit. When you define a multicast tunnel, you also specify the destination address of the multicast router at the tunnel's remote endpoint.

6

ROUTING WITH IPX

This chapter provides an overview of IPX routing, including:

- IPX Routing in the NetWare Environment
- How IPX Routing Works
- Elements of IPX Routing

IPX Routing in the NetWare Environment

The NetWare network operating system was developed and introduced to the market by Novell, Inc., in the early 1980s. Much of the NetWare networking technology was derived from XNS, a networking system developed by Xerox Corporation.

The NetWare operating system is based on a client/server architecture in which clients request services, such as file and printer sharing, from servers. The NetWare operating system uses the upper five layers of the Open Systems Interconnection (OSI) Reference Model.

The CoreBuilder™ 2500 system uses the following protocols for routing in a NetWare environment:

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

Figure 6-1 illustrates a simplified view of the better-known protocols of NetWare and their relationship to the OSI Reference Model.

Layers in the
OSI Reference Model

NetWare

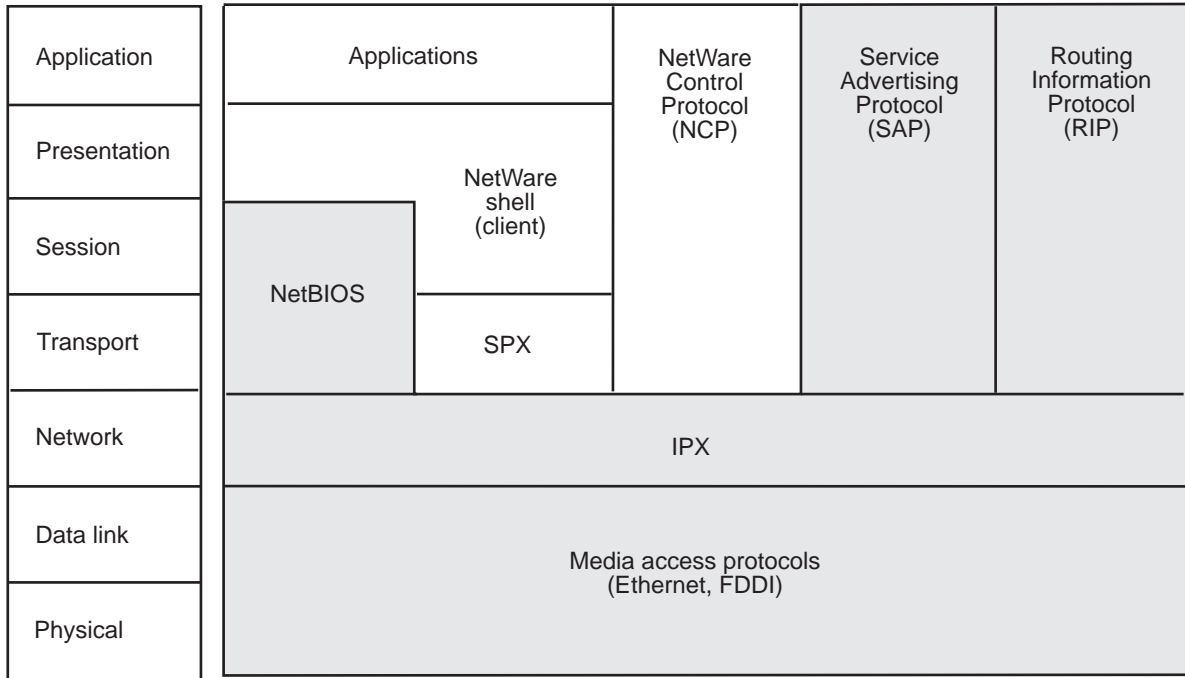


Figure 6-1 NetWare Protocols and the OSI Reference Model

Internet Packet Exchange (IPX)

IPX is the primary protocol used for routing in a NetWare environment. This connectionless, datagram protocol does not require an acknowledgment for each packet sent. Protocols above IPX provide packet acknowledgment or connection control.

IPX defines internetwork and intranode addressing schemes. Internetwork addressing is based on network numbers that are assigned to each interface in an IPX network. Intranode addressing is in the form of socket numbers. Because 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 create and maintain their dynamic routing tables.

With RIP, one router exchanges routing information with a neighboring router. When a router discovers any changes in the network layout, it broadcasts this information to any neighboring routers. IPX routers also send periodic RIP broadcast packets containing all of the routing information that the router possesses. These broadcasts synchronize all routers on the network and age those networks that might become inaccessible if a router disconnects abnormally from the network.

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. With this process, routers dynamically create and maintain a database (called a server table) of network service information. Clients on the network 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.

With SAP, a router exchanges information with neighboring SAP agents. When a router's SAP agent discovers a change in the network server layout, it immediately broadcasts this information to neighboring SAP agents. The router also periodically sends SAP broadcast packets that contain all server information that the SAP agent possesses. These broadcasts synchronize all servers on the network and age out any routers or servers that become inaccessible due to abnormal shutdown.

How IPX Routing Works

A router operates at the network layer of the OSI Reference Model. The router receives instructions to route packets from one segment to another from the network-layer protocol. IPX, with the help of RIP, performs network-layer tasks, including:

- Addressing packets
- Routing packets
- Switching packets

This section describes the “IPX Packet Format” and the process of “IPX Packet Delivery.”

IPX Packet Format

An IPX packet consists of a 30-byte header followed by packet data. The network, node, and socket addresses for both the destination and source are in the packet header.

Figure 6-2 shows the IPX packet format.

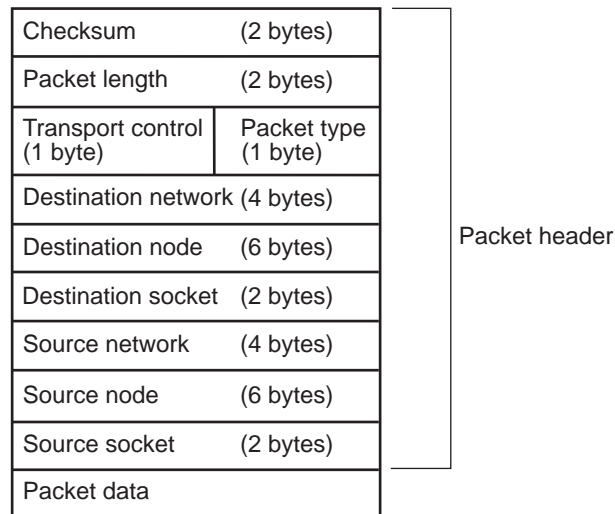


Figure 6-2 IPX Packet Format

The packet includes the following elements:

- **Checksum** — A 16-bit checksum that is set to 1s.
- **Packet length** — A 16-bit field that indicates the packet's length in bytes. This length includes both header and data. The length must be at least 30 bytes.
- **Transport control** — A 1-byte field that indicates how many routers a packet has passed through on its way to its destination. Packets are discarded when this value reaches 16. A network node sets this field to 0 before the node sends the IPX packet.
- **Packet type** — A 1-byte field that specifies the upper-layer protocol that receives the packet.
- **Destination network** — A 4-byte field that contains the destination node network number. When a sending node sets this field to 0, the system routes the packet as if the sending and destination nodes are on the same local segment.
- **Destination node** — A 6-byte field that contains the destination node physical address.
- **Destination socket** — A 2-byte field that contains the socket address of the packet's destination process.
- **Source network** — A 4-byte field that contains the source node network number. If a sending node sets this field to 0, the source's local network number is unknown.
- **Source node** — A 6-byte field that contains the source node physical address. Broadcast addresses are not allowed.
- **Source socket** — A 2-byte field that contains the socket address of the process that transmitted the packet.
- **Packet data** — A field that contains information for upper-layer network processes.

IPX Packet Delivery

Successful packet delivery depends on proper addressing and the network configuration. The packet's Media Access Control (MAC) protocol header and IPX header address handle packet addressing.

The sending node must have the destination's complete network address, including the destination network, node, and socket. After the sending node has the destination address, it can address the packet.

However, the way that the MAC header of the packet is addressed depends on whether a router separates the sending and destination nodes. See Figure 6-3.

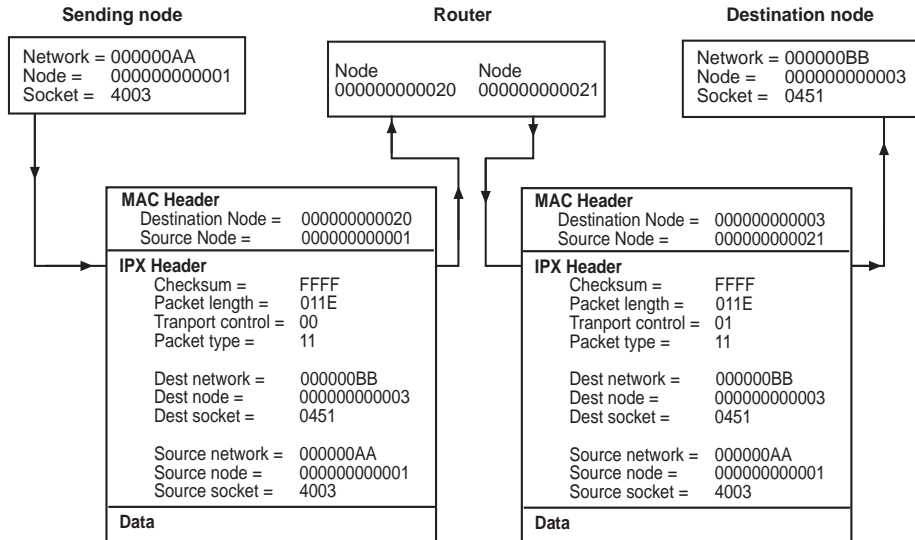


Figure 6-3 IPX Packet Routing

Sending Node's Responsibility

When sending and destination nodes have the same network number, the sending node addresses and sends packets directly to the destination node. If sending and destination nodes have different network numbers, as in Figure 6-3, the sending node must find a router on its own network segment that can forward packets to the destination node's network segment.

To find this router, the sending node broadcasts an RIP packet requesting the best route to the destination node's network number. The router that resides on the sending node's 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.

After the sending node determines an intermediate router's address, it can send packets to the destination node.



If the sending node is a router rather than a workstation, the node's internal routing tables supply the destination's network location. The destination router does not need to broadcast an RIP request.

Router's Responsibility

A router handles a received IPX packet in one of two ways:

- If the packet is destined for a network number to which the router is directly connected, the sending router:
 - Places the destination node address from the IPX header in the destination address field of the packet's MAC header
 - Places its own node address in the source address field of the packet's MAC header
 - Increments the transport control field of the IPX header and transmits the packet on the destination node segment
- If the packet is destined for a network number to which the router is not directly connected, the router sends the packet to the next router along the path to the destination node. The sending router:
 - Looks up the node address (in the routing information table) of the next router and places the address in the destination address field of the packet's MAC header. For more information about routing tables, see "Elements of IPX Routing," next.
 - Places its own node address in the source address field of the packet's MAC header
 - Increments the transport control field in the IPX header and sends the packet to the next router

Elements of IPX Routing

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

- Router interfaces
- Routing tables
- Service Advertising Protocol (SAP)

Router Interfaces

A router interface connects the router and the network number (address). In traditional routing models, the router interface is the same as the port, because only one interface exists per port. But in the CoreBuilder 2500 system, more than one port can connect to a network number. Therefore, the router interface is a relationship between one or more ports and the network number (address) in your IPX network.

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

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

Routing Tables

A routing table collects information about all intranetwork segments. This table allows a router to send packets toward their destinations over the best possible routes.

The table contains an entry for every network number that the router knows about. The router uses this information when the router is not directly connected to a packet's destination network. The routing information table provides the address of another router that *can* forward the packet toward its destination.

The routing table consists of the following elements:

- **Interface** — The interface number of the router that is used to reach a network segment
- **Addresses** — The network segments that the router knows about
- **Hops to network** — The number of routers that must be crossed to reach a network segment
- **Ticks to network** — An estimate of the time in seconds necessary to reach a network segment
- **Node** — The node address of the router that can forward packets to each network segment. (When this element is set to all 0s, the router is directly connected)
- **Aging timer** — The time in seconds since the network's last update

Figure 6-4 shows an example of a typical 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 6-4 Sample Routing Table

Generating Routing Table Information

The routing information table is updated statically or dynamically.

Static Update You manually configure a static route. Static routes are useful in environments where no routing protocol is used or where you want to override a routing protocol's generated route.

Because static routes do not automatically change in response to network topology changes, manually configure only a small number of reasonably stable routes. Static routes do not change until you change them, and they do not time out.

Dynamic Update A router uses RIP to exchange its routing table with other routers at regular intervals. This automatic method of learning routes helps you to keep up with a changing network environment and allows you to reconfigure routes quickly and reliably. Interior Gateway Protocols (IGPs), which operate within intranetworks, provide this automated learning. The CoreBuilder 2500 system uses RIP (one of the most widely used IGPs) to dynamically build routing tables.

RIP operates with active and passive network devices. *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. *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 the network number for each destination network and the number of hops to reach it. In RIP, each router through which a packet must travel to reach a destination counts as one network *hop*.

Selecting the Best Route

Large networks contain many possible routes to each destination. A router performs the following steps to find the best route toward a destination:

- If one route requires the lowest number of ticks, the router selects it as the best route.
- If multiple routes require the same lowest number of ticks, the router selects the route that requires the lowest number of hops as the best route.
- If multiple routes require the same lowest number of ticks and hops, the router may select any of them as the best route.

Service Advertising Protocol in IPX

With the Service Advertising Protocol (SAP), file, print, application, and gateway servers broadcast their addresses and services throughout the intranetwork. Services are added and removed dynamically. A server advertises itself at startup; and if a server is shut down, it broadcasts that its services are no longer available.

Internetwork Service Information

Using SAP, IPX routers maintain a database of internetwork service information. Client workstations use this data to determine what services are available on the network and to obtain the internetwork address of the nodes (servers) where they can access services.



Before it can initiate a session with a file server, a workstation must know a server's network address.

SAP Packet Structure

SAP packets support the following functions:

- Workstation requests for the name and address of the nearest server of a certain type
- Router requests for the names and addresses of all the servers, or of all servers of a certain type, on the intranetwork
- Responses to workstation or router requests
- Periodic server and router SAP broadcasts
- Changed server information broadcasts

Figure 6-5 illustrates the SAP packet structure. Note that the SAP packet is encapsulated within the IPX packet data area.

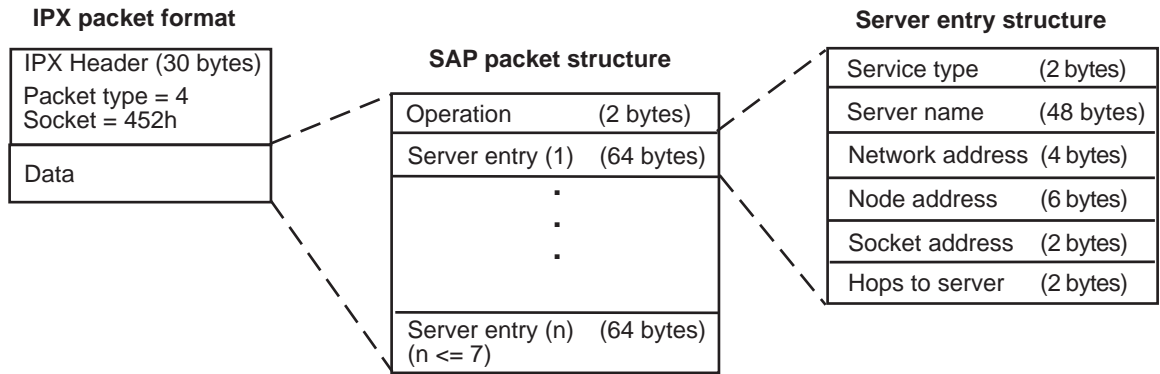


Figure 6-5 SAP Packet Structure

A SAP packet consists of the following fields:

- **Operation** — A 2-byte field that indicates the type of operation the SAP packet performs. You can set this field to one of the following values:
 - 1=Request
 - 2=Response
 - 3=Get nearest server request
 - 4=Get nearest server response
- **Server entry** — A 64-byte field of server information, which consists of the following subfields:
 - **Service type** — A 2-byte field that identifies the type of service the server provides



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

- **Server name** — A 48-byte name that is assigned to the server. The server name, in combination with the service type, uniquely identifies a server on an intranetwork.
- **Network address** — A 4-byte field that contains the server's network address
- **Node address** — A 6-byte field that contains the server's node address

- **Socket address** — A 2-byte field that contains the socket number that the server uses to receive service requests
- **Hops to server** — A 2-byte field that indicates the number of intermediate networks that must be passed through to reach this server. Each time a packet passes through an intermediate network, this field is incremented by 1.

Client systems do not use this SAP information directly. Rather, SAP agents within each router on the server's network segment collect this information. The SAP agents store this information in their server information tables. If a server also contains a SAP agent, the server's bindery stores the SAP information. Client systems then contact the nearest router or file server SAP agent to obtain server and service information.

Server and router SAP broadcasts reach SAP agents only on directly connected network segments. However, SAP agents periodically broadcast their server information tables, so that all SAP agents on the network know about all active servers. These server information tables are described in the next section.

Server Information Tables

Server information tables contain data about all active servers on the intranetwork. SAP agents use these tables to store information received in SAP broadcasts.

Figure 6-6 shows a sample server information table.

Server information 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 6-6 Server Information Table

This table contains the following data:

- **Interface** — The interface from which server information is received
- **Server name** — The name of the server
- **Server type** — The type of service the server provides
- **Network address** — The address of the network that contains the server
- **Node address** — The server's node address
- **Socket address** — The socket number through which the server receives service requests
- **Hops to server** — The number of intermediate networks that must be crossed to reach the server
- **Age of server** — The time in seconds since the server's last table update

The server information table is updated statically or dynamically.

Static Update You manually update 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/server protocol. Because static servers do not automatically change in response to network topology changes, manually configure only a small number of relatively stable servers.

Dynamic Update Servers are automatically added to and removed from the information table. This automatic SAP update helps you to keep up with changing network environments and allows servers to advertise their services and addresses quickly and reliably.

Server Information Maintenance

When a router's SAP agent receives a SAP broadcast response indicating a change in a server's configuration, the agent updates its server information table and informs other SAP agents. Examples of such a change are when a server is disconnected or becomes accessible through a better route.

The SAP agent immediately sends an update broadcast to all directly connected network segments except the segment from which the information was received. All future periodic broadcasts contain the change information.

SAP Aging Router SAP agents use a special aging mechanism to deal with a SAP agent that goes down suddenly without sending a DOWN broadcast. A hardware failure, power interruption, or power surge can cause this situation.

Each SAP agent maintains a timer for each entry in its server information tables. This timer tracks the elapsed time since the entry has been updated. This information is either new or changed, and the SAP agent immediately passes it on, so changes are quickly captured and stored throughout the intranetwork.

SAP Request Handling When a SAP agent receives a general request, it notifies the sending source about all servers known to the agent. This response includes the same information that is sent out in periodic SAP broadcasts. When the request is specific, the SAP agent notifies the sending source about all servers of the requested type.

7

ROUTING WITH OSPF

This chapter describes Open Shortest Path First (OSPF) routing, including:

- Elements of OSPF Routing
- How OSPF Routing Works

The OSPF link-state protocol dynamically responds to changes in network topology that occur within a group of networks and routers called an *autonomous system*. OSPF tracks the states of links and routers in each autonomous system, and when a change occurs, calculates new routes based on the new topology. The OSPF protocol responds to network topology changes with a minimum of administrator involvement and routing traffic.

All OSPF routers within an autonomous system build and synchronize databases of the autonomous system's network topology. Using its database, each router calculates the shortest path trees to every destination within the autonomous system. With this dynamic table of shortest paths, OSPF converges on an optimum route faster than other routing algorithms, such as the Routing Information Protocol (RIP).



Routers that use a distance-vector protocol like RIP periodically exchange all or a portion of their tables, but only to their neighbors. Routers using a link-state protocol like OSPF send small portions of their tables throughout the network by flooding.

Figure 7-1 illustrates a typical OSPF application.

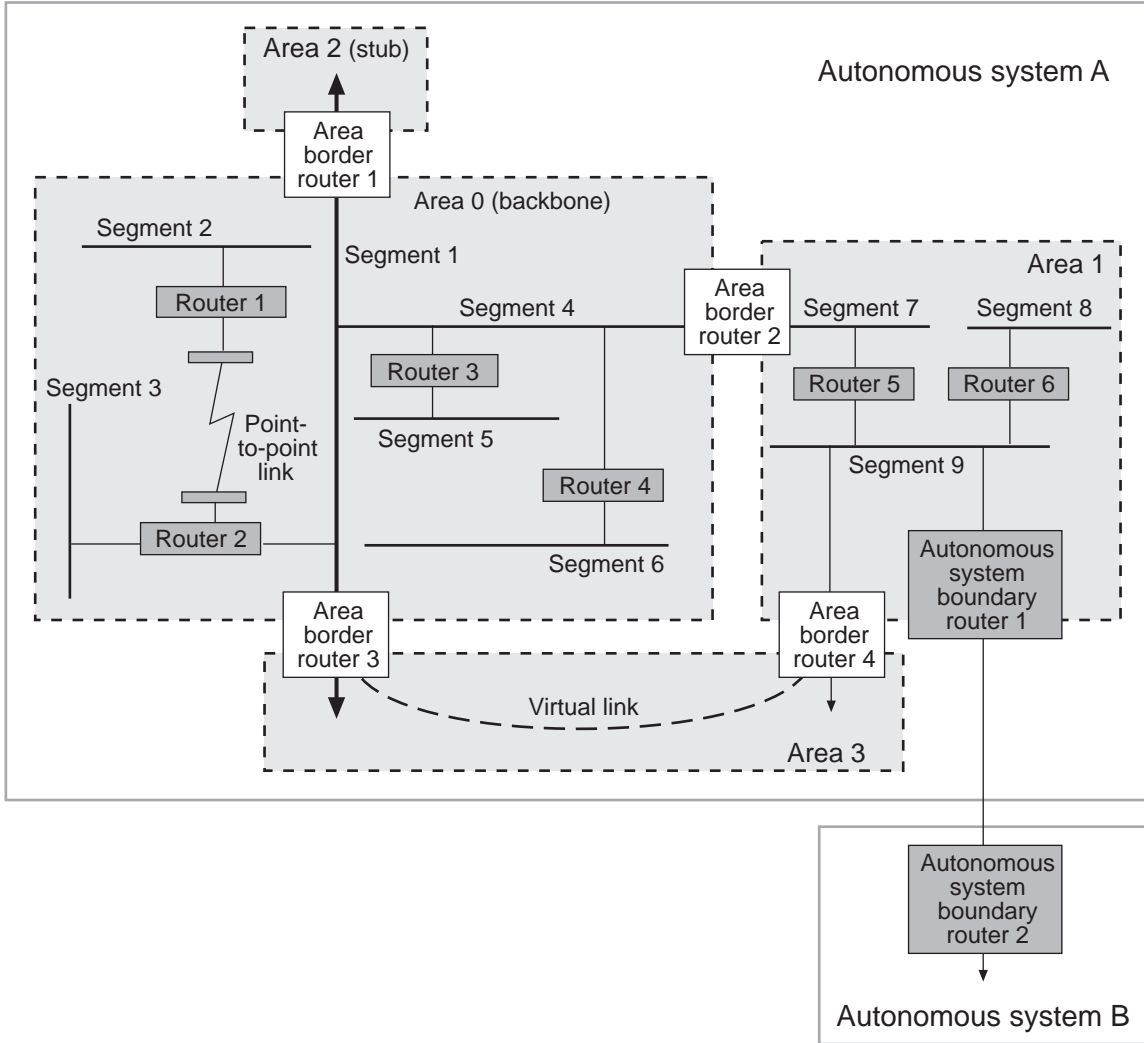


Figure 7-1 Sample OSPF Routing Application

Elements of OSPF Routing

OSPF routing uses the following network elements:

- Autonomous Systems
- Areas
- Neighbors
- Protocol Packets
- Router Types
- Interface Characteristics
- Stub Default Metrics
- Virtual Links

Autonomous Systems

An *autonomous system* consists of a set of OSPF routers that exchange routing information. The network shown in Figure 7-1 contains two autonomous systems.

Using identical topology databases, each router in an autonomous system calculates shortest-path routes from itself to every known destination in the autonomous system. The routers create their topology databases using the data in link state advertisements (LSAs) from other routers in the autonomous system.

Areas

You can reduce the amount of routing information that travels through a network, and the corresponding size of OSPF routers' topology databases, by subdividing OSPF autonomous systems into *areas*. The routers in an area maintain and use identical LSA databases.

The network shown in Figure 7-1 contains four OSPF areas within autonomous system A. There are three types of OSPF areas:

- **Transit** — An area through which network traffic can pass to reach other areas in the autonomous system. In Figure 7-1, the backbone area and areas 1 and 3 are transit areas.
- **Stub** — An area with only one entry or exit router. In Figure 7-1, area 2 is a stub area that is reached only through area border router 1.

- **Backbone** — A contiguous area within an autonomous system that is divided into more than one area. The system automatically defines the backbone area as 0.0.0.0. The backbone area distributes routing data between other areas in the autonomous system. By definition, the backbone area is also a transit area.



An area's network topology is not visible outside the area. Conversely, an area's systems cannot see detailed network structures outside the area. Because of this restriction of topological information, you can control traffic flow between areas and reduce routing traffic to below the levels that occur when the entire autonomous system is a single routing domain.

Area Border Routers

Each area (including the backbone area) includes all border routers connected to the area. In Figure 7-1, for example, you define:

- Area border routers 1, 2, and 3 as being in the backbone area 0
- Area border routers 2 and 4 as being in area 1
- Area border router 1 as being in area 2
- Area border routers 3 and 4 as being in area 3

Routers must communicate with each other through interfaces that are defined as being in the same area. An area, however, may contain virtual links from area border routers to the backbone area. For example, in Figure 7-1, area border routers 3 and 4 terminate a virtual link between area 1 and the backbone area. See “Virtual Links” later in this chapter for more details.



CAUTION: *Do not disconnect an area border router from the backbone area. This action may result in a loss of network topology information and routing capability. You must connect all area border routers to the backbone area using either physical or virtual links.*

Routing Databases

All routers connected to an area maintain identical routing databases about the area. Routers connected to multiple areas maintain a separate routing database for each attached area. For example, in Figure 7-1:

- Routers 1, 2, 3, and 4 maintain identical routing databases about backbone area 0.
- Routers 5 and 6 maintain identical routing databases about area 1.
- Area border router 1 maintains separate routing databases about backbone area 0 and area 2.
- Area border router 2 maintains separate routing databases about backbone area 0 and area 1.
- Area border router 3 maintains separate routing databases about backbone area 0 and area 3.
- Area border router 4 maintains separate routing databases about areas 1 and 3.
- Autonomous system boundary routers 1 and 2 maintain separate routing databases about autonomous systems A and B.

In the CoreBuilder™ 2500 implementation of the OSPF model, each area has the following configurable parameters:

- **Area ID** — The number, in the form of an IP address, that identifies the area to the OSPF autonomous system



A backbone area must have the area ID 0.0.0.0. Routers in the backbone area also must be able to communicate with each other through interfaces that are configured as in area 0.

- **Stub area** — An OSPF area that does not accept or distribute external address advertisements. Use the stub area designation to minimize topological data stored in the area's routers.
- **Range** — An address that covers a range of subnetwork addresses. A range address aggregates LSAs from all of its subnetwork addresses.
- **Default route metric** — The network cost for an OSPF default route. If the cost is greater than 0, the router advertises itself as the default router to the area.

Default The system default value for the default route metric is 0, which means that the router does not advertise itself as the area's default router.

Neighbors Neighbor routers are physically attached to the same network segment. A router attached to multiple network segments may have different sets of neighbors on each segment. For example, Figure 7-1 includes several sets of OSPF neighbor routers.

In backbone area 0:

- Router 2 and area border routers 1 and 3 are neighbors on segment 1 (the backbone network).
- Routers 1 and 2 are neighbors on a point-to-point link.
- Routers 3 and 4 and area border router 2 are neighbors on segment 4.
- No routers are neighbors on segments 2, 3, 5, and 6.

In area 1:

- Router 5 and area border router 2 are neighbors on segment 7.
- Routers 5 and 6, area border router 4, and autonomous system boundary router 1 are neighbors on segment 9.
- No routers are neighbors on segment 8.

In area 3:

- Area border routers 3 and 4 are neighbors on a virtual link between the backbone area 0 and area 1.



Routers use OSPF hello packets to learn neighbor addresses dynamically on interfaces that support multicast routing. Define static neighbors only on nonmulticast interfaces.

Protocol Packets

The OSPF protocol uses five types of packets:

- **Hello** — Router interfaces periodically transmit hello packets to identify and maintain communications with their neighbors.



In nonmulticast networks, routers find neighbors by sending unicast hello packets to other statically configured routers.

- **Database description** — Neighbor routers use database description packets to synchronize their link state summary databases.
- **Link state request** — To collect network topology data, routers transmit link state request packets to their neighbors on the segment.
- **Link state update** — On receiving a link state request packet, a router floods packets containing its LSA data into the area or autonomous system that it serves. The information contained in the packets depends on the router's location and function in the network.
- **Link state ack(nowledge)** — Routers use these packets to acknowledge receipt of link state update packets.



The router acknowledges receiving each LSA in the link state update packet.

Router Types

OSPF routers serve several different, often overlapping, functions:

- **Internal routers** — Internal routers connect only to networks that belong to the same area. An internal router runs one copy of the OSPF algorithm and maintains routing data only for its area.

In Figure 7-1, backbone area 0 and routers 1, 2, 3, and 4 are internal routers. In area 1, routers 5 and 6 are internal routers.

- **Backbone routers** — Backbone routers connect to the backbone network and are configured as belonging to the OSPF backbone area. Area border routers are always backbone routers, because you must configure them as being within the backbone area or connected to it by a virtual link.

In Figure 7-1, router 2 and area border routers 1, 2, 3, and 4 are backbone routers.

- **Area border routers** — Area border routers connect directly to networks in two or more areas. An area border router runs a separate copy of the OSPF algorithm and maintains separate routing data for each area connected to it (including the backbone area).

In Figure 7-1, four area border routers link the areas in autonomous system A.



Area border routers send configuration summaries for their attached areas to the backbone area, which distributes this information to other OSPF areas in the autonomous system.

- **Autonomous system boundary routers** — Autonomous system boundary routers exchange their autonomous system topology data with boundary routers in other autonomous systems. Every router inside an autonomous system knows how to reach the boundary routers for its autonomous system.

In Figure 7-1, two autonomous system boundary routers control traffic between two autonomous systems.

- **Designated routers (DRs)** — Designated routers advertise network link states for attached network segments. A link state advertisement lists all routers connected to a segment.



The DR exchanges routing data with all routers that are connected to its network segment.

- **Backup designated routers (BDRs)** — Backup designated routers are given a lower priority value than the DR and take over DR functions if the DR fails.

Router IDs

The OSPF router ID identifies a router to other routers within an autonomous system. OSPF uses three types of router identifiers, which take the form of an IP address:

- **Default** — An arbitrary ID that the system generates and uses as the default router ID
- **Interface** — The address of an IP interface on the router
- **Address** — An arbitrary user-defined ID in the form of an IP address

Interface Characteristics

You configure OSPF router interfaces by adding OSPF characteristics to existing IP Virtual LAN (VLAN) interfaces. See Chapter 9, “Administering VLANs,” for information about how to configure VLAN interfaces.

The OSPF interface has the following characteristics and statistics, which are discussed in the next sections:

- Mode
- Priority
- Area ID
- Cost
- Delay
- Hello Timer
- Retransmit Timer
- Dead Interval
- Password

Mode

The mode for an interface can be *off* or *active*. To run OSPF routing on the interface, set the mode to *active*.



To set the OSPF interface mode to active, you must enable IP routing. See Chapter 10 for information about enabling IP routing.

Default The default mode is *off*.

Priority

You assign the interface priority to an OSPF router to determine its status as a designated router. A router can function in one of three ways:

- **Designated router (DR)** — The router with the highest priority value, unless a designated router already exists on the network segment.
- **Backup designated router (BDR)** — The router with a lower priority than the DR; the BDR takes over DR functions if the DR fails.
- **Not a designated router** — Any router given a priority 0 or not elected DR or BDR. Priority 0 routers can never be elected as a DR or a BDR.

Default The default priority value is 1.

Area ID

The area ID associates the router interface with an OSPF area. See Chapter 13, “Administering OSPF Routing,” for more information.



CAUTION: Set the area ID to the same value for all routers on the network segment because they are in the same area.

Cost

The interface cost parameter reflects the line speed of the port. Although the system calculates a default cost based on the module media type, you can set the cost manually to a different value. In most cases, you can accept the default value set by the system.

Delay

The system adds the OSPF interface transmit delay (in seconds) to all LSAs that it sends out to the network.



Set the transmit delay according to the link speed; use a longer transmit delay for slower link speeds.

Default The default delay is 1 second.

Hello Timer

The hello timer (in seconds) determines how often the interface transmits hello packets to other routers. *Hello packets* notify other routers that the sending router is still active on the network. If a router does not send hello packets for the period of time specified by the dead interval, that router is considered inactive by its neighbors. For more information, see “Dead Interval,” later in this chapter.

Default The default value for the hello timer is 10 seconds.



CAUTION: Set the hello timer to the same value for all routers on the same network segment.

Retransmit Timer

The retransmit interval (in seconds) determines how often the interface retransmits LSAs to other routers.

Default The default value for the retransmit timer is 5 seconds.

Dead Interval

The dead interval determines how long neighbor routers wait for a hello packet before they determine that a neighbor is inactive. A router that receives a hello packet from a neighbor resets its dead interval timer for the neighbor.



CAUTION: Set the dead interval to the same value for all routers on the same network segment.

Default The default value for the dead interval is 40 seconds.

Password

You can set security passwords for specific OSPF interfaces.



CAUTION: Use the same password for all routers on the same network segment.

Default The default is no password.

Stub Default Metrics

The stub default metric value determines if an area border router generates the default route into a stub area. For example, in Figure 7-1, you configure area border router 1 to generate the default route into stub area 2.

Virtual Links

You can configure a virtual link between an area border router that links two nonbackbone areas and a second area border router that is connected to the backbone area 0. The nonbackbone area through which the virtual link runs is called a *transit area*. Virtual links are used to ensure that the OSPF backbone is contiguous. The endpoints of a virtual link must be area border routers (such as area border routers 3 and 4 in Figure 7-1). You must configure the virtual link on both routers. Each router's virtual link definition includes the other router's address and the transit area through which the routers connect (for example, area 3 in Figure 7-1).



You must configure a virtual link for any area border router that has an interface connected to a location outside the backbone area. For example, in Figure 7-1, area border router 4 has an interface connected to nonbackbone area 1, which makes the virtual link to the backbone area necessary. You can define up to 32 virtual links per router.



CAUTION: Do not configure a virtual link through a stub area.

How OSPF Routing Works

This section summarizes how the OSPF algorithm works for a router that meets these characteristics:

- Lies within an autonomous system area (an interior router)
- Is attached to a multiaccess network
- Is configured to be the designated router for its network

Starting Up

When the router starts, the interfaces that are configured to run OSPF begin in the *down* state. When the lower-level IP protocols indicate that an interface is available, the interface moves to the *waiting* state. It remains in this state until the designated router and backup designated router are chosen.

Finding Neighbors

The router sends out hello packets to locate its network neighbors. These packets also list the routers from which the sending router has *received* hello packets. When a router detects its own address in another router's hello packet, the two routers establish two-way communications as neighbors.

Establishing Adjacencies

If neighboring OSPF routers succeed in exchanging and synchronizing their link state databases, they appear as *adjacent* in all router and network link advertisements.

Electing the Backup Designated Router

OSPF selects a backup designated router for the network segment. This router takes over as the designated router if the current designated router fails.

The OSPF algorithm first eliminates all routers that have an assigned priority of 0. OSPF then selects the backup designated router from among the routers on the segment that have *not* declared themselves to be the designated router (based on their configuration settings). If some routers have already declared themselves to be the backup designated router, OSPF limits the selection to one of them.

OSPF selects the candidate router with the highest priority. If candidate routers have the same priority, OSPF selects the router with the highest router ID.

Electing the Designated Router

OSPF selects a designated router, which originates LSAs on behalf of the network segment. These advertisements list all routers (including the designated router) that are attached to the segment. The designated router also floods LSA packets throughout the segment to allow its neighbors to update their databases.

The OSPF algorithm first eliminates all routers that have an assigned priority of 0. OSPF then selects a designated router from among the routers that have declared themselves to be the designated router (based on their configuration settings). If no routers have declared their candidacy, the backup designated router becomes the designated router, and OSPF selects a new backup designated router.

OSPF selects the candidate router with the highest priority. If candidate routers have the same priority, OSPF selects the router that has the highest router ID.



The designated router also becomes adjacent to all other routers on the network segment by sending hello packets to them.

Calculating Shortest Path Trees

OSPF routers collect raw topological data from the LSAs that they receive. Each router then prunes this data down to a tree of the shortest network paths centered on itself. In a series of iterations, the router examines the total cost to reach each router or network node in its domain. By discarding all but the lowest-cost path to each destination, the router builds a shortest path tree to each destination, which it uses until the network topology changes.

Routing Packets

A packet's source and destination determine the routers that move it:

- **Intraarea** — When a packet's source and destination are in the same area, the packet is routed using internal router databases. No routers are used outside the area.
- **Interarea** — When a packet's source and destination are in different areas, the topology databases in the backbone area dictate the paths that are taken between areas.



You can use virtual links to influence the routes that are taken for interarea traffic.

- **To a stub area** — When a packet's destination is in a stub area (an area that does not accept external route advertisements), OSPF uses the area's predefined default route.




You configure default routing in area border routers that serve an OSPF stub area, such as area border router 1 in Figure 7-1.

- **To a different autonomous system** — When a packet's source and destination are in different autonomous systems, autonomous system boundary routers compute the routing paths using data obtained from another protocol, such as the Border Gateway Protocol. The boundary routers flood these external routes throughout the autonomous system.



Boundary routers do not flood external routes into stub areas.



ROUTING WITH APPLETALK TECHNOLOGY

This chapter provides an overview of AppleTalk routing, including these topics:

- About AppleTalk
- AppleTalk Network Elements
- AppleTalk Protocols
- About AARP

About AppleTalk

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

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

AppleTalk Network Elements

An AppleTalk network consists of different nodes and groups of networks. Nodes can include workstations, routers, printers, and servers that provide services for other computers, called *clients*.

This section describes the elements of an AppleTalk network:

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

AppleTalk Networks

A subnetwork in an AppleTalk intranet is a cable segment attached to a router. Each subnetwork is identified by a network number or range of network numbers. You assign 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 CoreBuilder™ 2500 system does not translate Phase 1 packets to Phase 2 packets, it does route packets to a Phase 1 network. The CoreBuilder 2500 system anticipates that a gateway exists between the two networks to translate the packets.

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

AppleTalk Nodes

A node in a AppleTalk network is any addressable device, including a workstation, printer, or router. Nodes are physically attached to a network. At initialization, each node in an AppleTalk network selects a unique AppleTalk address. 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, you can give the device a name. The name appears on the *Chooser* menu of the Macintosh with an associated icon. For example, the Chooser of the Macintosh can include a printer icon. When the user selects the printer icon, several printer names can appear in a list, such as `Laser1` or `Laser 2`. The Name Binding Protocol (NBP), described later in this chapter, translates these device names into AppleTalk addresses.

AppleTalk Zones

An AppleTalk zone is a logical collection of nodes on an AppleTalk intranet. 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 intranet.

Figure 8-1 illustrates the relationship between physical AppleTalk networks and logical AppleTalk zones.

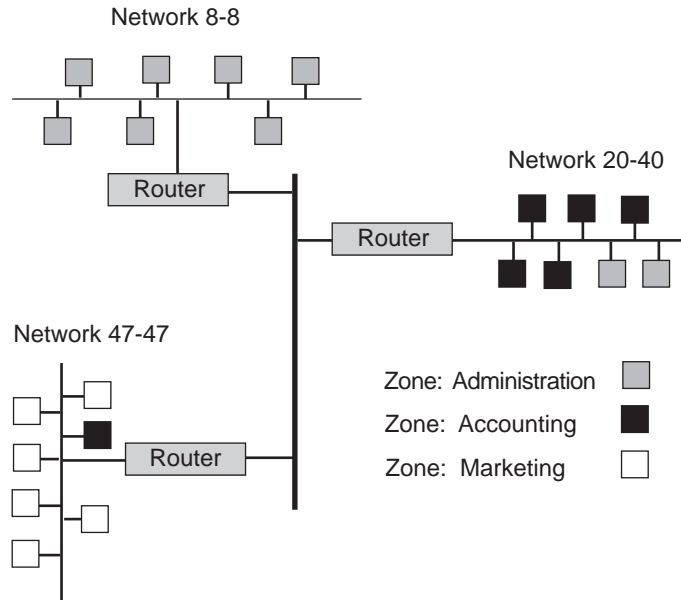


Figure 8-1 AppleTalk Networks and Zones

This example shows an AppleTalk intranet with three subnetworks: 47-47, 20-40, and 8-8. Three AppleTalk zones span these networks: 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 and all nodes in the Marketing zone. Network 8-8 consists of nodes in the Administration zone only.

Creating zones within a network reduces the amount of searching that a router must do to find a resource on the network. For example, to gain access to a printer on the network, instead of searching the whole network when you want to print a file to a certain printer, the router searches for it within a particular zone. You gain access to the printer more quickly within the zone because the zone includes fewer devices than the entire intranet.

Seed Routers

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

A nonseed router listens for a seed router and takes configuration information from the first one it detects. A nonseed router that obtains configuration data participates in the network as if it is a seed router.

AppleTalk Protocols

AppleTalk protocols ensure the flow of information through AppleTalk networks. Figure 8-2 shows a simplified view of AppleTalk protocols and their relationship to the OSI Reference Model. These protocols provide physical connectivity, end-to-end network services, and data delivery.

OSI Reference Model

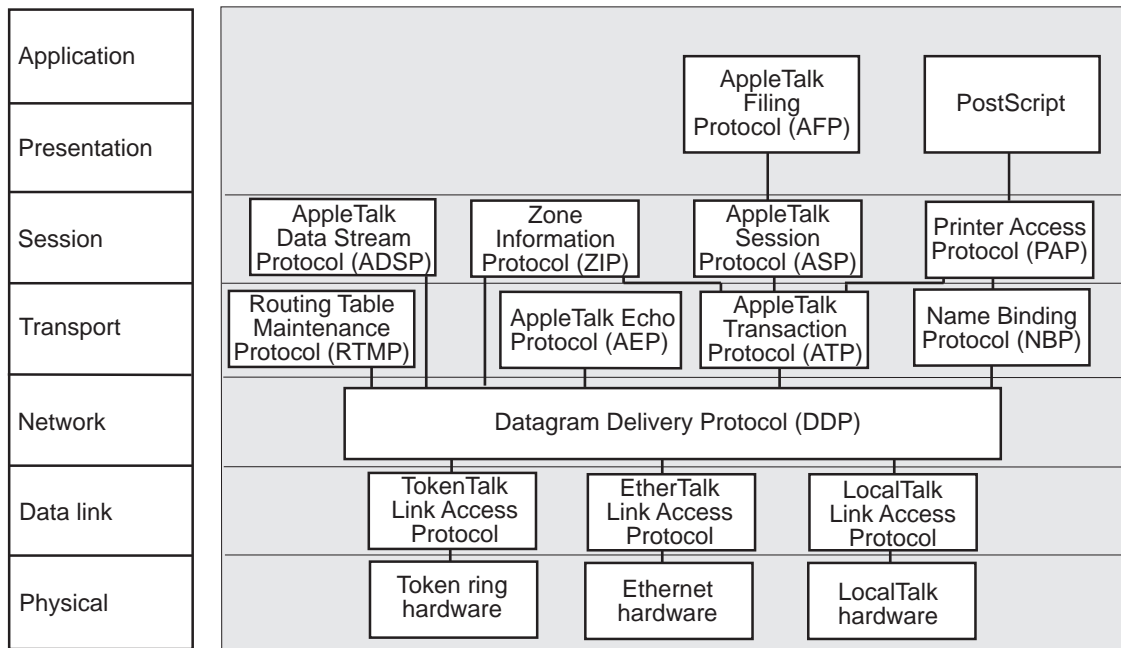


Figure 8-2 AppleTalk Protocols and the OSI Reference Model

The AppleTalk six-layer protocol suite does not fully comply with the OSI seven-layer model. However, AppleTalk provides many of the functions and services of OSI. AppleTalk has no specific protocols for the Application layer because the lower levels provide printer and file service.

**Physical Layer
Protocols**

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

Link Layer Protocols

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 (LAPs):

- TokenTalk LAP (TLAP)
- Ethernet LAP (ELAP)
- LocalTalk LAP (LLAP).



The AppleTalk Address Resolution Protocol (AARP), which translates hardware addresses to AppleTalk addresses, also exists at the data link layer because it is closely related to the Ethernet and token ring LAPs. AARP is usually included in the definition of each LAP, so it does not appear in the reference model. See “About AARP” on page 8-10 for more information about this protocol.

**Network Layer
Protocols**

The network layer accepts data from the layers above it and divides the data into packets to send over the network through the layers below it. The Datagram Delivery Protocol (DDP) transfers data in packets called *datagrams*.

Datagram delivery is the basis for building other AppleTalk services such as electronic mail. With DDP, AppleTalk runs as a process-to-process, best-effort delivery system in which the processes running in the nodes of interconnected networks exchange packets with each other.

**Transport Layer
Protocols**

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 these protocols and provides more detail for the protocols that you can view using the CoreBuilder 2500 Administration Console.

An AppleTalk intranet has four transport layer protocols:

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

Routing Table Maintenance Protocol (RTMP)

This protocol maintains information about AppleTalk addresses and connections between different networks. It specifies that each router:

- Learns new routes from other routers
- Deletes a route if the local router has not broadcast the route to the network for a certain period of time

Each router builds a routing table for dynamic routing operations in an AppleTalk intranet. Every 10 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*, or *really bad*)

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

- 1 After a specified 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 The router removes the entry of a nonresponding router with a *really bad* status.

The data in the routing table is cross-referenced to the Zone Information Table (ZIT). This table maps networks into zones. See "Session Layer Protocols" on page 8-8 for more information about the ZIT.

Figure 8-3 illustrates a simple AppleTalk network, and Table 8-1 shows the corresponding routing table.

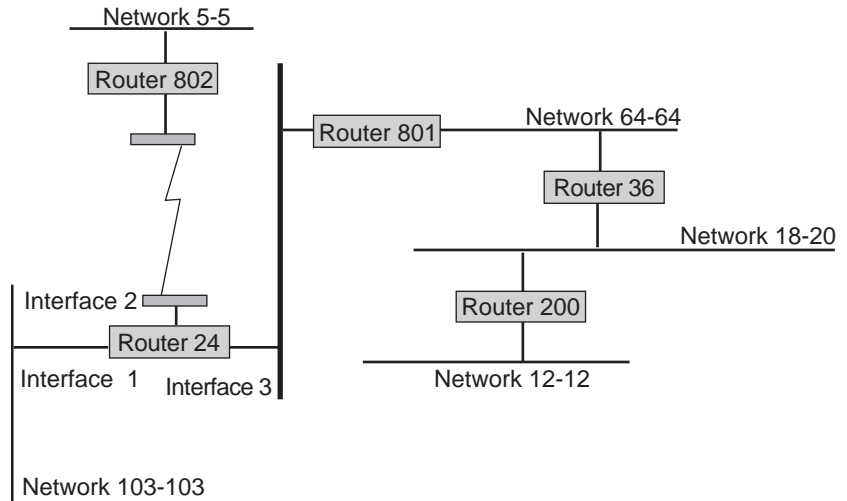


Figure 8-3 A Simple AppleTalk Network

Table 8-1 Routing Table for Router 24 in Figure 8-3

Network Range	Distance (hops)	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 view the AppleTalk routing tables in your network through the Administration Console.

AppleTalk Echo Protocol (AEP)

AppleTalk nodes use the AEP to send datagrams to other nodes in the network. The AEP datagram transmitted causes the destination node to return, or *echo*, the datagram to the sending node. This protocol determines whether a node is accessible before any sessions are started, and it enables users to estimate the round-trip delay time between nodes.

AppleTalk Transaction Protocol (ATP)

This protocol, along with the AppleTalk Data Stream Protocol (ADSP), ensures delivery of DDP packets to a destination without any losses or corruption.

Name Binding Protocol (NBP)

This protocol translates alphanumeric entity names to AppleTalk addresses. NBP maintains a table of node addresses and named entities within each node. Because each node also maintains its own list of named entities, the names directory within an AppleTalk network is not centralized. The names directory database is distributed among all nodes on the intranet.

Session Layer Protocols

An AppleTalk intranet has four session-layer protocols:

- AppleTalk Data Stream Protocol (ADSP)
- Zone Information Protocol (ZIP)
- AppleTalk Session Protocol (ASP)
- Printer Access Protocol (PAP)

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. Therefore, two nodes can communicate simultaneously. ADSP also includes flow control, so that a fast sender does not overwhelm a slow receiver.

Zone Information Protocol (ZIP)

ZIP works with RTMP to map network numbers to network zones for the entire AppleTalk intranet. Network zones are the logical groupings of AppleTalk networks. The table created by ZIP is called the *Zone Information Table (ZIT)*. You view the ZIT by network number or network zone from the Administration Console.

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, the router compares the zone name in the packet with zone names in the ZIT entries. The router then compares the network number in the matching ZIT entry with the network number in the RTMP table, to find the interface for routing the packet.

AppleTalk Session Protocol (ASP)

The ASP passes commands between a workstation and a server after they connect to each other. ASP ensures that the commands are delivered in the same order that they were sent and returns the results of these commands to the workstation.

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.

Presentation Layer Protocols


The presentation layer maintains information about files, formats, and translations between formats. An AppleTalk intranet has two protocols at the presentation layer: the AppleTalk Filing Protocol (AFP) and PostScript. AFP provides remote access to files on the network. PostScript is a graphic page 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. AARP maps 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, the node sends 10 AARP probe packets. The probe packets determine whether any other nodes on the network are using the selected address. If the address already exists, the initializing node randomly selects another address and sends another set of probe packets.

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, AppleTalk sends a request to the protocol address and adds the hardware address to the table when the destination node replies. You view this table, called the *AARP cache*, through the CoreBuilder 2500 Administration Console.



ADMINISTERING EXTENDED SWITCHING FEATURES

- Chapter 9** Administering VLANs
- Chapter 10** Administering IP Routing
- Chapter 11** Administering IP Multicast Routing
- Chapter 12** Administering IPX Routing
- Chapter 13** Administering OSPF Routing
- Chapter 14** Administering AppleTalk Routing



9

ADMINISTERING VLANs

This chapter describes how to display information about and configure VLANs, in these sections:

- Displaying VLAN Information
- Defining VLAN Information
- Modifying VLAN Information
- Removing a VLAN Definition

Displaying VLAN Information

You can display a summary of VLAN information or a detailed report. When you display a summary, you receive information about the protocols and ports assigned to each VLAN, plus the Layer 3 addresses that are used to manage flood domains for overlapping IP subnetworks. The detailed VLAN report includes the summary information plus additional utilization statistics.

To display VLAN information, from the top level of the Administration Console, enter:

```
bridge vlan summary
```

or

```
bridge vlan detail
```

The VLAN information appears in the format that you specified.

Top-Level Menu

```
system
ethernet
-----
fddi
atm
bridge
ip
ipx
appletalk
snmp
analyzer
script
logout
display
mode
ipFragmentation
ipxSnapTranslation
addressThreshold
agingTime
stpState
stpFollowList
stpPriority
stpMaxAge
stpHelloTime
stpForwarding
stpGroupAd
port
packetFilter
vlan
summary
detail
define
modify
remove
```

The following sample shows a summary display for several VLANs:

- VLAN summary

```

Index      Protocol  Identifier  Ports
  1         default      0         1-18
  2          IP        2          3

Index      Name                Layer 3
  1         none
  2      150.14         158.101.150.0 255.255.255.0

```

The following sample shows a detailed display for these VLANs:

- VLAN detail

```

Index      Protocol  Identifier  Ports
  1         default      0         1-18
  2          IP        2          3

Index      Name                Layer 3
  1         none
  2      150.14         158.101.150.0 255.255.255.0

Index      inPackets  inBytes  outPackets  outBytes
  1           54       7654         54         6897
  2          453      181028        453      181028

```

Table 9-1 describes these statistics.

Table 9-1 Field Attributes for VLAN Information

Field	Description
Index	A system-assigned index that identifies a VLAN
Protocol	The protocol suite of the VLAN
Identifier	A unique, user-defined (4-byte) integer for use by global management operations
Ports	The numbers of the ports that are assigned to the VLAN
Name	A 16-byte character string that identifies the members of the VLAN
Layer 3	Optional parameters (consisting of IP subnetwork and mask) that are used to set up flood domains for overlapping IP VLAN subnetworks
inPackets	The number of flooded packets that are received on the VLAN
inBytes	The number of flooded bytes that are received on the VLAN
outPackets	The number of flooded packets that are transmitted over the VLAN
outBytes	The number of flooded bytes that are transmitted over the VLAN

Defining VLAN Information

Top-Level Menu

```

system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  vlan
    display
    mode
    ipFragmentation
    ipxSnapTranslation
    addressThreshold
    agingTime
    stpState
    stpFollowLia
    stpPriority
    stpMaxAge
    stpHelloTim
    stpForward
    stpGroupAd
    port
    packetFilter
    vian
      summary
      detail
      define
      modify
      remove
  
```

To create a VLAN definition:

- 1 From the top level of the Administration Console, enter:
bridge vlan define
- 2 Enter the appropriate protocol: (**IP, IPX, Apple, XNS, DECnet, SNA, Vines, X.25, NetBIOS, default**)
- 3 Enter the integer of the VLAN interface identifier.
- 4 Enter the VLAN name. The system displays the port types.
- 5 Enter the numbers of the ports to assign to the VLAN or **a11** to assign all ports to the VLAN.
 - If you did not choose the IP protocol, you are finished.
 - If you chose the IP protocol, proceed to steps 6 and 7.
- 6 To use Layer 3 subnetwork addressing, enter **defined** and proceed to step 8. If you do not want to use Layer 3 subnetwork addressing, enter **undefined** and you are finished.
- 7 Enter the IP subnetwork address.
- 8 Enter the subnet mask.

Example:

```

Select menu option (bridge/vlan): define
Enter Protocol Suite (IP,IPX,Apple,XNS,DECnet,SNA,
Vines,X.25,NetBIOS,default): IP
Enter Integer VLAN Identifier: 1
Enter VLAN Name: SD Marketing
Ports 1-2=FDDI, 3-18=Ethernet
Enter port(s) (1-18|all): 1,3-5
Layer 3 Address (undefined, defined): defined
Enter IP Subnet Address: 158.111.122.0
Enter subnet mask [255.255.0.0] 255.255.255.0
  
```



You can define up to 32 VLANs on a single bridge. The CoreBuilder 2500 VLAN database allows up to 16 simultaneous VLAN protocols and reserves the first seven database entries for default, IP, IPX, and AppleTalk VLANs. The XNS and DEC protocols use three and five of the nine remaining entries, respectively. If you define DEC and XNS, plus two additional SNA, VINES, X25, or NetBIOS protocols, you receive a message that the system's database resources are exhausted.

Modifying VLAN Information

Top-Level Menu

```

system
ethernet
-----
fddi
atm
bridge
ip
ipx
appletalk
snmp
analyzer
script
logout
-----
display
mode
ipFragmentation
ipxSnapTranslation
addressThreshold
agingTime
stpState
stpFollowLi
stpPriority
stpMaxAge
stpHelloTim
stpForward
stpGroupAd
port
packetFilter
vlan
-----
summary
detail
define
modify
remove

```

To modify VLAN information:

- 1 From the top level of the Administration Console, enter:
bridge vlan modify
- 2 Enter the number of the VLAN interface index.
- 3 Enter the protocol for that VLAN: (**IP, IPX, Apple, XNS, DECnet, SNA, Vines, X.25, NetBIOS, default**).
- 4 Enter the VLAN identifier.
- 5 Enter the VLAN name.
The system displays the port types.
- 6 Enter the numbers of the ports to modify or **all** to modify all ports on the VLAN.
- 7 If you have selected the IP protocol and want to use the Layer 3 address information, enter **defined**

If you do not want Layer 3 addressing, enter **undefined**

Example:

```

Select menu option (bridge/vlan): modify
Select VLAN interface [1-2]: 2
Protocol Suite (IP,IPX,Apple,XNS,DECnet,SNA,
Vines,X.25,NetBIOS,default): IP
Integer VLAN Identifier [1]: 2
VLAN Name [Sales]:
Ports 1=FDDI, 2-17=FastEthernet
Enter port(s) (1-17|all) [1-5]:
Layer 3 Address (undefined,defined) [undefined]:

```

Removing a VLAN Definition

To remove a VLAN definition:

- 1 From the top level of the Administration Console, enter:
bridge vlan remove
- 2 Enter the indexes for the VLANs that you want to remove.

Example:

Select menu option (bridge/vlan): **remove**

Select VLAN index(es) (1-2|all): **1**

Top-Level Menu

```
system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  ethernet
    display
    mode
  ip
    ipFragmentation
    ipxSnapTranslation
    addressThreshold
    agingTime
    stpState
    stpFollowLip
    stpPriority
    stpMaxAge
    stpHelloTim
    stpForward
    stpGroupAd
    port
    packetFilter
  vlan
    summary
    detail
    define
    modify
    remove
```


10

ADMINISTERING IP ROUTING

This chapter describes how to set up and manage your CoreBuilder™ 2500 system to route Internet Protocol (IP) packets. The chapter discusses these tasks:

- Administering Interfaces
- Administering Routes
- Administering the ARP Cache
- Administering ATM ARP Servers (for CoreBuilder 2500 systems that have ATM modules)
- Administering UDP Helper
- Enabling and Disabling IP Routing
- Enabling and Disabling ICMP Router Discovery
- Configuring RIP (Routing Information Protocol)
- Pinging an IP Station
- Displaying IP Statistics

For more information about how IP routing works, see Chapter 4, “Routing with IP Technology.”

For more information about Open Shortest Path First (OSPF) protocol, see Chapter 7 and Chapter 13.

Administering Interfaces

You can define two types of IP interfaces through CoreBuilder 2500 Extended Switching software: IP VLAN interfaces and IP LIS interfaces. This section describes these interfaces and how to administer them.

IP VLAN Interfaces

An IP VLAN interface defines the relationship between an IP virtual LAN (VLAN) and the subnets in the IP network. Every IP VLAN interface has one IP VLAN associated with it. Each Ethernet and FDDI module has one interface defined for each subnetwork that is directly connected to it.



You must first define a VLAN, as described in Chapter 9, Administering VLANs, before you define an associated IP VLAN interface.

IP LIS Interfaces

A logical IP subnet (LIS) interface supports classical IP over ATM. You define LIS interfaces only for the ports on ATM modules. See Chapter 11 of the *CoreBuilder 2500 Operation Guide* for more information about the ATM protocol. See the *CoreBuilder 2500 Administration Console User Guide* for information about how to configure ATM ports.

Interface Characteristics


Each IP interface has the following information associated with it:

- **IP Address** — The address specific to your network. Choose an IP address from the range of addresses assigned to your organization. This address defines both the number of the network to which the interface is attached and its host number on that network.
- **Subnet Mask** — A 32-bit number that uses the same format and representation as an IP address. The *subnet mask* determines the bits in the IP address that are interpreted as the network number, the subnetwork number, and the host number. Each IP address bit that corresponds to a *1* in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a *0* is in the host part of the IP address.
- **Type** — *VLAN*, which supports routing between two VLANs, or *LIS*, which supports classical IP over ATM
- **State** — This status of the IP interface indicates whether the interface is available for communications.
- **VLAN index number** — When you select `vlan` as the interface type, the system prompts you for the VLAN index number, which indicates which bridge ports are associated with the IP interface. The system prompt includes a list of available VLAN indexes and the bridge ports associated with them.

- **LIS interface information** — When you select **lis** as the interface type, the system prompts you for LIS interface information. The information you enter depends on whether you define permanent virtual circuits (PVCs), switched virtual circuits (SVCs), or both on the LIS interface. See the *CoreBuilder 2500 Operation Guide* for more on PVCs and SVCs.
 - **If you define only SVCs**, enter an ATM ARP server address. This server maintains the IP-to-ATM address translation table. You can enter the maximum number of SVCs allowed on this interface. The minimum holding time determines the least amount of time that an SVC connection remains open. The inactivity timer determines how long the connection can remain open with no activity after the minimum holding time has expired. You also need to enter the ATM port number for this interface.
 - **If you define only PVCs** on the interface, you need to enter only the PVC numbers and the ATM port number. The other prompts do not appear because you need not enter an ATM ARP server address.
 - **If you define SVCs and PVCs**, enter all LIS interface information.

Displaying Interfaces

You can display a table that shows all of the IP interfaces that are configured for the system, including their parameter settings. The IP interface display screen includes the values that you select when you define the interface.



- 1 From the top level of the Administration Console, enter:
ip interface display
- 2 Enter the interface index numbers or **all**.

Select interface index(es) (1-2|all): **all**

Example:

```
IP forwarding is enabled, RIP is active, ICMP discovery is disabled.
```

Index	Type	IP address	Subnet mask	State	VLAN	index
1	VLAN	158.101.1.1	255.255.255.0	Down		2

Index	Type	IP address	Subnet mask	State	Port
2	LIS	158.101.112.1	255.255.255.0	Up	1

```
atmArpServer
47-0000-00-000000-0000-0000-00cc-080001200054-ff

maxSvcCount      inactivityTime    minHoldingTime
0                  1200              60
```

Defining an IP VLAN Interface

When you define an IP VLAN interface, you specify several interface characteristics, as well as the index of the VLAN that is associated with the interface.



You must first define a VLAN, as described in Chapter 9, Administering IP Routing, before you define an associated IP VLAN interface.

To define an IP VLAN interface:

- 1 From the top level of the Administration Console, enter:
ip interface define
- 2 You are prompted for the interface's parameters.
To accept the existing 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 you want to connect the interface.
- 5 Enter the type of IP interface: **vlan**
- 6 Enter the index of the VLAN that is associated with the interface.

Example:

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
IP VLANs:
  Index  Ports
     3    1-8
     4    9-12
Select VLAN index: 3
```



If you physically change the configuration of your system after you define IP interfaces, the ports designated for those interfaces may no longer be valid. Verify whether you need to reconfigure your interfaces.

Top-Level Menu

```
system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
interface
  route
  arp
  atmArpServer
  multicast
  routing
  icmpRouterDiscovery
  ospf
  rip
  ping
  statistics
  display
  define
  modify
  remove
  addPvc
  removePvc
```


Defining an IP LIS Interface

When you define an IP LIS interface, you specify several general IP interface characteristics and IP LIS characteristics.



Before you define an IP LIS interface with SVCs, be sure that you have defined an ATM ARP server, as described in the section “Administering ATM ARP Servers” later in this chapter. If the LIS interface has only PVCs, you do not need to define an ATM ARP server.


To define an IP LIS interface:

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


```
ip interface define
```
- 2 You are prompted for the interface's parameters.

To accept the existing 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 you want to connect the interface.
- 5 Enter the type of IP interface: **lis**
- 6 Enter the LIS information:
 - **For a LIS interface with SVCs** — Enter the ATM ARP server address, the maximum SVC count, the inactivity timer, the minimum holding time, and the ATM port associated with the interface. Press Return to accept the existing values shown in brackets.
 - **For a LIS interface with only PVCs** — Enter the ATM port and the PVCs associated with the interface. You can enter up to 51 PVCs for each interface. (The maximum number of PVCs on the CoreBuilder 2500 system is 64.)

Top-Level Menu



LIS interface example with both PVCs and SVCs:

```

Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter interface type (vlan,lis) [lis]:
Enter ATM arp server address
[00-0000-00-000000-0000-0000-0000-000000000000-00]:47-0000-00-000000-
00000000-00cc -000000000001-ff
Accept completed ATM address (yes,no) [yes]:
Enter max. SVC count (0=no max.0) [0]:
Enter inactivity time (0=infinite, 10-10000) seconds [1200]:
Enter min. holding time (0-10000) seconds [60]:
Select ATM port [1]:
Enter PVC(s) (VPI/VCI)[]: 1/32,1/200,1/3330
  
```

Modifying an Interface

You can change the configuration of an interface that you have already defined.



You can add one or more advertisement addresses or PVCs to an interface with the `ip rip addAdvertisement` and `ip interface addPvc` commands as well as with the `ip interface modify` command. If you add or change an advertisement address or PVC with the `modify` command, you must reenter all addresses or PVCs that are associated with the interface, not only the one that you want to add or change.


To modify an IP interface:

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

```
ip interface modify
```

You are prompted for the interface parameters.

- 2 Modify existing interface parameters by entering new values at the prompts. Press Return at the prompts to accept the values for the parameters that you do not want to modify.



Removing an Interface

You can remove an interface if you no longer route on the ports that are 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 index numbers of the interfaces you want to remove.


If you have defined one or more PVCs on the interface, the Administration Console displays a message indicating that the PVCs will be removed with the interface.

This example is a prompt for interface 2, which has one PVC associated with it:

```
1 PVC associated with the interface index 2. Do you wish to
continue (yes/no) [yes]:
```

Press Return to accept the default (**yes**) to delete the interface.


Deleted interfaces are removed immediately.



Adding a Permanent Virtual Circuit (PVC)

To add a PVC to an LIS interface:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:
ip interface addPvc
- 2 Enter the index interface number that you want to associate with the PVC.
- 3 Enter the Virtual Path Interface (VPI) and the Virtual Circuit Interface (VCI) pairs in this format: **VPI/VCI**. Separate additional entries with a comma.


Example:

```
Select interface index [1]: 1
Enter [VPI/VCI]: 2/20
```

Removing a Permanent Virtual Circuit (PVC)

You can remove one or more PVCs that are associated with an LIS interface.

Top-Level Menu



- 1 From the top level of the Administration Console, enter:
ip interface removePVC
- 2 Enter the index number of the interface and the VPI/VCI pair that you want to remove.

Administering Routes

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

Each routing table entry contains the following information:

- **Destination IP Address and Subnet Mask** — The elements that define the address of the destination network, subnetwork, or host. An incoming packet matches a route entry in the router table when the packet's destination address matches the destination address in the routing table entry. (The system only compares bits that correspond to the subnet mask in the routing table entry.)

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

- **Routing Metric** — The number of networks or subnetworks through which a packet must pass to reach its destination. This metric is included in RIP updates to allow routers to compare routing information that is received from different sources.
- **Gateway IP Address** — Tells the router where to forward packets whose destination address matches the route's IP address and subnet mask. The system forwards such packets to the indicated gateway.
- **Status** — The route descriptions, as listed in Table 10-1.

Table 10-1 Route Status

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

Default route


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

Displaying the Routing Table

You display a module's routing table to determine the routes that are configured and whether they are operational.

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

```
ip route display
```



This example shows part of a routing table:

```
IP routing is enabled, ICMP router discovery is disabled
```

Destination	Subnet mask	Metric	Gateway	Status
Default Route	--	3	158.101.112.254	Learned RIP
9.0.0.0	255.0.0.0	3	158.101.112.254	Learned RIP
10.0.0.0	255.0.0.0	3	158.101.112.254	Learned RIP
89.0.0.0	255.0.0.0	3	158.101.112.254	Learned RIP
129.213.0.0	255.255.0.0	3	158.101.112.254	Learned RIP

Depending on the system configuration, the table displays IP, RIP, and OSPF routing configurations and destination default routes.

Defining a Static Route

Before you define static routes, define at least one IP interface. (See "Defining an IP VLAN Interface" earlier in this chapter for more information.) 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 the system's periodic RIP updates. They are included in OSPF updates as type 1 externals.

To define a static route:

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

```
ip route static
```

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



Example:

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

Removing a Route

To remove a route:

Top-Level Menu

```
system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  interface
    route
    arp
    atmAr
    multi
    udpHe
    routin
    icmpR
    ospf
    rip
    ping
    statistics
    display
    static
    remove
    flush
    default
    noDefault
```

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

The route is immediately deleted from the routing table.

Flushing a Route

Flushing deletes all learned routes from the routing table.

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

```
ip route flush
```

All learned routes are immediately deleted from the routing table.

Top-Level Menu

```
system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  interface
    route
    arp
    atmAr
    multi
    udpHe
    routin
    icmpR
    ospf
    rip
    ping
    statistics
    display
    static
    remove
    flush
    default
    noDefault
```



If OSPF is enabled, flushing a route results in an OSPF shortest-path-first computation. The routing table recalculation takes a few seconds.

Setting the Default Route

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

To statically configure the default route:

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

The default route is immediately added to the routing table.

Top-Level Menu

```
system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  interface
    route
    arp
    atmAr
    multi
    udpHe
    routin
    icmpR
    ospf
    rip
    ping
    statistics
    display
    static
    remove
    flush
    default
    noDefault
```


Removing the Default Route



You can remove a module's default route.

Be careful when you remove a module's default route. If a system's routing table does not contain a default route — either statically configured or learned using RIP or OSPF — then the module cannot forward a packet that matches no routing table entries. When this occurs, the module drops the unroutable packet and sends an ICMP destination unreachable message to the host that sent the packet.

Top-Level Menu



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

ip route noDefault

The default route is immediately removed from the table.

Administering the ARP Cache


The CoreBuilder 2500 system uses the Address Resolution Protocol (ARP) to find the MAC addresses that correspond to the IP addresses of hosts and routers on the same subnetworks. An ARP cache is a table of known IP addresses and their corresponding MAC addresses.

Displaying the ARP Cache

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

ip arp display

Top-Level Menu



Sample ARP cache display:

```


IP routing is disabled, ICMP router discovery is disabled

IP address      I/F                Hardware address Circuit
158.101.43.33   1                  08-00-09-6d-b6-70  -/-
158.101.43.254 1                  08-00-02-0a-16-09  -/-
  
```

Removing an ARP Cache Entry

You can remove an entry from the ARP cache if the MAC address has changed.

Top-Level Menu



To remove an entry from the ARP cache:


- 1 From the top level of the Administration Console, enter:
ip arp remove
- 2 Enter the IP address that you want to remove.

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

Flushing the ARP Cache

You can delete all entries from the ARP cache if the MAC address has changed.

Top-Level Menu



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

ip arp flush

All ARP cache entries are immediately removed from the table.

Administering ATM ARP Servers

If you are running classical IP over ATM with SVCs, follow these rules:

- You must define an ATM ARP server for each LIS.
- Each LIS must connect to a single ATM network.
- Each LIS must belong to the same IP subnetwork.



The **atmArpServer** menu also includes the **arp** option, which allows you to administer the ATM ARP cache.

Displaying ATM ARP Servers

To display a list of ATM ARP servers, from the top level of the Administration Console, enter:

```
ip atmArpServer display
```

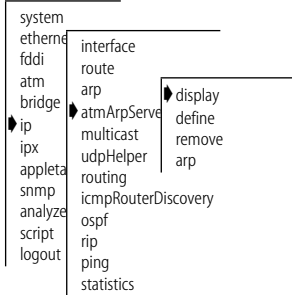
Example:

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

Index	Port	IP Address	Subnet Mask
1	1	158.101.1.1	255.255.255.0

```
ATM address
47-0000-00-000000-0000-0000-00cc-000000000000-ff
```

Top-Level Menu



Defining an ATM ARP Server

Determine the location of the ATM ARP server that you want to use. You can define the ATM ARP server externally on another CoreBuilder 2500 system or on an ATM switch, such as 3Com's CoreBuilder 7000 system.

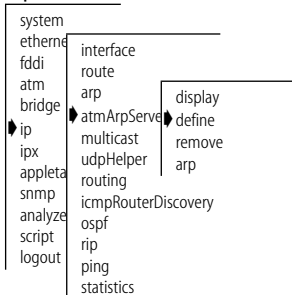
- 1 To define an ATM ARP server, from the top level of the Administration Console, enter:


```
ip atmArpServer define
```
- 2 Enter the number of the ATM port on which you want to define the ATM ARP server.
- 3 Enter the IP address of the ATM port that you want to define.
- 4 Enter the subnet mask. To accept the existing or default value shown in brackets, press Return.

Example:

```
Select ATM port [1]
Enter IP address: 158.101.20.30
Enter subnet mask [255.255.0.0]
```

Top-Level Menu



Removing an ATM ARP Server

To delete a currently defined ATM ARP server, from the top level of the Administration Console, enter:

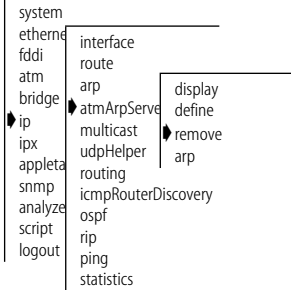
```
ip atmArpServer remove
```

The system prompts you for one or more of the index numbers associated with the ATM ARP servers that you want to remove. The ATM ARP server display shows the index number assigned to each ATM ARP server. The system also displays the current index numbers in brackets.

Example:

```
Select ATM ARP server index(es) [1-2,all]: 1
```

Top-Level Menu



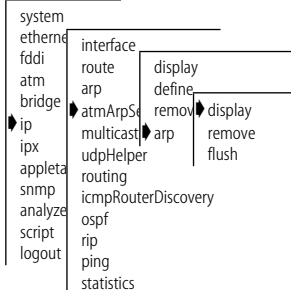
Displaying the ATM ARP Cache

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

```
ip atmArpServer arp display
```

Example:

Top-Level Menu

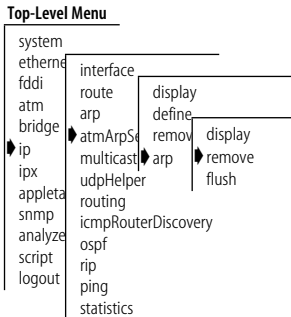


```
IP routing is enabled, RIP is active,
ICMP router discovery is disabled
```

IP address	ATM Address	Circuit
158.101.112.2	47-005-80-ffe100-0000-f21a-2130-80000212d0f-18	1/32
158.101.112.7	47-005-80-ffe100-0000-f22a-2130-80000211d01-18	1/33
158.101.116.7	47-005-81-ffe100-0000-f21a-2130-80000112d01-18	2/20
158.101.112.14	47-005-81-ffe100-0000-f21a-2130-80000112d01-18	2/22

Removing an ATM ARP Cache Entry

To remove an entry from the ATM ARP cache, from the top level of the Administration Console, enter:




ip atmArpServer arp remove

Enter the ATM address that you want to remove from the cache.

The address is immediately removed from the table.

Flushing the ATM ARP Cache

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



ip atmArpServer arp flush

All ATM ARP cache entries are immediately removed from the table.

Administering UDP Helper

Use UDP Helper 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. UDP Helper 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.

These ports for UDP services are mentioned in this section on UDP Helper:

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

With UDP Helper, you can configure the interval of time during which forwarding can occur 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 routing addresses that are defined for your system.

Top-Level Menu

```

system
ethernet
  interface
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  route
  arp
  atmAr
  multi
  udph
  routin
  icmpR
  ospf
  rip
  ping
  statistics
  display
  define
  remove
  hopCountLimit
  threshold
  interface

```

To display UDP Helper information, from the top level of the Administration Console, enter:

```
ip udphelper display
```

The following example shows a UDP Helper display:

```

BOOTP relay hop count limit is 4,BOOTP relay threshold is 0.
BOOTP is evenly distributing interfaces

```

```

UDP port           Forwarding address
   67                158.101.1.112

```

Defining a Port and IP Routing Address

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

Top-Level Menu

```

system
ethernet
  interface
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  route
  arp
  atmAr
  multi
  udph
  routin
  icmpR
  ospf
  rip
  ping
  statistics
  display
  define
  remove
  hopCountLimit
  threshold
  interface

```

To define port numbers and IP routing addresses:

- 1 From the top level of the Administration Console, enter:
ip udphelper define
- 2 Enter the port numbers and IP routing addresses that you want to define.

Removing a Port Number and IP Routing Address

To remove a port number and IP routing address that is defined for UDP Helper:

Top-Level Menu

```

system
ethernet
  interface
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  route
  arp
  atmAr
  multi
  udph
  routin
  icmpR
  ospf
  rip
  ping
  statistics
  display
  define
  remove
  hopCountLimit
  threshold
  interface

```

- 1 From the top level of the Administration Console, enter:
ip udphelper remove
- 2 Enter the UDP port number that you want to remove.
- 3 Enter the IP routing address that you want to remove.

The port numbers and IP routing addresses are immediately removed.

Setting the BOOTP Hop Count Limit


You can set the maximum *hop count*, that is, the maximum number of times that packets will be forwarded within the router. The range is 0 through 16 hops.

Default The default hop count setting is 4.

To set the hop count limit:

- 1 From the top level of the Administration Console, enter:
ip udpHelper hopCountLimit
- 2 Enter the BOOTP relay hop count limit.

Top-Level Menu




Setting the BOOTP Relay Threshold

You can set the maximum time limit for routing a packet. If you use 0 as the value, the router does not process the value in the seconds field. If you use a nonzero value, the router uses that value along with the hop count value to determine whether to forward a UDP packet.

To set the BOOTP relay threshold:

- 1 From the top level of the Administration Console, enter:
ip udpHelper threshold
- 2 Enter the BOOTP relay threshold.

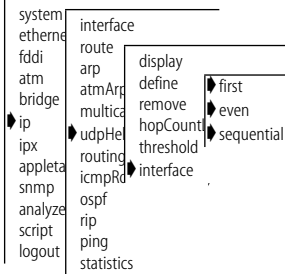
Top-Level Menu



Configuring Overlapped Interfaces

Overlapped IP interfaces define multiple logical interfaces for a single physical port. You can specify how UDP Helper forwards packets from overlapped IP interfaces using the `udpHelper` interface menu option.

Top-Level Menu



To configure UDP Helper to support overlapped IP interfaces:

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

```
ip udpHelper interface
```

2 Enter one of these interface configuration options:

- **first** — The system uses the first overlapped IP interface as the source network for forwarded packets.
- **even** — The system hashes the client's MAC address to determine the source network for forwarded packets. This procedure evenly distributes the interface used among those on the network.
- **sequential** — The system assigns each overlapped IP interface, in turn, as the source network for forwarded packets.

The system implements your selection immediately. You can view the UDP Helper configuration after you configure the forwarding address.

Enabling and Disabling IP Routing

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

Default The default IP routing setting is *disabled* on the system.


To enable or disable IP routing:

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

```
ip routing
```

2 Enter the IP routing state (**enabled** or **disabled**).

Top-Level Menu



Enabling and Disabling ICMP Router Discovery

The Internet Control Message Protocol (ICMP) Router Discovery protocol (RFC 1256) allows an appropriately configured end station to locate one or more routers on the LAN to which it is attached. The end station then installs a default route to each of the routers that is running ICMP Router Discovery. You do not need to manually configure a default route. While IP traffic may initially be directed to any of the routers on the LAN, ICMP redirect messages subsequently channel IP traffic to the correct router.

Only certain end stations, such as Solaris workstations, can be configured to work with the ICMP Router Discovery protocol. Refer to the documentation for your workstation to determine whether you can configure it to work with this protocol.

To enable ICMP Router Discovery, from the top level of the Administration Console, enter:

```
ip icmpRouterDiscovery
```

Enter the ICMP router discovery mode (**enabled** or **disabled**).

Default This protocol is *disabled* by default.

Top-Level Menu

```
system
ethernet interface
fddi route
atm arp
bridge atmArpServer
ip multicast
ipx udphelper
appletalk routing
snmp icmpRouterDiscovery
analyze ospf
script rip
logout ping
statistics
```

Configuring RIP

With Extended Switching software you can configure IP Routing Information Protocol (RIP) on individual interfaces, rather than turning it on or off for all interfaces.

Displaying the RIP Interface Configuration

To display RIP interface configuration information, from the top level of the Administration Console, enter:

```
ip rip display
```

Example RIP display:

IP routing is enabled, ICMP router discovery is disabled

RIP interface information:

Index	Mode	Cost	PoisonReverse	AdvertisementAddress
1	active	1	enabled	158.101.112.255
2	off	1	disabled	158.101.142.255

Top-Level Menu

```
system
ethernet interface
fddi route
atm arp
bridge atmArpServer
ip multicast
ipx udphelper
appletalk routing
snmp icmpR
analyze ospf
script rip
logout ping
statistics
```

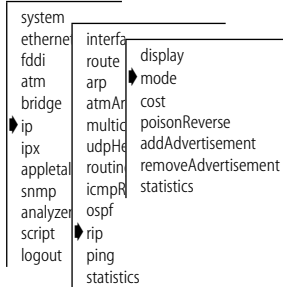
Setting the RIP Mode

You can select an RIP mode that is appropriate for each interface. RIP can operate in any of three modes:

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

Default The default RIP mode is *passive*.

Top-Level Menu



To set the RIP mode:

- 1 From the top level of the Administration Console, enter:
ip rip mode
- 2 Enter the RIP mode (**off**, **passive**, or **active**).
- 3 To accept the existing value in brackets, press Return at the prompt.

Example:

```
Select interface(s) (1-2|all):2
Enter RIP mode (off,passive,active): active
```

Setting the RIP Interface Cost

The system uses the RIP interface cost (a number between 1 and 15) to calculate route metrics. Unless your network has special requirements, assign a cost of 1 to all interfaces.

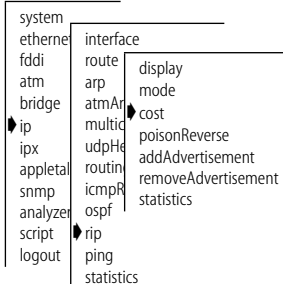
To set the RIP interface cost:

- 1 From the top level of the Administration Console, enter:
ip rip cost
- 2 Enter the cost value.

Example:

```
Select interface(s) (1-2|all): 2
Enter RIP cost [1]:
```

Top-Level Menu



Setting the Poison Reverse Mode

The poison reverse mode allows you to prevent a routing loop if a route is advertised on the interface on which it came in. Poison reverse has two modes:

- **enabled** — On the interface on which a route came in, RIP advertises that the route is unavailable. Because the originating router does not overwrite its routing table entry, a routing loop does not occur.
- **disabled** — RIP does not advertise that a route is unavailable.

To set the poison reverse mode, from the top level of the Administration Console, enter:

```
ip rip poisonReverse
```

Example:

```
Select interfaces (1-2|all):2
```

```
Enter RIP poison reverse mode (disabled,enabled) [enabled]:
```

Default The default poison reverse mode is *enabled*.

Adding an RIP Advertisement Address

The system uses this IP address when it advertises routes to other stations on the same subnetwork. In particular, the system uses this address for sending RIP updates.

Default By default, the system uses a directed advertisement (all **1s** in the host field).

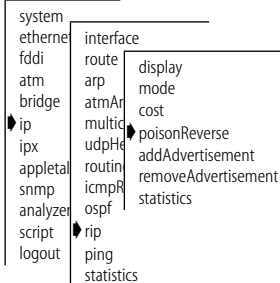
To add an RIP advertisement address to an IP interface:

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

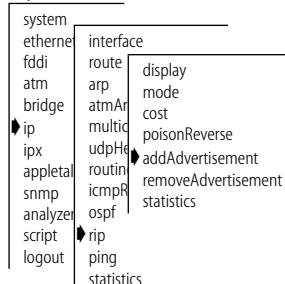
```
ip rip addAdvertisement
```

- 2 Enter the RIP advertisement addresses that you want to add.

Top-Level Menu



Top-Level Menu



Removing an RIP Advertisement Address

You can remove an RIP advertisement address from an IP interface.

Top-Level Menu

```

system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  interface
    route
    arp
    atmAr
    multic
    udph
    routin
    icmpR
    ospf
    rip
    ping
    statistics
    display
    mode
    cost
    poisonReverse
    addAdvertisement
    removeAdvertisement
    statistics
  
```

- 1 From the top level of the Administration Console, enter:
ip rip removeAdvertisement
- 2 Enter the RIP advertisement addresses that you want to remove.

Displaying RIP General Statistics

To display statistics about RIP activity, from the top level of the Administration Console, enter:

Top-Level Menu

```

system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  interface
    route
    arp
    atmAr
    multic
    udph
    routin
    icmpR
    ospf
    rip
    ping
    statistics
    display
    mode
    cost
    poisonReverse
    addAdvertisement
    removeAdvertisement
    statistics
  
```

ip rip statistics

The following sample shows an RIP statistics display:

```

RIP general statistics
                routeChanges                queries
                225                          23
  
```

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 that you specify. Ping then waits for an ICMP echo reply packet. Possible responses from ping:

- 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 that you want to ping.

```
IP Address: 192.9.200.40
```

You may receive one of the following types of responses:

```
192.9.200.40 is alive
```

OR

```
no answer from 192.9.200.40
```

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

```
Network is unreachable
```

Top-Level Menu

```
system
ethernet interface
fddi route
atm arp
bridge atmArpServer
ip multicast
ipx udphelper
appleta routing
snmp icmpRouterDiscovery
analyze ospf
script rip
logout ping
statistics
```

Displaying IP Statistics

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

```
ip statistics
```

The examples shows statistics displays for IP, UDP, and ICMP general statistics: The accompanying tables describe the fields.

Top-Level Menu

```
system
ethernet
fddi
atm
bridge
ip
ipx
appletalk
snmp
analyze
script
logout
  interface
  route
  arp
  atmArpServer
  multicast
  udphelper
  routing
  icmpRouterDiscovery
  ospf
  rip
  ping
  statistics
```

IP general statistics

inReceived	2327812806	inHdrErrors	0	inAddrErrors	0
forwDatagrams	2327800881	unkProtos	0	inDiscards	0
inDelivers	17933	outRequests	50005	outDiscards	0
outNoRoutes	0	reasmReqs	0	reasmOks	0
reasmFails	0	fragOks	0	fragFails	0
fragCreates	0	osReceives	0	osTransmits	0
rtDiscards	0				

UDP general statistics

inDatagrams	noPorts	inErrors
14431	0	0
outDatagrams		
43266		

ICMP general statistics

messages	inErrors	inDestUnreach
81	0	81
inTimeExcds	inParmProbs	inSrcQuenchs
0	0	0
inRedirects	inEchos	inEchoReps
0	0	0
inTimeStamps	inTimeStampsReps	inAddrMasks
0	0	0
inAddrMaskReps	outMsgs	outErrors
0	128	0
outDestUnreach	outTimeExcds	outParmProbs
128	0	0
outSrcQuenchs	outRedirects	outEchos
0	0	0
outEchoReps	outTimeStamps	outTimeStampReps
0	0	0
outAddrMasks	outAddrMaskReps	
0	0	

Table 10-2 describes the general IP statistics.

Table 10-2 IP Statistics

Field	Description
forwDatagrams	Number of datagrams that the IP station tried to forward
fragCreates	Number of IP datagram fragments that were generated as a result of fragmentation on this system
fragFails	Number of IP datagrams discarded because they needed to be fragmented but could not be (because, for example, their Don't Fragment bit was set)
fragOks	Number of IP datagrams that were successfully fragmented
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
inDiscards	Number of packet receive discards
inHdrErrors	Number of datagrams that the IP station discarded because the IP header contained errors
inReceived	Total number of IP datagrams that were received, including those with errors
osReceives	Number of packets received that are destined to higher-level protocols such as telnet, DNS, TFTP, and FTP
osTransmits	Number of packets that were sent through the router by higher-level protocols such as telnet, DNS, TFTP, and FTP
outDiscards	Number of packet transmit discards
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
reasmFails	Number of packet reassembly failures
reasmReqs	Number of packet reassembly requests
reasmOks	Number of successful packet reassemblies
rtDiscards	Number of packets that were discarded because of system resource errors
unkProtos	Number of packets whose protocol is unknown

Table 10-3 describes the UDP statistics.

Table 10-3 UDP Statistics

Field	Description
inDatagrams	Number of UDP packets received and addressed to the router or broadcast address
inErrors	Number of received UDP or ICMP packets that contain header errors
noPorts	Number of UDP packets received but addressed to an unsupported UDP port
outDatagrams	Number of UDP packets sent by the router

Table 10-4 describes the ICMP statistics.

Table 10-4 ICMP Statistics

Field	Description
inAddrMaskReps	Number of ICMP address mask reply frames received
inAddrMasks	Number of ICMP address mask request packets received
inDestUnreach	Number of ICMP destination unreachable packets received
inEchoReps	Number of ICMP echo reply packets received
inEchos	Number of ICMP echo request packets received
inParmProbs	Number of ICMP parameter problem frames received
inRedirects	Number of ICMP redirect packets received
inSrcQuenchs	Number of ICMP source quench packets received
inTimeExcds	Number of ICMP time exceeded packets received
inTimeStamps	Number of ICMP time stamp request packets received
inTimeStampsReps	Number of ICMP time stamp reply packets received
messages	Number of ICMP packets received
outAddrMaskReps	Number of ICMP address mask reply packets sent
outAddrMasks	Number of ICMP address mask request packets sent
outDatagrams	Number of UDP packets sent by the router
outDestUnreach	Number of ICMP destination unreachable packets sent
outEchoReps	Number of ICMP echo reply packets sent
outEchos	Number of ICMP echo request packets sent
outErrors	Number of ICMP packets that were sent and then dropped because of system resource errors

(continued)

Table 10-4 ICMP Statistics (continued)

Field	Description
outMsgs	Number of ICMP packets that were sent
outParmProbs	Number of ICMP parameter problem packets that were sent
outRedirects	Number of ICMP redirect packets that were sent
outSrcQuenchs	Number of ICMP source quench packets that were sent
outTimeExcds	Number of ICMP time exceeded packets that were sent
outTimeStampReps	Number of ICMP time stamp reply packets that were sent
outTimeStamps	Number of ICMP time stamp request packets that were sent

11

ADMINISTERING IP MULTICAST ROUTING

This chapter describes how to set up your CoreBuilder™ 2500 system to use IP multicast routing, including information about these tasks:

- Enabling and Disabling DVMRP
- Enabling and Disabling IGMP
- Administering IP Multicast Interfaces
- Administering Multicast Tunnels
- Displaying Routes
- Displaying the Multicast Cache

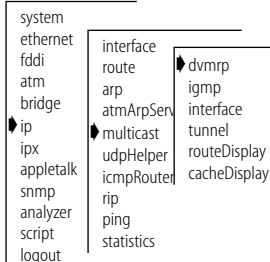


Before you define any IP multicast interfaces, you must first define IP interfaces and routes, as described in Chapter 10.

For more information about how IP multicast routing works, see Chapter 5, “Routing with IP Multicast.”

Enabling and Disabling DVMRP

Top-Level Menu



Default

DVMRP (Distance Vector Multicast Routing Protocol) is similar to the IP Routing Information Protocol (RIP). Multicast routers exchange distance vector updates that contain lists of destinations and the distance in hops to each destination. They maintain this information in a routing table.

To run multicast routing, you enable DVMRP, which enables it on all IP interfaces that have not been disabled. To enable or disable DVMRP, from the top level of the Administration Console, enter:

```
ip multicast dvmrp
```

The interface prompts you to enable or disable DVMRP.

The default is *disabled*.

Example:

```
Enter DVMRP mode (disabled, enabled) [disabled]: enabled
```

Enabling and Disabling IGMP

The Internet Group Membership Protocol (IGMP) enables a router or switch to determine whether group members exist in a subnetwork, or *leaf*, of the Spanning Tree. The protocol uses two search methods to make this determination:

- **Query mode** — The router or switch with the lowest IP address in the LAN broadcasts a query to all other members of the subnetwork to determine whether they are also in the group. End stations respond to the query with IGMP packets, which report the multicast group to which they belong.
- **Snooping mode** — A router or switch performs dynamic multicast filtering based on IGMP snooping, which ensures that multicast packets flood only to the appropriate ports within a routing interface.

When you select the IGMP option, the interface prompts you to enable or disable IGMP snooping mode and IGMP query mode. Both are *enabled* by default. Under most conditions, IGMP snooping mode and IGMP query mode should remain enabled.

To enable or disable IGMP, from the top level of the Administration Console, enter:

```
ip multicast igmp
```

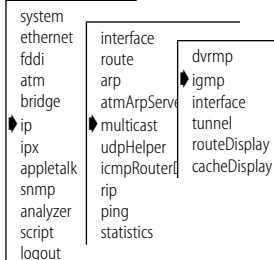
The interface prompts you to enable or disable IGMP snooping and query modes.

Example:

```
Enter IGMP snooping mode (disabled, enabled) [enabled]:  
enabled
```

```
Enter IGMP query mode (disabled, enabled) [enabled]:  
enabled
```

Top-Level Menu



Administering IP Multicast Interfaces

With the IP multicast interface selections, you can enable and disable multicast characteristics on previously defined IP interfaces.

Multicast Interface Characteristics


A multicast interface has three characteristics:

- **DVMRP Metric Value** — Determines the cost of a multicast interface. Use a higher cost to decrease the likelihood that multicast packets will be routed over the interface. The default value is *1*.
- **Time To Live (TTL) Threshold** — Determines whether the interface forwards multicast packets to other switches and routers in the subnetwork. If you do not want the interface to forward packets, use an interface TTL greater than the packet TTL. If you want the interface to forward all packets, use the default value of *1*.
- **Rate Limit** — Determines how fast multicast traffic can travel over the interface in packets per second. If multicast traffic exceeds the rate limit, the system drops packets to maintain the set rate. If you want no rate limit, use the default value of *0*. If you want to reduce traffic over the interface, use a low rate limit.

Displaying Multicast Interfaces

To display a multicast interface:

Top-Level Menu



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

```
ip multicast interface display
```

- 2 Enter the index numbers of the interfaces to display.

This example shows a multicast interface configuration:

```
DVRMP is disabled, IGMP snooping is enabled
```

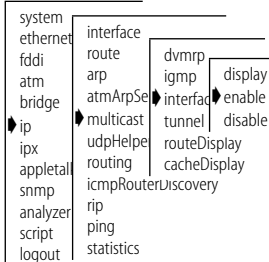
Index	Local Address	Metric	Threshold	RateLimit	State
1	158.101.112.32	1 pkts	1 in:0	0 pkts out:0	querier leaf

Enabling Multicast Interfaces

Multicast routing is enabled on all existing IP interfaces unless you have specifically disabled it. Use the option described in this section to change the characteristics of an existing interface or to enable an interface that you previously disabled.

To enable a multicast interface or to modify the multicast characteristics of an existing IP interface:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:
ip multicast interface enable
- 2 Enter the index numbers of the interfaces that you want to enable.
- 3 Enter the DVMRP metric value of the chosen interfaces.
- 4 Enter the Time To Live (TTL) threshold of the chosen interfaces.
- 5 Enter the rate limit of the chosen interfaces.

Example:

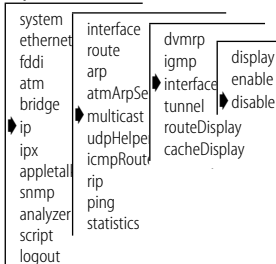
```

Enter an IP interface index [1]: 2
Enter Interface DVMRP metric [1]: 1
Enter Interface TTL threshold [1]:
Enter interface rate limit in KBits/sec [0]:
  
```

Disabling Multicast Interfaces

To disable multicast routing on an interface:

Top-Level Menu



- 1 From the top level of the Administration Console, enter:
ip multicast interface disable
- 2 Enter the index number of the IP interface to disable.
Enter an IP interface index {1-2}: 1

The interface is disabled.

Administering Multicast Tunnels

A multicast tunnel allows multicast packets to cross several unicast routers to a destination router that supports multicast routing. A tunnel has two end points. The local end point is associated with an interface on the CoreBuilder 2500 router.

When you define the tunnel, specify the associated interface on the local CoreBuilder 2500 router, specify the IP address of the remote multicast router, and then specify the characteristics of the tunnel. Tunnel characteristics are the same as those of an interface. Also specify the IP address of the remote multicast router.

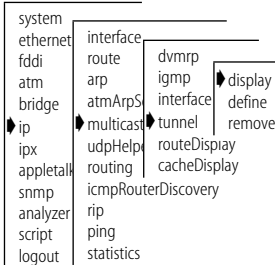


Not all multicast configurations require a tunnel. The only configurations that require a tunnel are those that require a connection between two multicast internetworks through one or more unicast routers.

Displaying Multicast Tunnels

To display the IP multicast tunnels on the router, from the top level menu of the Administration Console, enter:

Top-Level Menu



ip multicast tunnel display

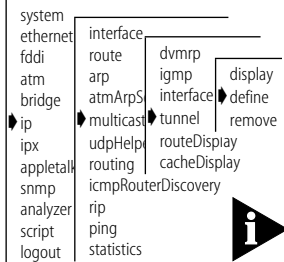
This example shows an IP multicast tunnel configuration:

Index	Local Address	Remote Address	Metric	Threshold	RateLimit	State
1	158.101.112.204	137.39.229.98	2	255	500	
		pkts in:320069	pkts out:0			
		peers 137.39.229.98	(3.8)	(0xe)		

Defining a Multicast Tunnel

To define a multicast tunnel:

Top-Level Menu



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

```
ip multicast tunnel define
```

- 2 Enter the index numbers of the interfaces that you want to associate with a multicast tunnel.

- 3 Enter the IP address of the destination multicast router.

The IP address of the destination multicast router must be a remote address. The destination router cannot be directly connected to the same subnetworks as the local IP address.

- 4 Enter the DVMRP metric value of the tunnel.

- 5 Enter the Time To Live (TTL) threshold of the tunnel.

- 6 Enter the rate limit of the tunnel.

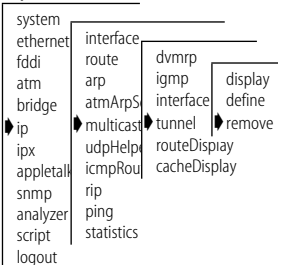
Example:

```
Enter an IP interface index [1]: 2
Enter remote IP address: 192.9.200.40
Enter tunnel DVMRP metric [1]: 1
Enter tunnel TTL threshold [1]:
Enter tunnel rate limit [0]:
```

Removing a Multicast Tunnel

To remove an IP multicast tunnel:

Top-Level Menu



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

```
ip multicast tunnel remove
```

- 2 Enter the index numbers of the interfaces that are associated with the tunnel that you want to remove.


```
Enter multicast tunnel index [1]: 1
```

The tunnel is removed.

Displaying Routes

To display all available routes in the IP multicast routing table:

Top-Level Menu



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

ip multicast routeDisplay

- 2 The system displays the DVMRP and IGMP status:

DVMRP is disabled, IGMP snooping is enabled

This example route display shows available multicast routes.

Multicast Routing Table (2598 entries)

Origin-Subnet	From-Gateway	Metric	Tmr	In-If	Out-Ifs
157.88.29.1/32	137.39.229.98	18	25	T1	I1
137.39.2.254/32	137.39.229.98	5	25	T1	I1
131.215.125.236/32	137.39.229.98	14	25	T1	I1
130.118.106.254/32	137.39.229.98	10	25	T1	I1
129.127.118.12/32	137.39.229.98	10	25	T1	I1
129.127.110.12/32	137.39.229.98	10	25	T1	I1
129.127.110.11/32	137.39.229.98	13	25	T1	I1
129.127.110.5/32	137.39.229.98	10	25	T1	I1
129.95.63.12/32	137.39.229.98	13	25	T1	I1
129.95.63.11/32	137.39.229.98	31	25	T1	I1*
129.95.63.9/32	137.39.229.98	13	25	T1	I1
129.95.63.8/32	137.39.229.98	13	25	T1	I1
129.95.63.6/32	137.39.229.98	13	25	T1	I1
129.95.63.2/32	137.39.229.98	13	25	T1	I1
129.95.48.4/32	137.39.229.98	13	25	T1	I1
129.95.48.3/32	137.39.229.98	13	25	T1	I1
129.95.48.2/32	137.39.229.98	13	25	T1	I1

Table 11-1 describes the fields in the route display.


Table 11-1 Field Attributes for Multicast Route Display

Field	Description
Origin-Subnet	Source address and the number of bits in the subnetwork
From-Gateway	Interface address of the gateway
Metric	Hop count
Tmr	Time, in seconds, since the routing table entry was last reset
In-If*	Interface number on which that gateway is connected. Traffic is expected to originate from this interface. T represents the tunnel; P denotes that a prune has been sent to this tunnel.
Out-Ifs*	The set of interfaces out of which the traffic will be flooded. Ix represents the interface.

*In-If and Out-Ifs Together, these attributes define a Spanning Tree configuration. The system disables interfaces that comprise loops.

Displaying the Multicast Cache

Top-Level Menu



The multicast cache contains the IP source address and destination address for packets observed on the system. The multicast cache shows how information is routed over interfaces and ports in your system.

To display all learned routes in the multicast cache:

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

```
ip multicast cacheDisplay
```

2 At the prompt, enter the multicast source address.

To display all multicast traffic, enter 255.255.255.0.

3 At the prompt, enter the multicast group address.

To display all multicast traffic, enter 255.255.255.0.

The DVMRP status and IGMP status appear on the screen.

Example:

```
Enter multicast source address [131.188.0.0]:
Enter multicast group address [244.2.0.1]:
```

DVMRP is enabled, IGMP snooping is enabled

This example shows a multicast cache configuration:

Multicast Routing Cache Table (125 entries)

Origin	Mcast-group	CTmr	Age	PTmr	In-If	Out-Ifs
>202.242.133.128/26	224.2.0.1	7m	11m	6m	T1P	I1p
202.242.133.139	2 packets					
>128.84.247/24	224.2.0.1	2m	36m	2m	T1P	I1p
128.84.247.53	43 packets					
128.84.247.156	33 packets					
>128.138.213/24	224.2.0.1	3m	2h	2m	T1P	I1p
128.138.213.1	23 packets					
>128.206.212/24	224.2.0.1	92s	36m	60s	T1P	I1p
128.206.212.69	8 packets					
>131.136.234/24	224.2.0.1	3m	57m	3m	T1P	I1p
131.136.234.103	12 packets					
>138.39.25/24	224.2.0.1	103s	4h	71s	T1P	I1p
138.39.25.48	46 packets					
>192.5.28/24	224.2.0.1	80s	2h	48s	T1P	I1p
192.5.28.43	178 packets					
>199.94.220/24	224.2.0.1	104s	1h	72s	T1P	I1p
199.94.220.184	10 packets					
>199.104.80/24	224.2.0.1	3m	32m	3m	T1P	I1p
199.104.80.5	4 packets					

Table 11-2 describes the fields in the cache configuration display.

Table 11-2 Field Attributes for the Cache Configuration Display

Field	Description
Origin	The source of the incoming packets. Entries preceded by an angle bracket (>) indicate a multicast subnetwork. Entries without an angle bracket are multicast routers within the subnetwork immediately preceding them in the table.
Mcast-group	The destination multicast group
CTmr	Cache timer, which is the amount of time that a cache entry remains in the cache
Age	The number of seconds (s), minutes (m), or hours (h) that the cache entry has existed
PTmr	The time remaining, in seconds (s), minutes (m), or hours (h), before sending another prune to prune the Spanning Tree
In-If	The interface number where that gateway is connected. Traffic is expected to originate from this interface. T represents the tunnel; P denotes that a prune has been sent to this tunnel.
Out-Ifs	The set of interfaces out of which the traffic will be flooded. Ix represents the interface.

12

ADMINISTERING IPX ROUTING

This chapter describes how to set up your CoreBuilder™ 2500 system to use the Internet Packet Exchange (IPX) protocol to route packets. The chapter discusses these tasks:

- Administering Interfaces
- Administering Routes
- Administering Servers
- Setting IPX Forwarding
- Setting the RIP Mode
- Setting the Enhanced RIP Mode
- Setting RIP Triggered Updates
- Setting the SAP Mode
- Setting SAP Triggered Updates
- Displaying Statistics

For information about how IPX works, see Chapter 6, “Routing with IPX.”

Administering Interfaces

An IPX interface defines the relationship between an IPX Virtual LAN (VLAN) and the IPX network. Every IPX interface has one IPX VLAN associated with it. Each switching module has one IPX interface defined for each subnetwork directly connected to it.



Before you define an associated interface, define a VLAN, as described in Chapter 9, “Administering VLANs.”

An IPX interface has the following information associated with it:

- **IPX network address** — You set this 4-byte address. Make each address unique within the network.
- **Cost** — The system uses the cost, a number between 1 and 15, to calculate route metrics. Unless your network has special requirements, such as the need for redundant paths, assign a cost of 1 to each interface.
- **Frame format** — IPX routing uses four Ethernet frame formats and two FDDI frame formats. The Ethernet frame formats are Ethernet Type II, Novell 802.3 raw, 802.2 LLC, and 802.3 SNAP. The FDDI frame formats are FDDI 802.2 and FDDI SNAP.

The two FDDI frame formats correspond to the Ethernet 802.2 LLC and 802.3 SNAP frame formats. If you select either of these Ethernet frame formats, the corresponding FDDI frame format is automatically selected for shared Ethernet and FDDI ports.

- **State** — The status of the IPX interface indicates whether the interface is available for communications (`UP`) or unavailable (`DOWN`).
- **VLAN index** — The VLAN index indicates the bridge ports that are associated with the IPX interface. When the interface prompts you for this option, a list of available VLAN indexes and the ports associated with them appears.

Displaying IPX Interfaces

Use this command to display a table that shows all IPX interfaces and their parameter settings currently configured for the system.

Top-Level Menu

```

system
ethernet
fdi
atm
bridge
ip
ipx
appletalk
snmp
analyzer
script
logout
  
```

To display IPX interface information, from the top level of the Administration Console, enter:

```
ipx interface display
```

Example with IPX interface, IPX forwarding, RIP, and SAP information:

```
IPX forwarding is enabled, RIP is active, SAP is active,
RIP Trig enabled, SAP Trig enabled
```

Index	IPX address	Cost	Format	State	VLAN index
1	45469f30	1	802.2	Up	2
2	5d41a110	1	802.2	Up	1
3	6d321a22	1	802.2	Up	4

Defining an Interface

When you define an interface, you define the interface's IPX address, cost, format, and any associated IPX VLAN index.



Before you define the IPX interface to associate with a VLAN, define the IPX VLAN. See Chapter 9, "Administering VLANs."

To define an IPX interface:

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

```
ipx interface define
```

You are prompted for the interface parameters. Press Return to accept the existing values, shown in brackets.

- 2 Enter the IPX network address of the interface.
- 3 Enter the cost of the interface.
- 4 Enter the frame format of the interface.
- 5 Enter the index of the IPX VLAN that is associated with this interface.

Example:

```
Enter IPX Address: 0x45469f30
```

```
Enter Cost [1]: 1
```

```
Enter Frame Format (Ethernet II, 802.2, Raw 802.3, SNAP): 802.2
```

```
IPX VLANs:
```

Index	Ports
3	1-8
4	9-12

```
Select VLAN index: 3
```

Top-Level Menu


```

system
ethernet
fdi
atm
bridge
ip
ipx
appletalk
snmp
analyzer
script
logout
  
```

Modifying an Interface

To change the configuration of a defined IPX interface:

Top-Level Menu



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

```
ipx interface modify
```


- 2 At the prompt for each interface parameter, enter the new values for parameters that you want to modify.

Press Return to accept the existing values in brackets for parameters that you do not want to modify.

Removing an Interface

Remove an interface when you no longer perform routing on the ports that are associated with the interface.

Top-Level Menu



To remove an IPX interface definition:

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

```
ipx interface remove
```

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

The interface is removed.

Administering Routes


Your system 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 in the router's routing table.
- **Hops** — The number of routers that must be crossed to reach the network segment. The maximum number of routers a packet can cross is 15. Exception: An IPX NetBIOS packet can cross no more than 7 routers.
- **Tics** — An estimate of how long the packet will take 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** — The number of seconds that have elapsed since the last time the router sent a packet.

Displaying the Routing Table

You can display the routing tables for the system, determine which routes are configured, and find out whether they are operational.

Top-Level Menu



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

ipx route display

This example shows IPX forwarding, RIP, SAP, and routing information.

```
IPX forwarding is enabled, RIP is active, SAP is active,
RIP Trig enabled, SAP Trig enabled
```

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

Before you define static routes on the system, define at least one IPX interface. See "Defining an Interface" on page 12-3 for more details. Static routes remain in the routing 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 have a maximum of 16 static routes.


To define a static route:

- 1 From the top level of the Administration Console, enter:
- ipx route static**
- 2 Enter the 4-byte IPX network address of the route.
 - 3 Enter the cost of the route.
 - 4 Enter the interface number of the route.
 - 5 Enter the node address of the route.

Example:

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


Top-Level Menu



Removing a Route

To remove a route from the IPX routing table:

Top-Level Menu



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

```
ipx route remove
```


- 2 Enter the 4-byte IPX network address.

The route is immediately deleted.

Flushing Routes

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

Top-Level Menu



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

```
ipx route flush
```

All learned routes are immediately deleted.

Administering Servers

Your system maintains a table of servers that reside on other IPX networks. Either use the Service Advertising Protocol (SAP) to exchange server information automatically or make static entries in this server 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 forwards packets to the segment.
- **Socket** — The 2-byte socket address of the server that receives service requests.
- **Hop** — The number of networks that must be crossed to reach the server. The maximum number is 15.
- **Age** — The number of seconds that have elapsed since the last time that a server in the table sent a packet.

Displaying the Server Table

You can display the system server table to determine which servers are learned and whether they are operational.

Top-Level Menu

```

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

```

interface
route
server
forwarding
rip
sap
statistics
  
```

```

display
static
remove
flush
  
```

To display the server table, from the top level of the Administration Console, enter:

```
ipx server display
```

This example shows information about known servers, IPX forwarding, RIP, and SAP.

```

IPX forwarding is enabled, RIP is active, SAP is active,
RIP Trig enabled, SAP Trig enabled
  
```

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

Before you define static servers on the system, you must define at least one IPX interface. See "Defining an Interface" on page 12-3 for more details. 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 8 static servers. To define a static server:

Top-Level Menu

```

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

```

interface
route
server
forwarding
rip
sap
statistics
  
```

```

display
static
remove
flush
  
```

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

```
ipx server static
```

- 2 Enter the interface number of the server.

- 3 Enter the service type of the server.

- 4 Enter the service name of the server.

- 5 Enter the IPX network address of the server.

- 6 Enter the socket value of the server.

- 7 Enter the node address of the server.

- 8 Enter the number of hops to the server.

Example:


```

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 from the IPX server table:

Top-Level Menu



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


ipx server remove

- 2 Enter the service type of the server.
- 3 Enter the service name of the server.
The server is immediately deleted.

Flushing Servers

Flushing deletes all dynamically learned servers from the server table.

Top-Level Menu



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

ipx server flush

All learned servers are immediately deleted.


Setting IPX Forwarding

Use this command to control whether the system forwards or discards IPX packets addressed to other routers. When you enable IPX forwarding, the system acts as a normal IPX router, forwarding IPX packets from one network to another when required. When you disable IPX forwarding, the system discards any IPX packets that are not addressed directly to one of its defined IPX interfaces.

Default The default IPX forwarding setting is *disabled*.

To enable or disable IPX forwarding:

Top-Level Menu



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

ipx forwarding

- 2 Enter the IPX forwarding state (**enabled** or **disabled**). Press Return to accept the existing value in brackets.

Setting the RIP Mode

Use this command to select an RIP mode that is appropriate for your network. RIP can operate in any of three modes:

- **off** — The IPX module does not process any incoming RIP packets and does not generate any RIP packets of its own.
- **passive** — The system processes all incoming RIP packets, but it does not broadcast periodic or triggered RIP updates or respond to RIP requests.
- **active** — The system processes all incoming RIP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered RIP updates.

Default The default RIP mode is *off*.

To set the RIP mode:


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

```
ipx rip mode
```

- 2 Enter the RIP mode (**off**, **passive**, or **active**).

Press Return to accept the existing value in brackets.

Top-Level Menu



Setting the Enhanced RIP Mode

Standard IPX RIP packets include up to 50 route advertisements, but some routers allow up to 68 route advertisements. Enhanced RIP mode increases the number of advertisements in an RIP packet that the system accepts. Enhanced RIP mode gives the system greater interoperability with routers that do not follow standard IPX advertisement guidelines.

Default The default enhanced RIP mode is *disabled*.

To enable or disable enhanced RIP mode:


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

```
ipx rip enhanced
```

- 2 Enter the enhanced RIP state (**enabled** or **disabled**).

Press Return to accept the existing value in brackets.

Top-Level Menu



Setting RIP Triggered Updates

The `ipx rip triggered` command directs the IPX protocol when to broadcast newly learned routes.

Default The default setting is *enabled*.

To configure IPX RIP broadcast timing:

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

```
ipx rip triggered
```


- 2 Enter the triggered update mode:

- **disabled** — IPX broadcasts routes 3 seconds after learning them.
- **enabled** — IPX broadcasts routes immediately after learning them.

Example:

```
Enter Trigger Update mode for RIP (disabled,enabled) [enabled]:
```

Top-Level Menu



Setting the SAP Mode

Use this command to select a SAP mode that is appropriate for your network. SAP can operate in any of these modes:

- **off** — The system does not process any incoming SAP packets and does not generate any SAP packets of its own.
- **passive** — The system processes all incoming SAP packets, but it does not broadcast periodic or triggered SAP updates or respond to SAP requests.
- **active** — The system processes all incoming SAP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered SAP updates.

Default The default mode is *off*.

To set the SAP mode:


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

```
ipx sap mode
```

- 2 Enter the SAP mode (**off**, **passive**, or **active**).

Press Return to accept the existing value in brackets.

Top-Level Menu



Setting SAP Triggered Updates

The `ipx sap triggered` command directs the IPX protocol when to broadcast newly learned SAP server addresses.

Default The default setting is *enabled*.

To configure IPX SAP broadcast timing:

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

`ipx sap triggered`


2 Enter the triggered update mode:

- **disabled** — IPX broadcasts SAP server addresses 3 seconds after learning them.
- **enabled** — IPX broadcasts SAP server addresses immediately after learning them.

Example:

Enter Trigger Update mode for SAP (disabled,enabled) [enabled]:

Top-Level Menu



Displaying Statistics

Use this command to display IPX summary, RIP, SAP, and forwarding statistics from the Administration Console.

Displaying IPX Summary Statistics

To display IPX summary statistics, from the top level of the Administration Console, enter:

`ipx statistics summary`

This example shows IPX summary statistics:

IPX forwarding is enabled, RIP is active, SAP is active,
RIP Trig enabled, SAP Trig enabled

Received	Transmitted	Dropped	Msg Pool	Empty
1170878	565099	0		0

Top-Level Menu




Table 12-1 describes the IPX summary statistics.

Table 12-1 IPX Summary Statistics

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

Displaying IPX RIP Statistics

To display IPX RIP statistics, from the top level of the Administration Console, enter:

```
ipx statistics rip
```

This example shows IPX RIP statistics:

```
IPX forwarding is enabled, RIP is active, SAP is active,
RIP Trig enabled, SAP Trig enabled
```

RIP Received	RIP Transmitted	RIP dropped
106195	7929	0
RIP Responses	RIP Requests	RIP Entries
100552	5643	2

Table 12-2 describes the IPX RIP statistics.

Table 12-2 IPX RIP Statistics

Field	Description
RIP Received	Number of IPX RIP packets that were received
RIP Transmitted	Number of IPX RIP packets that were transmitted
RIP Dropped	Number of IPX RIP packets that were 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

Top-Level Menu

system	interface	
ethernet	route	summary
fdi	server	rip
atm	forwarding	sap
bridge	rip	forwarding
ip	sap	
ipx	statistics	
appletalk		
snmp		
analyzer		
script		
logout		

Displaying IPX SAP Statistics

To display IPX SAP statistics, from the top level of the Administration Console, enter:

```
ipx statistics sap
```

This example shows IPX SAP statistics:

Top-Level Menu

system		
ethernet	interface	
fdi	route	
atm	server	summary
bridge	forwarding	rip
ip	rip	sap
ipx	sap	forwarding
appletalk	statistics	
snmp		
analyzer		
script		
logout		

```
IPX forwarding is enabled, RIP is active, SAP is active,
RIP Trig enabled, SAP Trig enabled
```

```

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
```

Table 12-3 describes the IPX SAP statistics.

Table 12-3 IPX SAP Statistics

Field	Description
SAP Received	Number of IPX SAP packets that were received
SAP Transmitted	Number of IPX SAP packets that were transmitted
SAP Dropped	Number of IPX SAP packets that were 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 that were processed

Displaying IPX Forwarding Statistics

To display IPX forwarding statistics, from the top level of the Administration Console, enter:

Top-Level Menu

```

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

```

interface
route
server
forwarding
rip
sap
statistics
  
```

```

summary
rip
sap
forwarding
  
```

ipx statistics forwarding

This example shows IPX forwarding statistics:

```

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
0                        0
  
```

Table 12-4 describes the IPX forwarding statistics.

Table 12-4 IPX Forwarding Statistics

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

13

ADMINISTERING OSPF ROUTING

This chapter describes how to set up your CoreBuilder™ 2500 system to use the Open Shortest Path First (OSPF) protocol to route packets.

Use this chapter to perform the following tasks:

- Administering Areas
- Setting the Default Route Metric
- Configuring OSPF Interfaces
- Displaying the Link State Database
- Administering Neighbors
- Setting OSPF Router IDs
- Administering Memory Partitions
- Administering Stub Default Metrics
- Administering Virtual Links
- Displaying OSPF Statistics

For more information about how OSPF works, see Chapter 7 “Routing with OSPF.”

Administering Areas

An *OSPF area* is a logical, user-defined group of networks, hosts, and directly attached routers that use a common copy of the OSPF routing algorithm. Different areas run individual copies of the OSPF algorithm and maintain unique routing databases.

With the commands in this section, you define, remove, and modify areas as well as display existing areas. After you define an area with the `define` command, you add network segments and interfaces to the area with the `area addRange` command and the `interface areaID` command.

Displaying Areas


Use this command to display a list of existing areas according to their area identification (ID) numbers.

To display areas:

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

```
ip ospf areas display
```

Top-Level Menu



This example shows existing areas:

Area definition

Indx	AreaID	IP Address	Mask	Advertise	Stub
1	158.101.0.0				n
2	0.0.0.2	158.102.0.0	255.255.0.0	y	n

Table 13-1 Field Attributes for the Areas Display

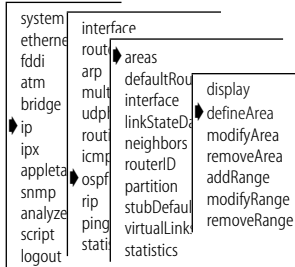
Field	Description
Indx	The entry index for the area
Area ID	The area identifier
IP Address	The network portion of IP address range
Mask	The subnet mask
Advertise	Whether or not the range is advertised
Stub	Whether or not the area is a stub area

Defining Areas

Use this command to define an area. Each OSPF area is a logical group of network entities, including network segments, routers, and nodes. Each area has the following parameters:

- **Area ID** — A number, which is in the form of an IP address, that identifies the area to the OSPF autonomous system.
- **Stub Area** — Indicates whether this is a stub area.

When you define an area, the system assigns it an area index number, based on the area ID. If the area ID is 0.0.0.1, for example, the system assigns a sequential index number of 1 to that area.

Top-Level Menu

- 1 From the top level of the Administration Console, enter:
ip ospf areas defineArea
- 2 Enter the area identification number.
- 3 Specify whether this is a stub area. The default is *no*. Press Return or Enter to accept the existing value in brackets.

Example:

```

Enter Area ID []: 0.0.0.2
Is this a Stub Area (yes,no) [no]:

```

Modifying Areas

Use this command to modify an area range.

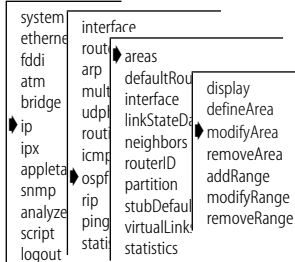
- 1 From the top level of the Administration Console, enter:
ip ospf areas modifyArea
- 2 Enter the index number of the area that you want to modify. Press Return or Enter to accept the existing value in brackets.
- 3 Enter the area ID.
- 4 Specify whether this is a stub area.

Example:

```

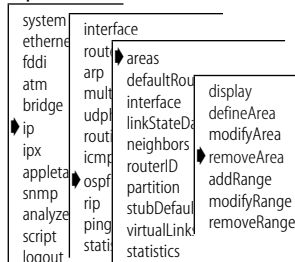
Select area:{1-2} 2
Enter Area ID [0.0.0.2]: 0.0.0.3
Is this a Stub Area (yes,no) [no]:

```

Top-Level Menu**Removing Areas**

Use this command to remove an existing area.

- 1 From the top level of the Administration Console, enter:
ip ospf areas removeArea
- 2 Enter the index number of the area to remove.

Top-Level Menu

Adding Network Ranges

A range is a network segment that can include multiple network nodes. You can add a range to a previously defined OSPF area. When you add a range, you specify only the network portion of the IP address.

To add a network range:

- 1 From the top level of the Administration Console, enter:
ip ospf areas addRange
- 2 Enter the index number of the area to which you want to add the range.
- 3 Enter the IP address of the range to add to the area.
- 4 Enter the subnet mask. Press Return or Enter to accept the current mask in brackets.
- 5 Specify whether to advertise the range on the network. The default is yes.

Example:

```
Select area {1-2}: 2
Enter IP address: 158.101.0.0
Enter subnet mask [255.255.0.0]:
Advertise this area range (yes,no) [yes]:
```

Modifying Network Ranges

Use this command to modify information that is associated with a previously defined range.

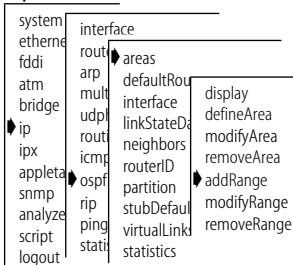
To modify a range:

- 1 From the top level of the Administration Console, enter:
ip ospf areas modifyRange
- 2 Enter the number of the area that contains the range that you want to modify.
- 3 Enter the IP address of the range to modify.
- 4 Enter the subnet mask. Press Return or Enter to accept the current subnet mask in brackets.
- 5 Specify whether to advertise the range on the network. The default is yes.

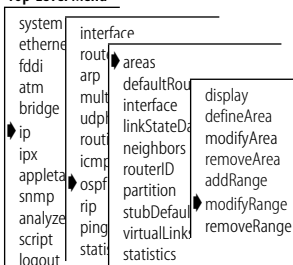
Example:

```
Select area {1-2}: 2
Enter IP address of range to modify: 158.101.0.0
Enter subnet mask [255.255.0.0]:
Advertise this area range (yes,no) [yes]:
```

Top-Level Menu



Top-Level Menu



Removing Network Ranges

This command removes a previously defined range.

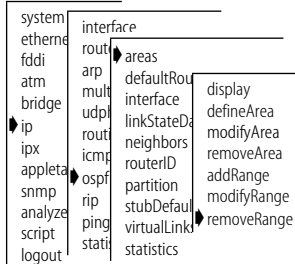
To remove a range from an area:

- 1 From the top level of the Administration Console, enter:
- 2 Enter the address to remove.
- 3 Enter the IP address of the range to delete.

Example:

```
Select area {1-2}: 2
Enter IP address of range to delete: 158.101.0.0
```

Top-Level Menu



Setting the Default Route Metric

The default route metric value indicates the cost for a default route. If the cost is greater than 0, the router advertises itself as the default router to the area.

Default

The default metric value is 0, which indicates no advertisement.

Displaying Default Route Metrics

Use this command to display the current default route metric value on the router.

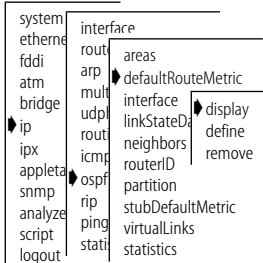
To display the default route metric, from the top level of the Administration Console, enter:

```
ip ospf defaultRouteMetric display
```

A message indicating the cost for the route appears:

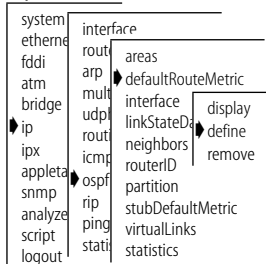
```
Default route metric = 1
```

Top-Level Menu



Defining Default Route Metrics

Top-Level Menu



Use this command to define the default route metric for the router.

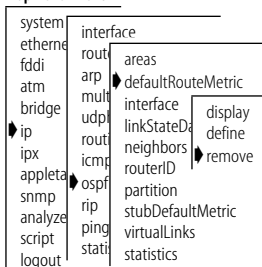
To define the default route metric:

- 1 From the top level of the Administration Console, enter:
ip ospf defaultRouteMetric define
- 2 At the prompt, enter the default route metric value, or press Return or Enter to accept **1** as the value.

Allowable values: **1** through **65535**.

Removing Default Route Metrics

Top-Level Menu



To remove a default route metric from an interface, from the top level of the Administration Console, enter:

ip ospf defaultRouteMetric remove

The designated default route metric is removed immediately.

Configuring OSPF Interfaces


This section describes how to configure OSPF interfaces by adding OSPF characteristics to existing IP virtual LAN (VLAN) interfaces. See Chapter 9, "Administering VLANs," for information about how to configure an IP VLAN interface. You can configure the following OSPF characteristics on existing IP VLAN interfaces:

- Mode
- Priority
- Area ID
- Cost
- Transmit delay
- Hello timer
- Retransmit timer
- Dead interval
- Password

Displaying OSPF Interface Information

The OSPF interface summary and detail displays contain information about the system's OSPF interface configuration.

Top-Level Menu



To display OSPF interface configuration information, from the top level of the Administration Console, enter:

```
ip ospf interface summary
```

OR

```
ip ospf interface detail
```

The following sample shows an OSPF interface summary:

```
OSPF interface summary information
```

Indx	Pri	AreaID	Xmit Cost	Xmit Delay	Hello Intvl	Rxmit Intvl	Dead Intvl	Password
1	1	0.0.0.0	1	1	10	5	40	No Password
2	1	158.101.0.0	1	1	10	5	40	No Password

The following sample shows an OSPF interface detail:

```
OSPF interface detail information
```

Indx	IP Address	State	DR	BDR	Notes
1	158.101.112.113	Off	0.0.0.0	0.0.0.0	
2	158.101.142.1	BDR	158.101.142.254	158.101.142.1	RouterID

Table 13-2 lists the fields for the OSPF displays.

Table 13-2 Field Attributes for the OSPF Interface Displays

Field	Description
Indx	Interface entry index, which corresponds to the IP interface index
Pri	OSPF router priority for the interface
Area ID	OSPF area to which the interface belongs
Xmit Cost	Interface transmit cost
Xmit Delay	Interface transmit delay
Hello Intvl	OSPF hello packet transmit interval for the interface
Rxmit Intvl	LSA retransmit interval
Dead Intvl	Time interval before OSPF declares that a neighbor is dead

(continued)

Table 13-2 Field Attributes for the OSPF Interface Displays (continued)

Field	Description
State	Interface state: <ul style="list-style-type: none"> ■ <code>Off</code> = OSPF is not enabled on the interface. ■ <code>Down</code> = The interface is down, but OSPF is enabled on it. ■ <code>Loopback</code> = The interface is a loopback interface. ■ <code>Waiting</code> = The router is trying to determine the identity of the DR and BDR on the network. ■ <code>PTP</code> = The interface is operational and connects to either a point-to-point network or a virtual link. The router attempts to form adjacency with the neighboring router. ■ <code>DRother</code> = The interface is on a multiaccess network where this router is not the DR or BDR. ■ <code>Backup</code> = The router is the BDR on the attached network. ■ <code>DR</code> = The router is the DR on the attached network.
DR	Router ID of the designated router (DR)
BDR	Router ID of the backup designated router (BDR)
Notes	When <code>RouterID</code> appears, the interface address is being used as the OSPF router ID.

Displaying OSPF Interface Statistics

Use this command to display statistics associated with specific OSPF interfaces.

To display IP interface statistics:

- 1 From the top level of the Administration Console, enter:
ip ospf interface statistics
- 2 Select an IP interface.

Top-Level Menu

system	interface	
ethernet	route	areas
fdi	arp	default
atm	multi	interface
bridge	udpl	statistics
ip	icm	linkSta
ipx	ospf	neighb
apple	rip	routerl
snmp	ping	partiti
analyze	stati	stubDe
script		virtuall
logout		statisti
		summary
		detail
		mode
		priority
		arealD
		cost
		delay
		hello
		retransmit
		dead
		password

This example shows OSPF interface statistics:

OSPF interface statistics

```

receiveHello      transmitHello     receiveDD
         50             49             3

transmitDD        receiveLSR        transmitLSR
         2             1             2

receiveLsAck      transmitLsAck     receiveLSU
         5             5             5

transmitLSU       computeDR         adjacencyUp
         1             2             1

adjacencyDown     transmitError     receiveError
         0             0             0

mismatchHello     mismatchDead      mismatchMask
         0             0             0

mismatchAreaId    mismatchAreaType  receivedUnknown
         0             0             0

authError         packetXsumError   lsaXsumError
         0             0             0

```

Table 13-3 describes the fields for the interface statistics display.

Table 13-3 Field Attributes for Interface Statistics Display

Field	Description
receiveHello	Number of hello packets that were received
transmitHello	Number of hello packets that were transmitted
receiveDD	Number of database description packets that were received
transmitDD	Number of database description packets that were transmitted
receiveLSR	Number of LSA request packets that were received
transmitLSR	Number of LSA request packets that were transmitted
receiveLsAck	Number of LSA acknowledgments that were received
transmitLsAck	Number of LSA acknowledgments that were transmitted
receiveLSU	Number of link state update packets that were received
transmitLSU	Number of link state update packets that were transmitted
computeDR	Number of times that the designated router was computed
adjacencyUp	Number of times that OSPF adjacencies have been formed
adjacencyDown	Number of times that OSPF adjacencies have gone down

(continued)

Table 13-3 Field Attributes for Interface Statistics Display (continued)

Field	Description
transmitError	Number of general transmit errors
receiveError	Number of general receive errors
mismatchHello	Number of hello packet interval mismatches that were detected
mismatchDead	Number of router dead interval mismatches that were detected
mismatchMask	Number of subnet mask mismatches that were detected
mismatchAreaID	Number of interface area ID mismatches that were detected
mismatchAreaType	Number of interface area type mismatches that were detected
receivedUnknown	Number of unknown LSAs that were received
authError	Number of authentication errors
packetXsum	Number of packet checksum errors since interface came up
lsaXsumError	Number of LSA checksum errors that were detected

Setting Modes

Use this command to set the OSPF mode for each interface. The mode can be *off* or *active*. To run OSPF routing, set the mode to *active*.

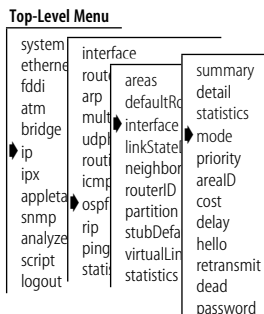


To set the OSPF interface mode to active, you must enable IP routing. See Chapter 10 “Administering IP Routing,” for information about how to enable IP routing.

Default The default is *off*.

To set the interface mode:

- 1 From the top level of the Administration Console, enter:
ip ospf interface mode
- 2 Enter the interface.
- 3 Enter the interface mode: **off** or **active**.



Setting Priorities

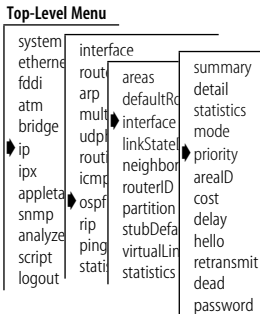
The interface priority is a value that you assign to an OSPF router to determine its status as a designated router. A router can function in one of three ways:

- **Designated router (DR)** — The router that has the highest priority value (unless a designated router already exists on the subnetwork)
- **Backup designated router (BDR)** — A router that has a lower priority value
- **Not a designated router** — A router that has a priority value of 0

Default The default priority value is 1.

To set the interface priority:

- 1 From the top level of the Administration Console, enter:
ip ospf interface priority
- 2 Enter an IP interface.
- 3 Enter the priority value.



Setting Area IDs

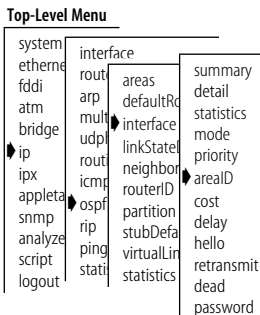
The interface area ID associates the interface that you specify with an OSPF area. See “Defining Areas” on page 13-2 for more information about OSPF areas.



CAUTION: Set the area ID to the same value for all routers on the network segment because they are in the same area.

To associate an interface with an OSPF area:

- 1 From the top level of the Administration Console, enter:
ip ospf interface areaID
- 2 Enter an IP interface.
- 3 Enter the area ID number, in the form of an IP address.



Setting Costs

The interface cost reflects the line speed of the port. Although the system calculates a default cost value based on the module media type, you can manually change the cost to a different value. In most cases, you can accept the system default value.

To set the interface cost:

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

```
ip ospf interface cost
```

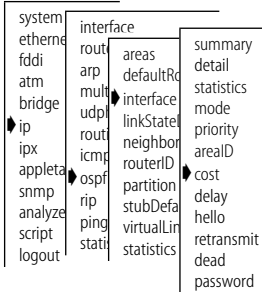
- 2 Enter an IP interface.

- 3 Enter the cost value for the interface.

The default is calculated by the system.

Allowable values: **1** through **65535**.

Top-Level Menu



Setting Transmit Delays

This command sets the OSPF interface transmit delay, in seconds. The system adds the value of the transmit delay to all link state advertisements (LSAs) that it sends out to the network. Set the transmit delay according to the link speed: use a longer transmit delay time for slower link speeds.

Default The default delay is *1* second.

To set the interface delay:

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

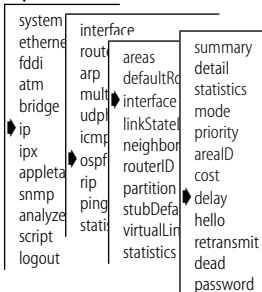
```
ip ospf interface delay
```

- 2 Enter an IP interface.

- 3 Enter the interface transmit delay value.

Range: **1** through **65535** seconds.

Top-Level Menu



Setting Hello Timers

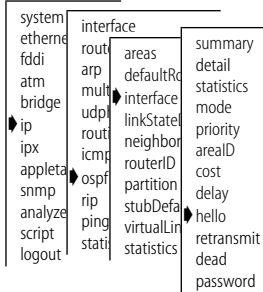
The interface hello timer determines how often, in seconds, the interface transmits hello packets to neighbor routers on the network. Hello packets notify other routers that the sending router is still active on the network. If a router does not send hello packets for a period of time specified by the dead interval, the router is considered inactive by its neighbors. See “Setting Dead Intervals” on page 13-14 for more information.

Default The default value for the hello timer is 10 seconds.



CAUTION: Set the hello timer to the same value for all routers on the network segment because they are in the same area.

Top-Level Menu



To set the hello timer:

- 1 From the top level of the Administration Console, enter:
ip ospf interface hello
- 2 Enter an IP interface.
- 3 Enter the hello timer value, in seconds.

Allowable values: 1 through 65535.

Setting Retransmit Timers

This command specifies the LSA retransmit interval for each interface. You set the retransmit interval in seconds.

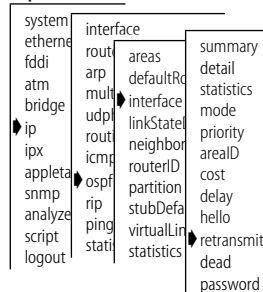
Default The default value for the retransmit timer is 5 seconds.

To set the retransmit time:

- 1 From the top level of the Administration Console, enter:
ip ospf interface retransmit
- 2 Enter an IP interface.
- 3 Enter the LSA retransmit time, in seconds.

Allowable values: 1 through 65535.

Top-Level Menu



Setting Dead Intervals

The value of the dead interval determines how long neighbor routers wait for a hello packet before they determine that the transmitting router is inactive. Each time a router receives a hello packet from a neighbor, the router resets the dead interval timer for that neighbor. See “Setting Hello Timers” on page 13-13 for more information.



CAUTION: Set the dead interval to the same value for all routers on the network segment.

Default The default value for the dead interval is 40 seconds.

To set the value of the dead interval:

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


```
ip ospf interface dead
```

- 2 Enter an IP interface.

- 3 Enter the value of the dead interval, in seconds.

Allowable values: 1 through 65535.

Top-Level Menu



Setting Passwords

Use this command to set a security password for a specific OSPF interface.



CAUTION: Use the same password for all routers on the network segment because they are in the same area.

Default The default is no password.

To enter an IP interface password:

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

```
ip ospf interface password
```

- 2 Enter an IP interface.

- 3 Enter the password.

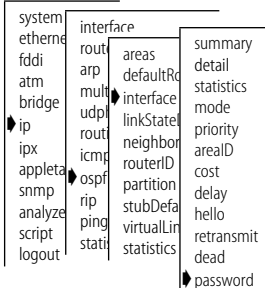
You can use up to eight ASCII characters.

Example:

```
Select interface(s) (1-2|all): 2
```

```
Enter interface password [none]: xyzzy
```

Top-Level Menu



Use the password **none** to remove a previously assigned password.

Displaying the Link State Database

The link state database contains information about different link state advertisements (LSAs).



An asterisk (*) after the router ID in a display indicates that the LSA originated locally.

Displaying a Database Summary

This display summarizes all LSAs in the link state database.

To display a link state database summary:

- 1 From the top level of the Administration Console, enter:
ip ospf linkStateData databaseSummary
- 2 Enter the area ID.
- 3 Enter the area mask.
- 4 Enter the Link State ID (LSID).
- 5 Enter the LSID mask.

Example:

```
Enter Area ID [0.0.0.0]: 0.0.0.2
Enter Area mask [0.0.0.0]:
Enter LSID [0.0.0.0]:
Enter LSID mask [0.0.0.0]:
```

The following sample shows a link state database summary:

Area ID	Checksum Summation	LSA Count	Router LSAs	Network LSAs	Summary LSAs	External LSAs
--	003175AE	107	--	--	--	107
158.101.0.0	00017729	3	2	1	0	--

Table 13-4 describes the fields for the summary display.

Table 13-4 Field Attributes for Link State Database Summary Display

Field	Description
Checksum Summation	Total of all LSA checksums
LSA Count	Number of LSAs
Router LSAs	Number of router link LSAs
Network LSAs	Number of network link LSAs
Summary LSAs	Number of summary link LSAs
External LSAs	Number of external link LSAs

Top-Level Menu

system	interface
ethernet	router
fdi	areas
atm	arp
bridge	default
ip	interface
ipx	linkStateData
appletalk	network
snmp	router
analyze	summary
script	external
logout	partitions
	stubDefaultMetric
	virtualLinks
	statistics

Displaying Router LSAs

Use this command to show the router LSAs in the link state database.

To display a link state router information summary:

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

```
ip ospf linkStateData router
```

- 2 Enter the area ID.
- 3 Enter the area mask.
- 4 Enter the LSID.
- 5 Enter the LSID mask.

The following sample shows link state database router information:

OSPF router links

LSID	Router ID	LS Seq	LS Age	Flags
158.101.142.1	158.101.142.1*	8000002E	335	ASBR
Link Type	Link ID	Link Data		Metric
Transit Net	158.101.142.254	158.101.142.1		1

LSID	Router ID	LS Seq	LS Age	Flags
158.101.142.254	158.101.142.254	80000019	330	
Link Type	Link ID	Link Data		Metric
Transit Net	158.101.142.254	158.101.142.254		1

Top-Level Menu




Table 13-5 describes the fields for the link state database router display.

Table 13-5 Field Attributes for a Link State Database Router Display

Field	Description
LSID	ID of the router that originated the LSI
Router ID	ID of the remote router
LS Seq	Hexadecimal sequence number of the LSA (used to detect older LSAs)
LS Age	Time in seconds since LSA originated
Flags	<ul style="list-style-type: none"> ■ <code>V</code> = Router is the endpoint of an active virtual link that uses the area as a transmit area ■ <code>ASBR</code> = Router is an autonomous system boundary router ■ <code>ABR</code> = Router is an area border router
Link Type	<ul style="list-style-type: none"> ■ <code>PTP</code> = Connection is point-to-point to another router ■ <code>Transit</code> = Connection is to a transit network (one with more than one OSPF router on it) ■ <code>Stub</code> = Connection is to a stub network ■ <code>Virtual link</code> = Connection is to a far-end router that is the endpoint of a virtual link
Link ID	<ul style="list-style-type: none"> ■ <code>PTP</code> = Router ID of neighboring router ■ <code>Transit</code> = Address of designated router ■ <code>Stub</code> = IP network/subnetwork number ■ <code>Virtual link</code> = Router ID of neighboring router
Link Data	<ul style="list-style-type: none"> ■ <code>PTP</code> = MIB II index value for an unnumbered point-to-point interface ■ <code>Transit</code> = IP interface address of designated router ■ <code>Stub</code> = Network IP address mask ■ <code>Virtual link</code> = IP interface address of neighboring router
Metric	Cost of the link

Displaying Network LSAs

Use this command to show the network LSAs in the link state database.

To display network link state advertisement information:

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

```
ip ospf linkStateData network
```

- 2 Enter the ID of the OSPF area.

- 3 Enter the area mask.

- 4 Enter the LSID.

- 5 Enter the LSID mask.

This example shows network link state advertisement information:

```
OSPF network links
```

```
LSID          Router ID      LS Seq      LS Age      Network Mask
158.101.142.254 158.101.142.254 80000003   975        255.255.255.0
Attached Routers
158.101.142.254
158.101.142.1
```

Table 13-6 shows the attributes for the link state database network display.

Table 13-6 Field Attributes for Link State Database Network Display

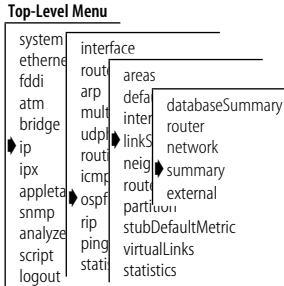
Field	Description
LSID	Interface address of designated router
Router ID	Originating router ID
LS Seq	Hexadecimal sequence number of the LSA (used to detect older LSAs)
LS Age	Time in seconds since LSA originated
Network Mask	IP address mask for the network
Attached Routers	List of routers that are fully adjacent to the designated router (DR). The ID of the DR is also listed here.

Top-Level Menu

system	interface
ethernet	route
fdi	areas
atm	arp
bridge	defa
ip	inter
ipx	links
appleta	network
snmp	summary
analyze	external
script	part
logout	stubDefaultMetric
	virtualLinks
	statistics

Displaying Summary Network LSAs

Use this command to summarize all network LSAs in the link state database.



To display network link state summary information:

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

ip ospf linkStateData summary

2 Enter the ID of the OSPF area.

3 Enter the area mask.

4 Enter the LSID.

5 Enter the LSID mask.

This example shows link state database network summary information:

OSPF summary links

LSID (Type 3)	Router ID	LS Seq	LS Age	Network Mask	Metric
193.202.111.0	193.202.111.254*	80001648	20	255.255.255.0	1
LSID (Type 3)	Router ID	LS Seq	LS Age	Network Mask	Metric
193.202.111.128	193.202.111.254*	80001136	20	255.255.255.128	11

Table 13-7 shows the fields for the link state database network summary.

Table 13-7 Field Attributes for Link State Database Network Summary Display

Field	Description
LSID	<ul style="list-style-type: none"> Type 3 = IP network number Type 4 = Router ID of ASBR's OSPF
Router ID	Originating router ID
LS Seq	Hexadecimal sequence number of the LSA (used to detect older LSAs)
LS Age	Time in seconds since LSA was originated
Network Mask	<ul style="list-style-type: none"> Type 3 = IP address mask of destination network Type 4 = Not used; must be 0
Metric	Cost to reach the network

Displaying External Network LSAs

To display network link state external information:

Top-Level Menu

system	interface
ethernet	route
fddi	arp
atm	mult
bridge	udp
ip	route
ipx	icmp
appletan	ospf
snmp	rip
analyze	ping
script	stat
logout	statistics
	areas
	defa
	inter
	linkS
	netw
	neig
	sum
	rou
	exte
	part
	stub
	virtu
	links
	stat
	istics
	Summary
	router
	network
	summary
	external
	part
	stub
	Default
	Metric
	virtual
	Links
	statistics

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

```
ip ospf linkStateData external
```

2 Enter the ID of the OSPF area.

3 Enter the area mask.

4 Enter the LSID.

5 Enter the LSID mask.

This example shows link state database external information:

OSPF external links

LSID	Router ID	LS Seq	LS Age	Network Mask
9.0.0.0	158.101.142.1*	80000004	295	255.0.0.0
Fwd Address	Metric	Type	RouteTag	
158.101.142.1	4	1	0	
LSID	Router ID	LS Seq	LS Age	Network Mask
89.0.0.0	158.101.142.1*	80000004	295	255.0.0.0
Fwd Address	Metric	Type	RouteTag	
158.101.142.1	4	1	0	

Table 13-8 shows the attributes for the link state database external display.

Table 13-8 Field Attributes for the Link State Database External Display

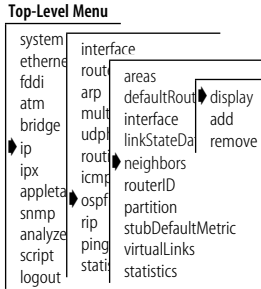
Field	Description
LSID	IP network number.
Router ID	Originating router ID.
LS Seq	Hexadecimal sequence number of the LSA (used to detect older LSAs).
LS Age	Time in seconds since LSA originated.
Network Mask	IP address mask for the advertised destination.
Fwd Address	Forwarding address for data traffic to the advertised destination.
Metric	Cost to reach advertised destination.
Type	<ul style="list-style-type: none"> ■ Type 1 = normal link state metric. ■ Type 2 = metric is larger than any local link state path.
RouteTag	Not used by OSPF. These 32 bits may be used to communicate other information between boundary routers. Application systems usually define their own route tags.

Administering Neighbors

Neighbor routers are physically attached to the same network segment and exchange OSPF routing tables.

Displaying Neighbors

Use this command to display information about the currently defined neighbors in an OSPF area.



To display information about OSPF neighbors, from the top level of the Administration Console, enter:

```
ip ospf neighbors display
```

This example shows OSPF neighbors:

OSPF neighbor information

Indx	Neighbor Addr	Router ID	State	Pri	RxQ	SumQ	ReqQ	Flags
2	158.101.142.254	158.101.142.254	Full	1	0	0	0	D+DR

Table 13-9 Field Attributes for Neighbors Display

Field	Description
Indx	Interface index that a neighbor belongs to.
Neighbor Addr	Interface address of neighbor.
Router ID	Router ID of neighbor's OSPF.
State	Neighbor's adjacency: <ul style="list-style-type: none"> ■ Down = No recent data received from neighbor, connection is down. ■ Attempt = Only used on nonbroadcast networks. No recent data received from neighbor (will attempt to contact). ■ Init = Have recently seen hello packet from neighbor, however two-way communication has not been established. ■ Two-way = Bidirectional communication has been established. ■ ExStart = Taking initial step to create adjacency between neighboring routers. ■ Exchange = Database descriptions are being exchanged. ■ Loading = LSA databases are being exchanged. ■ Full = Neighboring routers are fully adjacent.
Pri	Router priority of neighbor's OSPF.
RxQ	Number of LSAs in local retransmit queue to the neighbor.

(continued)

Table 13-9 Field Attributes for Neighbors Display (continued)

Field	Description
SumQ	Number of LSAs in LSA summary queue for the neighbor.
ReqQ	Number of LSAs being requested from neighbor.
Flags	Neighbor identification flags: <ul style="list-style-type: none"> ■ D = dynamic neighbor. ■ S = static neighbor. ■ BDR = backup designated router. ■ DR = designated router. Example: [S, BDR] + [D, DR] is a static neighboring backup designated router and a dynamic neighboring designated router.

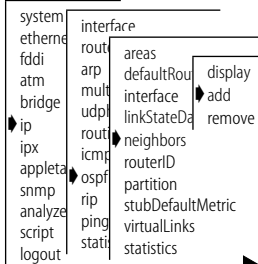
Adding Neighbors

You can add a neighbor static IP address to an existing interface.



The CoreBuilder 2500 system learns neighbor addresses dynamically on interfaces that support multicast routing. Define static neighbors only on nonmulticast interfaces.

Top-Level Menu



To add a neighbor:

- 1 From the top level of the Administration Console, enter:
ip ospf neighbors add
- 2 Enter the interface to which to add the OSPF neighbor.
- 3 Enter the static IP address of the neighbor.




The system deletes a dynamically learned neighbor that is inactive for more than 5 minutes.

Removing Neighbors

Use this command to remove a static or dynamic neighbor address from an existing interface.

Top-Level Menu



To remove a neighbor:

- 1 From the top level of the Administration Console, enter:
ip ospf neighbors remove
- 2 Enter the IP interface.
- 3 Enter the address of the neighbor that you want to remove.



The system relearns dynamic neighbors.

Setting OSPF Router IDs

The OSPF router ID identifies the router to other routers within an autonomous system. Three types of router identifiers are available; all three take the form of an IP address:

- **Default** — An arbitrary ID that the system generates and uses as the default router ID
- **Interface** — The index of an IP interface on the router
- **Address** — An ID that you define in the form of an IP address



OSPF routing must be inactive (off) before you can add or modify an OSPF router ID. To set the OSPF mode to off, see “Setting Modes” on page 13-10 for details. After you add the router ID, set the OPSF mode to active on the interface.

To set the router ID:

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

```
ip ospf routerID
```

The system displays the current router ID and the router ID type.

- 2 Enter the router ID type (**default**, **interface**, or **address**), or press Return or Enter to accept the currently defined type.
- 3 Enter the interface or address for the router ID.

Example:

```
Current OPSF router id = 158.101.142.1 (interface)
Enter router ID type {default,interface,address}(address)
Enter router ID [158.101.112.113]: 158.101.112.111
```

Top-Level Menu

system	interface	
ethernet	route	areas
fdi	arp	defaultRouteMetric
atm	multicast	interface
bridge	udphelper	linkStateData
ip	routing	neighbors
ipx	icmpRoute	routerID
appleata	ospf	partition
snmp	rip	stubDefaultMetric
analyze	ping	virtualLinks
script	statistics	statistics
logout		

Administering Memory Partitions

You can specify how much memory the OSPF protocol can use for its data processing and storage from the Administration Console.

Displaying Memory Partitions

You can display a module's OSPF memory allocation. From the top level of the Administration Console, enter:

Top-Level Menu

```

system
ethernet interface
fdi      rout
atm      arp
bridge  mult
ip       udpl
ipx      rout
appleta icmp
snmp     ospf
analyze rip
script  ping
logout  stati

```

```
ip ospf partition display
```

This example shows an OSPF memory partition summary:

```

Current partition maximum size = 500000 (bytes).
Configured partition maximum size = 500000 (bytes).
Allocated partition size = 100000 (bytes).

```

This display shows three partition parameters:

- **Current partition maximum size** (500000 in this example) — The OSPF memory limit implemented at the last system reboot
- **Configured partition maximum size** (500000 in this example) — The last value that you entered, which becomes the current partition maximum size at system reboot
- **Allocated partition size** (100000 in this example) — The module's current working memory. OSPF dynamically allocates memory in 100,000-byte chunks up to the current partition maximum size.



If the system shows a partition size of 0, the OSPF protocol is using the CoreBuilder 2500 system memory partition.

Modifying Memory Partitions

Use this command to change the OSPF memory allocation. This change takes effect at system reboot. To modify a module's OSPF memory allocation:

Top-Level Menu

```

system
ethernet interface
fdi      rout
atm      arp
bridge  mult
ip       udpl
ipx      rout
appleta icmp
snmp     ospf
analyze rip
script  ping
logout  stati

```

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

```
ip ospf partition modify
```

- 2 Enter the new partition size (in bytes).

Example:

```

Maximum partition size is 8443220 bytes
Enter new partition maximum size (in bytes) [500000]: 600000
New partition size will take effect after reboot.

```



The maximum partition size (8443220 in this example) shows how much total memory is available to define as the OSPF maximum partition.

Administering Stub Default Metrics

The stub default metric value determines whether the router generates the default route into the stub areas of the network. This value applies to area border routers that have attached stub areas.

Displaying Stub Default Metrics

To display the stub default metric value, enter the following command from the top level of the Administration Console:

```
ip ospf stubDefaultMetric display
```

A message appears indicating the current value:

```
Stub default metric = 20
```

Top-Level Menu

```
system
etherne
fddi
atm
bridge
ip
ipx
appleta
snmp
analyze
script
logout
interface
rout
arp
mult
udp
rout
icm
ospf
rip
ping
stati
areas
default
interfa
linkSta
neighbors
routerID
partition
stubDefaultMetric
virtualLinks
statistics
```

Defining Stub Default Metrics

Use this command to set the stub default metric value on an area border router with an attached stub area.

Default

The default value is 1.

To add a stub default metric:

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

```
ip ospf stubDefaultMetric define
```

- 2 Enter the stub default metric value or press Return or Enter to select the current value.

Allowable values: 1 through 65535.

Top-Level Menu

```
system
etherne
fddi
atm
bridge
ip
ipx
appleta
snmp
analyze
script
logout
interface
rout
arp
mult
udp
rout
icm
ospf
rip
ping
stati
areas
default
interfa
linkSta
neighbors
routerID
partition
stubDefaultMetric
virtualLinks
statistics
```

Removing Stub Default Metrics

Use this command to disable the stub default metric value on the router.

To disable the stub default metric, from the top level of the Administration Console, enter:

```
ip ospf stubDefaultMetric remove
```

The stub default metric is removed immediately.

Top-Level Menu

```
system
etherne
fddi
atm
bridge
ip
ipx
appleta
snmp
analyze
script
logout
interface
rout
arp
mult
udp
rout
icm
ospf
rip
ping
stati
areas
default
interfa
linkSta
neighbors
routerID
partition
stubDefaultMetric
virtualLinks
statistics
```

Administering Virtual Links

Virtual links establish connections from your local area to other autonomous systems in the network. You can define, remove, modify, and display the virtual links on your system.

Displaying Virtual Links

The virtual links display shows the virtual links associated with the interface you specify.

Top-Level Menu

system	interface	
ethernet	router	
fddi	arp	display
atm	defaultRoute	define
bridge	interface	modify
bridge	linkStateData	remove
ip	neighbors	
ipx	routerID	
appletalk	icmp	
snmp	ospf	partition
analyze	rip	stubDefaultMetric
script	ping	virtualLinks
logout	status	statistics

To display virtual links, from the top level of the Administration Console, enter:

```
ip ospf virtualLinks display
```

This example shows virtual links:

Virtual links

Indx	Interface Address	Router ID	Rxmit Intvl	Router Address	Link State	Link Cost
1	193.202.111.1	0.0.0.8	50	193.202.111.194	PTP	2

Table 13-10 shows the attributes for the virtual links display.

Table 13-10 Field Attributes for Virtual Links Display

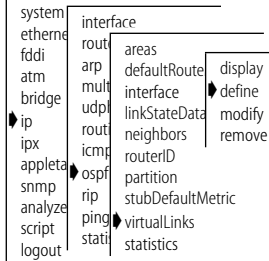
Field	Description
Indx	Index of the local interface that the virtual link is connected to
Interface Address	Local interface address
Router ID	Router ID of the remote router's OSPF
Rxmit Intvl	LSA retransmit interval for the virtual link
Router Address	Interface address of remote router (changes dynamically)
Link State	Virtual link state
Link Cost	Cost of virtual link (computed dynamically)

Defining Virtual Links

Use this command to create a new virtual link to a destination router. You must configure a virtual link for each area border router that has an interface outside the backbone area.



You can define up to 32 virtual links per router.

Top-Level Menu

To define a virtual link:

- 1 From the top level of the Administration Console, enter:
ip ospf virtualLinks define
- 2 Enter the IP interface index.
- 3 Enter the router ID of the destination router.
- 4 Enter the retransmit interval.

Example:

```
Select IP interface index {1-2}: 1
Enter destination router ID []: 192.42.95.254
Enter retransmit interval [50]: 100
```

Modifying Virtual Links

Use this command to modify a virtual link.

To modify virtual link parameters:

- 1 From the top level of the Administration Console, enter:
ip ospf virtualLinks modify
- 2 Enter the IP address of the destination router.
- 3 Enter the retransmit interval, in seconds.

Example:


```
Select IP interface index {1-2}: 1
Enter destination router ID [192.42.95.254]: 192.42.95.150
Enter retransmit interval [50]: 75
```

Removing Virtual Links

Use this command to remove a virtual link.

To remove a virtual link, from the top level of the Administration Console, enter:

```
ip ospf virtualLinks remove
```

Top-Level Menu

Displaying OSPF Statistics

Use this command to display general OSPF statistics. To display general OSPF statistics, from the top level of the Administration Console, enter:

```
ip ospf statistics
```

Top-Level Menu

system	interface
ethernet	route
fdi	areas
atm	arp
bridge	defaultRouteMetric
ip	interface
ipx	mult
appleta	udpl
snmp	linkStateData
analyze	route
script	neighbors
logout	icm
	routerID
	ospf
	partition
	rip
	stubDefaultMetric
	ping
	virtualLinks
	statist
	statistics

This example shows OSPF statistics:

```
OSPF general statistic
```

```

SPFComputations      memoryFailures      LSAsTransmitted
          0              0              0
          20             0             328

      LSAsReceived  routeUpdateErrors      rcvErrors
          0              0              0
          28             0              8

extLsaChanges      softRestarts
          0              0
          109           0

```

Table 13-11 shows the attributes of the OSPF statistics display.

Table 13-11 Field Attributes for OSPF Statistics Display

Field	Description
SPFComputations	Number of shortest-path-first computations done.
memoryFailures	Number of nonfatal memory-allocation failures.
LSAsTransmitted	Number of link state advertisements transmitted.
LSAsReceived	Number of link state advertisements received.
routeUpdateErrors	Number of nonfatal routing table update failures.
rcvErrors	Number of general receive errors.
extLsaChanges	Number of external LSA changes made to database.
softRestarts	Number of OSPF router soft restarts because of insufficient memory resources (implies a fatal memory-allocation failure). Solution: change the OSPF memory partition, add memory, or reconfigure the network topology to generate smaller OSPF databases.

14

ADMINISTERING APPLE TALK ROUTING

This chapter describes how to set up your CoreBuilder™ 2500 system to use the AppleTalk protocol to route packets. For more information about how AppleTalk routing works, see Chapter 8, “Routing with AppleTalk Technology.”

This chapter describes the following tasks:

- Administering Interfaces
- Administering Routes
- Administering the AARP Cache
- Displaying the Zone Table
- Configuring Forwarding
- Configuring Checksum
- Pinging an AppleTalk Node
- Viewing AppleTalk Statistics

Administering Interfaces

An AppleTalk interface defines the relationship between an AppleTalk virtual LAN (VLAN) and the AppleTalk network. Every AppleTalk interface has one AppleTalk VLAN associated with it. Each module has one AppleTalk interface defined for each subnetwork that is directly connected to it.



Before you define an associated AppleTalk interface, define a VLAN, as described in Chapter 9.



You can configure a maximum of 32 interfaces per router.


An AppleTalk interface has several elements 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 that you enter. This information includes 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 detect. After the nonseed interface obtains a network configuration, it can participate in network routing.
- **Network Range** — A range of numbers used to designate a network segment's identity. This element allows the physical segment between two CoreBuilder 2500 systems to support a range of multiple networks.
- **Address** — The AARP address based on the network range and the network node (1 through 253).
- **Zone** — The default zone name, and up to 15 additional defined zones.
- **State** — The status of the AppleTalk interface, which indicates whether the interface is available for communications (`up`) or unavailable (`down`).
- **VLAN Index** — The number of the AppleTalk VLAN associated with the AppleTalk interface. The VLAN index indicates the bridge ports that are associated with the AppleTalk interface. When an Administration Console menu prompts you for this option, it displays a list of defined VLAN indexes and their ports.

Displaying AppleTalk Interfaces

Use this command to display all of the AppleTalk interfaces and their parameter settings that are configured for the system.

Top-Level Menu



To display the AppleTalk interfaces defined on the router, from the top level of the Administration Console, enter:

```
appletalk interface display
```

This example shows an AppleTalk interface:

```
AppleTalk forwarding is enabled.
```

Index	Network Range	Address	State	VLAN index
1	20112-20112	20112.27	enabled	3
2	20124-20124	20124.1	enabled	4
3	20125-20125	20125.1	enabled	6

Defining an Interface

When you define an interface, you define the interface's network range, zone name, and the VLAN index associated with the interface. Before you define the AppleTalk interface to associate with an AppleTalk VLAN, define the VLAN.

To define an AppleTalk interface:

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

```
appletalk interface define
```

The following message appears:

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

When you are prompted for interface parameters, press Return or Enter to accept the existing value in brackets.

- 2 Enter **n** (for no) or **y** (for yes).
- 3 Enter the start of the network range associated with the interface.
- 4 Enter the end of the network range associated with the interface.
- 5 Enter the default zone name.



Clients that have not been configured to use a particular zone use the default zone name.

- 6 Enter the zone name.
- 7 After you enter all the zone names, type **q**.
- 8 Enter the index of the AppleTalk VLAN that is associated with this interface.




You can enter up to 16 zone names per interface.

Example:

```
Enter start of network range: 10000
Enter end of network range: 10100
Enter default zone: engineering
Enter zone name: q
Appletalk VLANs:
  Index    Ports
    3      1-8
    4      9-12
Select VLAN index: 3
```


Top-Level Menu



Removing an Interface

Use this command to remove an interface when you no longer perform routing on the ports that are associated with the interface.

Top-Level Menu



To remove an interface:

- 1 From the top level of the Administration Console, enter:
appletalk interface remove
- 2 Enter the index numbers of the interfaces you want to remove.

The interfaces are no longer defined on the router.

Administering Routes

Your system maintains a table of routes to other AppleTalk networks. The Routing Table Maintenance Protocol (RTMP) automatically generates the routing table. RTMP defines rules for:

- Information contained within each routing table
- Exchanging information between routers so that the routers can maintain their routing tables


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

Use this command to display the system's routing tables to determine the routes that are configured and whether they are operational.

Top-Level Menu



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

appletalk route display

This example shows a routing table:

AppleTalk 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, from the top level of the Administration Console, enter:

```
appletalk route flush
```

The dynamically learned routes are removed from the routing table.

Top-Level Menu

system	
ethernet	interface
fddi	▶ route
atm	display
bridge	aarp
ip	▶ flush
ipx	zone
▶ appletalk	forwarding
snmp	checksum
analyzer	ping
script	statistics
logout	

Administering the AARP Cache

Use AARP to map hardware addresses to an AppleTalk protocol address. AppleTalk uses dynamically assigned 24-bit addresses, unlike the statically assigned 48-bit addresses that Ethernet uses.

To make the address mapping process easier, AARP uses an Address Mapping Table (AMT). The AMT maintains the most recently used addresses. 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 also registers a node's dynamically assigned address on the network, as follows:

- AARP randomly assigns an address.
- To determine whether another node is already using the address, AARP broadcasts AARP probe packets to this 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 AARP/AMT address cache
- Remove entries from the cache
- Flush the cache


Displaying the AARP Cache

Use this command to display the AARP cache for the system to determine the routes that are configured and whether they are operational.

To display the AARP cache, from the top level of the Administration Console, enter:

```
appletalk aarp display
```

Top-Level Menu




This example shows an AARP cache:

AppleTalk 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

Removing an Address from the Cache

To remove an AARP cache entry:



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

```
appletalk aarp remove
```

- 2 Enter the AARP address at the prompt.

The entry is removed.


Flushing All Cache Entries

To flush all AARP cache entries, from the top level of the Administration Console, enter:

```
appletalk aarp flush
```

All addresses are removed from the AARP cache.

Top-Level Menu



Displaying the Zone Table

By logically grouping AppleTalk nodes into zones, you can more easily navigate through the network. You group nodes using 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, from the top level of the Administration Console, enter:


```
appletalk zone network
```

or

```
appletalk zone zone
```

Depending on the command that you enter, the table displays information by network or by zone. The following examples show both types of zone display:

Top-Level Menu



Zone Table by Network Numbers

AppleTalk 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
```

Zone Table by Zones

AppleTalk 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
```

Configuring Forwarding

Top-Level Menu

```

system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  interface
    route
    aarp
    zone
    forwarding
    checksum
    ping
    statistics

```

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 subnetwork to another when required. When you disable AppleTalk forwarding, the router discards any AppleTalk packets not addressed directly to one of its defined interfaces.

To configure AppleTalk forwarding:

- 1 From the top level of the Administration Console, enter:
appletalk forwarding
- 2 At the prompt, enter **enabled** or **disabled**

Configuring Checksum

AppleTalk uses checksums to detect errors in data transmissions. A checksum totals all data bytes and adds this sum to the end of the data packet. The receiving station computes a verification checksum from the incoming data and compares the new checksum with the value sent with the data. If the values do not match, the transmission contains an error. You can enable or disable checksum generation and verification.

To enable or disable checksum generation/verification:

- 1 From the top level of the Administration Console, enter:
appletalk checksum
- 2 At the checksum generation prompt, enter **enabled** or **disabled**
- 3 At the checksum verification prompt, enter **enabled** or **disabled**

Top-Level Menu

```

system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
  interface
    route
    aarp
    zone
    forwarding
    checksum
    ping
    statistics


```

Pinging an AppleTalk Node

The AppleTalk Echo Protocol (AEP) sends a datagram (an Echo Request) from one node to another. The destination node returns, or *echoes*, the datagram to the sender (using an Echo Reply). This process allows you to determine whether or not a node is accessible before starting a session.

To ping an AppleTalk node:

- 1 From the top level of the Administration Console, enter:
appletalk ping
The system prompts you 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 for 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, from the top level of the Administration Console, enter:

appletalk statistics ddp

The following sample shows DDP statistics:

AppleTalk 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			




Table 14-1 describes DDP statistics.

Table 14-1 AppleTalk DDP Statistics

Field	Description
inReceives	Total number of packets that were received, including those with errors
inForwards	Total number of packets that were forwarded, including those with errors
inLocals	Number of DDP datagrams for which this entity was their final DDP destination
inNoRoutes	Number of DDP datagrams that were dropped for unknown routes
inNoClients	Number of DDP datagrams that were dropped for unknown DDP types
inTooShorts	Number of input DDP datagrams that were 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 that were dropped because they exceeded the maximum DDP datagram size
inShortDdps	Number of input DDP datagrams that were dropped because this entity was not their final destination and their type was short DDP
inCsumErrors	Number of DDP datagrams that were dropped because of a checksum error
inBcastErrors	Number of DDP datagrams for which this DDP entity was their final destination, and that were dropped because of a broadcast error
inTooFars	Number of input datagrams that were dropped because this entity was not their final destination and their hop count would exceed 15
inDiscards	Number of DDP Datagrams that were thrown out during routing
outLocals	Number of host-generated DDP datagrams

Displaying RTMP Statistics

To display RTMP statistics, from the top level of the Administration Console, enter:

Top-Level Menu

system		
ethernet	interface	
fdi	route	summary
atm	aarp	rtmp
bridge	zone	zip
ip	forwarding	nbp
ipx	checksum	
appletalk	ping	
snmp	statistics	
analyzer		
script		
logout		

```
appletalk statistics rtmp
```

The following example shows RTMP statistics:

```
AppleTalk forwarding is enabled.
```

```

           inDatas      inRequests      outDatas      outRequests
           7204          0          4865          6

           routeEqChgs  routeLessChgs  routeDeletes  routeOverflows
           0             0             0             0

           inVersionErrs  inOtherErrs
           0              119

```

Table 14-2 describes RTMP statistics.

Table 14-2 AppleTalk 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
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 that 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 Statistics

To display ZIP statistics, from the top level of the Administration Console, enter:

Top-Level Menu

```

system
ethernet
  fddi
  atm
  bridge
  ip
  ipx
  appletalk
  snmp
  analyzer
  script
  logout
interface
  route
  aarp
  zone
  forwarding
  checksum
  ping
  statistics
summary
  rtmp
  zip
  nbp

```

```
appletalk statistics zip
```

This example shows ZIP statistics:

AppleTalk 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		
0	0		

Table 14-3 describes the ZIP statistics that you can view.

Table 14-3 AppleTalk ZIP Statistics

Field	Description
inQueries	Number of ZIP queries that were received
inReplies	Number of ZIP replies that were received
inExReplies	Number of ZIP extended replies that were received
inGniRequests	Number of ZIP GetNetInfo request packets that were received
inGniReplies	Number of ZIP GetNetInfo reply packets that were received
inLocalZones	Number of Zip GetLocalZones requests packets that were received
inZoneLists	Number of Zip GetZoneLists requests packets that were received
inObsoletes	Number of ZIP Takedown or ZIP Bringup packets that were received
inZoneCons	Number of times that a conflict has been detected between this entity's zone information and another entity's zone information
inZoneInvs	Number of times that 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 that were received and then rejected for any error
outQueries	Number of ZIP queries that were sent
outReplies	Number of ZIP replies that were sent
outExReplies	Number of ZIP extended replies that were sent
outGniRequests	Number of ZIP GetNetInfo packets that were sent
outGniReplies	Number of ZIP GetNetInfo reply packets that were sent out of this port
outLocalZones	Number of transmitted ZIP GetLocalZones reply packets
outZoneLists	Number of transmitted ZIP GetZoneList reply packets
outzoneInvs	Number of times that 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 that this entity had to broadcast a ZIP GetNetInfo reply because the GetNetInfo request had an invalid address

Displaying NBP Statistics

The NBP translates between numeric IP addresses and alphanumeric AppleTalk entity names.

To display NBP statistics, from the top level of the Administration Console, enter:

```
appletalk statistics nbp
```

This example shows NBP statistics:

```
AppleTalk forwarding is enabled.
```

```

          inLkupReqs      inBcastReqs      inFwdReqs      inLkupReplies
                3093                611                5951                0

          inErrors
                0


```


Table 14-4 describes NBP statistics.

Table 14-4 AppleTalk NBP Statistics

Field	Description
inLkupReqs	Number of NBP Lookup Requests received
inBcastReqs	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


Top-Level Menu





RMON AND THE COREBUILDER 2500 SYSTEM

Chapter 15 Remote Monitoring (RMON) Technology



REMOTE MONITORING (RMON) TECHNOLOGY

This chapter provides an overview of RMON and describes the specific CoreBuilder™ 2500 RMON implementation.

This chapter includes these sections about RMON:

- Overview
- RMON Benefits
- RMON in the CoreBuilder 2500
- RMON Groups
- 3Com Transcend RMON Agents
- RMON Management Information Base (MIB)

Overview

The Remote Monitoring (RMON) Management Information Base (MIB) provides a way to monitor and analyze a LAN from a remote location. The Internet Engineering Task Force (IETF) defines RMON in the document RFC 1757. A typical RMON implementation has two components:

- **Probe** — Connects to a LAN segment, examines all the LAN traffic on that segment, and keeps a summary of statistics (including historical data) in its local memory.
- **Management Console** — Communicates with the probe and collects summarized data from it. The console does not need to reside on the same network as the probe. It can manage the probe through SNMP or through out-of-band connections.

The RMON specification consists almost entirely of the definition of the MIB. The RMON MIB contains standard MIB variables that are defined to collect comprehensive network statistics that alert a network administrator to significant network events. If the embedded RMON agent operates full time, it collects data on the correct port when the relevant network event occurs.

RMON Benefits

From a network management console, traditional network management applications poll network devices such as switches, bridges, and routers at regular intervals. The console gathers statistics, identifies trends, and highlights network events. The console polls network devices constantly to determine if the network is within its normal operating conditions.

As network size and traffic levels grow, however, the network management console can become overburdened by the amount of data it must collect. Frequent console polling also generates significant network traffic that itself can create problems for the network.

An RMON implementation offers solutions to both of these problems:

- The RMON probe examines the network without affecting the characteristics and performance of the network.
- The RMON probe reports by exception rather than by gathering constant or frequent information. That is, the RMON probe informs the network management console directly if the network enters an abnormal state. The console can then use more information gathered by the probe, such as historical information, to diagnose the abnormal condition.

RMON in the CoreBuilder 2500

CoreBuilder 2500 Extended Switching software offers full-time embedded RMON support using SNMP for seven of the RMON groups. When combined with the Roving Analysis Port (RAP) function, RMON support for these groups provides a comprehensive and powerful mechanism for managing your network.



You can gain access to the RMON capabilities of the CoreBuilder 2500 system only through SNMP applications such as Transcend[®] Enterprise Manager software, not through the serial interface or telnet. For more information about the details of managing 3Com devices using RMON and Transcend tools, see the user documentation for the Transcend Enterprise Manager software.

RMON Groups

The CoreBuilder 2500 system supports seven of the RMON groups that the IETF has defined. These seven groups are described in Table 15-1.

Table 15-1 RMON Groups Supported in the CoreBuilder 2500

Group	Group Number	Purpose
Statistics	1	Maintains utilization and error statistics for the segment being monitored
History	2	Gathers and stores periodic statistical samples from the statistics group
Alarm	3	Allows you to define thresholds for any MIB variable and trigger alarms
Host	4	Discovers new hosts on the network by keeping a list of source and destination physical addresses that are seen in good packets
HostTopN	5	Allows you to prepare reports that describe the hosts that top a list sorted by one of their statistics
Matrix	6	Stores statistics for conversations between pairs of addresses
Event	9	Allows you to define actions (generate traps, log alarms, or both) based on alarms

The CoreBuilder 2500 system also supports the RMON/FDDI extension groups that the AXON Enterprise-specific MIB specifies. See Table 15-2.

Table 15-2 RMON/FDDI Extension Groups Supported in the CoreBuilder 2500

Group	Group Number	Purpose
axFDDIStatistics	axFDDI group 1	Maintains utilization and error statistics for the monitored segment
axFDDIHistory	axFDDI group 2	Gathers and stores periodic statistical samples from the statistics group

Statistics and axFDDIStatistics Groups

The statistics and axFDDIStatistics groups record frame statistics for Ethernet and FDDI interfaces. The information available per interface segment includes these statistics:

- Number of received octets
- Number of received packets
- Number of received broadcast packets
- Number of received multicast packets
- Number of received packets with CRC or alignment errors (for Ethernet frames)
- Number of received packets with FCS or invalid data length (for FDDI frames)
- Number of received packets that are undersized but otherwise well-formed
- Number of received packets that are oversized but otherwise well-formed (for Ethernet frames)
- Number of detected transmit collisions

Byte sizes include the 4-byte FCS but exclude the framing bits. Table 15-3 lists the number of the packet length counters.

Table 15-3 Supported Frame Sizes for Ethernet and FDDI

Frame Lengths (Bytes)	
Ethernet	FDDI
64	22 or fewer
65 to 127	23 to 63
	64 to 127
128 to 511	128 to 511
512 to 1023	512 to 1023
1024 to 1518	1024 to 2047
	2048 to 4095

History and axFDDIHistory Groups

The history and axFDDIHistory groups record periodic statistical samples for Ethernet and FDDI interfaces and store them for later retrieval. The information available per interface for each time interval includes these statistics:

- Number of received octets
- Number of received packets
- Number of received broadcast packets
- Number of received multicast packets
- Number of received packets with CRC or alignment errors
- Number of received packets that are undersized but otherwise well-formed
- Number of received packets that are oversized but otherwise well-formed
- Number of detected transmit collisions
- Estimate of the mean physical layer network utilization

Alarm Group

The CoreBuilder 2500 system supports the following RMON alarm mechanisms:

- Counters
- Gauges
- Integers
- Timeticks

These RMON MIB objects yield alarms when the network exceeds predefined limits. The most frequently used objects are *counters*, although the other objects may be used in much the same way. The balance of this chapter illustrates RMON functions using counters.

Counters hold and update the number of times an event occurs on a port, module, or switch. *Alarms* monitor the counters and report when counters exceed their set threshold.

Counters are useful when you compare their values at specific time intervals to determine rates of change. The time intervals can be short or long, depending on what you measure.

Occasionally, counters can produce misleading results. Because counters are finite, they are useful for comparing rates. When counters reach a predetermined limit, they *roll over* (that is, return to 0). A single low counter value may accurately represent a condition on the network. On the other hand, the same value may simply indicate a rollover.



When you disable a port, the application may not update some of its associated statistics counters.

An alarm calculates the difference in counter values over a set time interval and remembers the high and low values. When the value of a counter exceeds a preset threshold, the alarm reports this occurrence.

Using Transcend Enterprise Manager or any other SNMP network management application, you can assign alarms to monitor any counter, gauge, timetick, or integer. See the documentation for your management application for details about setting up alarms.

Setting Alarm Thresholds

Thresholds determine when an alarm reports that a counter has exceeded a certain value. You can set alarm thresholds manually through the network, choosing any value for them that is appropriate for your application. The network management software monitors the counters and thresholds continually during normal operations to provide data for later calibration.

Figure 15-1 shows a counter with thresholds set manually.

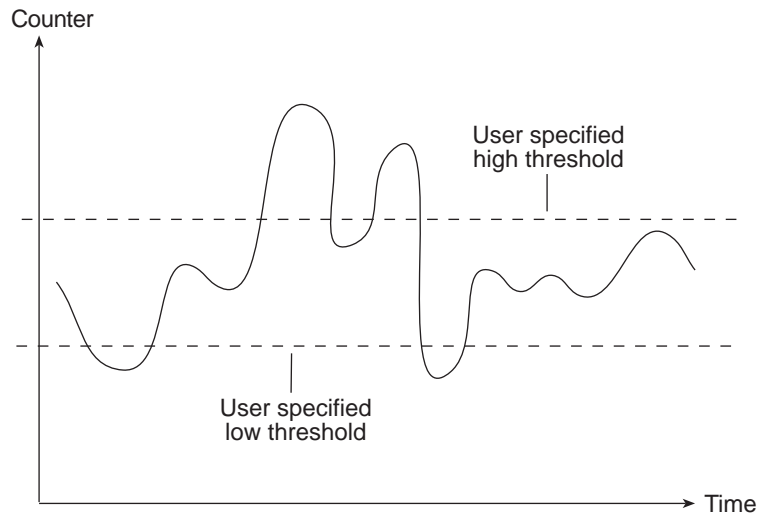


Figure 15-1 Manually Set Thresholds

You can associate an alarm with the high threshold, the low threshold, or both. The actions that occur because of an alarm depend on the network management application.

RMON Hysteresis Mechanism

The RMON hysteresis mechanism prevents small fluctuations in counter values from causing alarms. Alarms occur only when one of these events occurs:

- The counter value exceeds the high threshold after previously falling below the low threshold. (An alarm does not occur if the value has not fallen below the low threshold before rising above the high threshold.)
- The counter value falls below the low threshold after previously exceeding the high threshold. (An alarm does not occur if the value has not risen above the high threshold before falling below the low threshold.)

For example, in Figure 15-1, an alarm occurs the first time that the counter exceeds the high threshold, but not the second time. At the first instance, the counter is rising from below the low threshold. In the second instance, the counter is not rising from below the low threshold.

Host Group The host group detects hosts on the network by their physical MAC addresses. The host group records the following statistics for each host:

- Number of received packets
- Number of transmitted packets
- Number of received octets
- Number of transmitted octets
- Number of transmitted broadcast packets
- Number of transmitted multicast packets

These statistics, indexed by the relative order in which the hosts are discovered, appear in *hostTimeTable*.

HostTopN Group The HostTopN group prepares reports describing hosts that top a list that is sorted in order of one of their statistics. Information from this group includes these statistics:

- Number of received packets
- Number of transmitted packets
- Number of received octets
- Number of transmitted octets
- Number of transmitted broadcast packets
- Number of transmitted multicast packets

Matrix Group The matrix group records the following statistics about conversations between sets of addresses:

- Number of packets transmitted from the source address to the destination address
- Number of octets, excluding errors, transmitted from the source address to the destination address
- Number of bad packets transmitted from the source address to the destination address

Event Group The event group logs alarms or traps network event descriptions. Although alarm group thresholds trigger most events, other RMON groups may define event conditions.

3Com Transcend RMON Agents

RMON requires one probe per LAN segment. Because a segment is a portion of the LAN that is separated by a bridge or router, the cost of implementing many probes in a large network can be high.

To solve this problem, 3Com has built an inexpensive RMON probe into the Transcend SmartAgent software in each CoreBuilder 2500 system. With this probe, you deploy RMON widely around the network at a cost of no more than that of traditional network monitors.

Placing probe functionality inside the CoreBuilder 2500 system has these advantages:

- You can integrate RMON with normal device management.
- The CoreBuilder 2500 system can manage conditions proactively.

The CoreBuilder 2500 system associates statistics with individual ports and then takes action based on these statistics. For example, the system can generate a log event and send an RMON trap if errors on a port exceed a threshold set by the user.



To manage RMON, you must assign an IP address to the CoreBuilder 2500 system. See the CoreBuilder 2500 Administration Console User Guide for information about how to assign an IP address.

Figure 15-2 shows an example of a CoreBuilder RMON implementation. The CoreBuilder 2500 system in this figure has two Fast Ethernet connections in addition to the 10BASE-T connections.

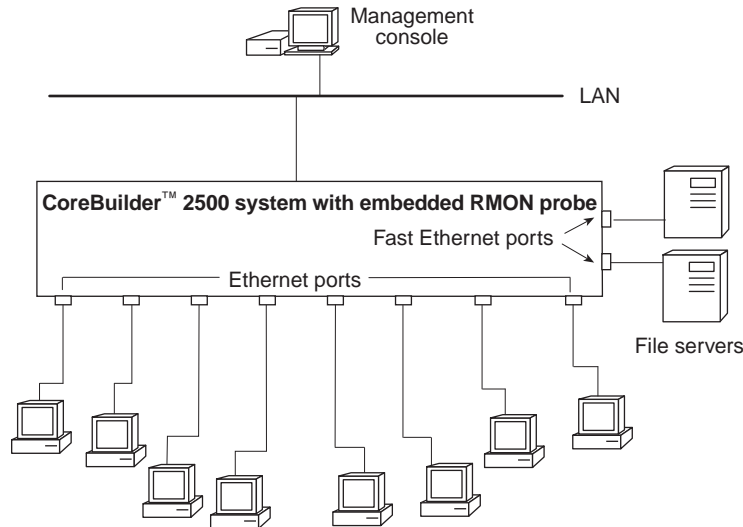


Figure 15-2 Embedded RMON Implemented on the CoreBuilder 2500 System

RMON Management Information Base (MIB)

A MIB is a structured set of data that describes the way that the network is functioning. The management software, known as the *agent*, gains access to this set of data and extracts the information it needs. The agent can also store data in the MIB.

The organization of a MIB allows a network management package based on the Simple Network Management Protocol (SNMP), such as the Transcend Enterprise Manager application suite, to manage a network device without having a specific description of that device. 3Com ships SNMP MIB files with CoreBuilder 2500 Extended Switching System software as ASN.1 files.

MIB Objects

The data in the MIB consists of objects that represent features of the equipment that an agent can control and manage. Examples of objects in the MIB include a port that you can enable or disable and a counter that you can read.

A counter is a common type of MIB object used by RMON. A counter object may record the number of frames that are transmitted onto the network. The MIB may contain an entry for the counter object something like the one in Figure 15-3.

```
etherStatsPkts OBJECT-TYPE
    SYNTAX      Counter
    ACCESS      read-only
    STATUS      mandatory
    DESCRIPTION
        This is a total number of packets
        received, including bad packets,
        broadcast packets, and multicast
        packets.
 ::= { etherStatsEntry 5 }
```


Figure 15-3 Example of an RMON MIB Counter Object

The counter object information includes these items:

- The name of the counter. In Figure 15-3, the counter is called *etherStatsPkts* (Ethernet, Statistics, Packets).
- Access level. In Figure 15-3, access is read-only.
- The number of the counter's column in the table. In Figure 15-3, the counter is in column 5 of the *etherStatsEntry* table.


The name of the table where the counter resides is *3CometherStatTable*, although this name does not appear in the display.


To manage a network, you do not need to know the contents of every MIB object. Most network management applications, including Transcend Enterprise Manager software, make the MIB transparent. However, by knowing how different management features are derived from the MIB you can better understand how to use the information they provide.



APPENDIX

Appendix A Technical Support





TECHNICAL SUPPORT

3Com provides easy access to technical support information through a variety of services. This appendix describes these services.

Information contained in this appendix is correct at time of publication. For the very latest, 3Com recommends that you access the 3Com Corporation World Wide Web site.

Online Technical Services

3Com offers worldwide product support 24 hours a day, 7 days a week, through the following online systems:

- World Wide Web site
- 3Com FTP site
- 3Com Bulletin Board Service (3Com BBS)
- 3ComFactsSM automated fax service

World Wide Web Site

Access the latest networking information on the 3Com Corporation World Wide Web site by entering the URL into your Internet browser:

<http://www.3com.com/>

This service provides access to online support information such as technical documentation and software library, as well as support options ranging from technical education to maintenance and professional services.

3Com FTP Site Download drivers, patches, and software across the Internet from the 3Com public FTP site. This service is available 24 hours a day, 7 days a week.

To connect to the 3Com FTP site, enter the following information into your FTP client:

- Hostname: **ftp.3com.com** (or **192.156.136.12**)
- Username: **anonymous**
- Password: **<your Internet e-mail address>**



A user name and password are not needed with Web browser software such as Netscape Navigator and Internet Explorer.

3Com Bulletin Board Service

The 3Com BBS contains patches, software, and drivers for 3Com products. This service is available through analog modem or digital modem (ISDN) 24 hours a day, 7 days a week.

Access by Analog Modem

To reach the service by modem, set your modem to 8 data bits, no parity, and 1 stop bit. Call the telephone number nearest you:

Country	Data Rate	Telephone Number
Australia	Up to 14,400 bps	61 2 9955 2073
Brazil	Up to 14,400 bps	55 11 5181 9666
France	Up to 14,400 bps	33 1 6986 6954
Germany	Up to 28,800 bps	4989 62732 188
Hong Kong	Up to 14,400 bps	852 2537 5601
Italy	Up to 14,400 bps	39 2 27300680
Japan	Up to 14,400 bps	81 3 3345 7266
Mexico	Up to 28,800 bps	52 5 520 7835
P.R. of China	Up to 14,400 bps	86 10 684 92351
Taiwan, R.O.C.	Up to 14,400 bps	886 2 377 5840
U.K.	Up to 28,800 bps	44 1442 438278
U.S.A.	Up to 28,800 bps	1 408 980 8204

Access by Digital Modem

ISDN users can dial in to the 3Com BBS using a digital modem for fast access up to 56 Kbps. To access the 3Com BBS using ISDN, use the following number:

1 408 654 2703

3ComFacts Automated Fax Service

The 3ComFacts automated fax service provides technical articles, diagrams, and troubleshooting instructions on 3Com products 24 hours a day, 7 days a week.

Call 3ComFacts using your Touch-Tone telephone:

1 408 727 7021

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:

- Product model name, part number, and serial number
- A list of system hardware and software, including revision levels
- Diagnostic error messages
- 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 obtain assistance from the 3Com online technical resources or from your network supplier, 3Com offers technical telephone support services. To find out more about your support options, please call the 3Com technical telephone support phone number at the location nearest you.

When you contact 3Com for assistance, have the following information ready:

- Product model name, part number, and serial number
- A list of system hardware and software, including revision levels
- Diagnostic error messages
- Details about recent configuration changes, if applicable

Below is a list of worldwide technical telephone support numbers:

Country	Telephone Number	Country	Telephone Number
Asia Pacific Rim			
Australia	1 800 678 515	New Zealand	0800 446 398
China		Singapore	800 6161 463
From anywhere in China:	86 21 6350 1590	S. Korea	
From Shanghai:	10 800 3656	From anywhere in S. Korea:	82 2 3455 6455
Hong Kong	800 933 486	From Seoul:	00798 611 2230
India	61 2 9937 5085	Taiwan	0080 611 261
Indonesia	001 800 61 009	Thailand	001 800 611 2000
Japan	0031 61 6439	Pakistan	61 2 9937 5085
Malaysia	1800 801 777	Philippines	1235 61 266 2602
Europe			
From anywhere in Europe, call:	+31 (0)30 6029900 phone +31 (0)30 6029999 fax		
From the following European countries, you may use the toll-free numbers:			
Austria	06 607468	Netherlands	0800 0227788
Belgium	0800 71429	Norway	800 11376
Denmark	800 17309	Poland	0800 3111206
Finland	0800 113153	Portugal	05 05313416
France	0800 917959	South Africa	0800 995014
Germany	0130 821502	Spain	900 983125
Hungary	00800 12813	Sweden	020 795482
Ireland	1 800 553117	Switzerland	0800 55 3072
Israel	177 3103794	U.K.	0800 966197
Italy	1678 79489		
Latin America			
Argentina	541 312 3266	Colombia	571 629 4847
Brazil	55 11 523 2725, ext. 422	Mexico	01 800 849 2273
North America			
	1 800 NET 3Com (1 800 638 3266)		

Returning Products for Repair

Before you send a product directly to 3Com for repair, you must first obtain a Return Materials Authorization (RMA) number. Products sent to 3Com without RMA numbers 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
Asia, Pacific Rim	65 543 6342	65 543 6348
Europe, South Africa, and Middle East	011 44 1442 435860	011 44 1442 435718
From the following European countries, you may call the toll-free numbers; select option 2 and then option 2:		
Austria	06 607468	
Belgium	0800 71429	
Denmark	800 17309	
Finland	0800 113153	
France	0800 917959	
Germany	0130 821502	
Hungary	00800 12813	
Ireland	1 800 553117	
Israel	177 3103794	
Italy	1678 79489	
Netherlands	0800 0227788	
Norway	800 11376	
Poland	0800 3111206	
Portugal	05 05313416	
South Africa	0800 995014	
Spain	900 983125	
Sweden	020 795482	
Switzerland	0800 55 3072	
U.K.	0800 966197	
Latin America	1 408 326 2927	1 408 764 6883
U.S.A. and Canada	1 800 876 3266, option 2	1 408 764 7120

INDEX

Numbers

3Com bulletin board service (3Com BBS) A-2
3Com URL A-1
3ComFacts A-3

A

AARP (AppleTalk Address Resolution Protocol) 8-10

cache

administering 14-6
displaying 14-6
flushing all entries from 14-8
removing an entry from 14-7

address

classes 4-3
IP 10-2
IP to MAC, translating 10-11
MAC 3-3
network 3-3
RIP advertisement 10-21

Address Resolution Protocol. *See* ARP

adjacencies, OSPF 7-12

administering OSPF areas

defining 13-2
displaying 13-2
modifying 13-3
removing 13-3

administering OSPF default route metrics

defining 13-6
displaying 13-5
removing 13-6
setting 13-5

administering OSPF interfaces

configuring 13-6

displaying
configuration 13-7
statistics 13-8

setting

area ID 13-11
cost 13-12
dead interval 13-14
hello timer 13-13
mode 13-10
password 13-14
priority 13-11

retransmit timer 13-13

transmit delay 13-12

administering OSPF link state databases

configuring 13-15

displaying

external network LSAs 13-20

network LSAs 13-18

router LSAs 13-16

summary 13-15

summary network LSAs 13-19

administering OSPF memory partitions

displaying 13-24

modifying 13-24

administering OSPF neighbors

adding 13-22

displaying 13-21

removing 13-22

administering OSPF network ranges

adding 13-4

modifying 13-4

removing 13-5

administering OSPF router IDs, setting 13-23

administering OSPF statistics, displaying 13-28

administering OSPF stub default metrics 13-25

defining 13-25

displaying 13-25

removing 13-25

administering OSPF virtual links

defining 13-26

displaying 13-26

modifying 13-27

removing 13-27

Administration Console, menu descriptions 1-2

ADSP (AppleTalk Data Stream Protocol) 8-8

AEP (AppleTalk Echo Protocol) 8-8

alarm thresholds, RMON

examples of 15-7

setting 15-6

AppleTalk

address resolution protocol (AARP) 8-10

checksum 14-10

configuring forwarding 14-10

data stream protocol (ADSP) 8-8

DDP statistics 14-11

echo protocol (AEP) 8-8

- interface
 - defining 14-3
 - displaying 14-2
 - parts of 14-2
 - removing 14-4
 - main menu 1-8
 - name binding protocol (NBP) 8-8
 - NBP statistics 14-16
 - network layer 8-5
 - networks 8-2
 - nodes 8-2
 - pinging 14-11
 - physical layer 8-5
 - printer access protocol (PAP) 8-9
 - protocols
 - about 8-1
 - and OSI levels 8-4
 - routes
 - administering 14-4
 - displaying 14-4
 - flushing 14-5
 - routing 8-1
 - routing table maintenance protocol (RTMP) 8-6
 - routing tables 8-7
 - RTMP statistics 14-13
 - session layer protocol (ASP) 8-9
 - statistics 14-11
 - transaction protocol (ATP) 8-8
 - ZIP statistics 14-14
 - zone information protocol (ZIP) 8-8
 - zone information table (ZIT) 8-8
 - zones 8-2, 14-8
- area border routers, OSPF 7-4, 7-8
 - areas, OSPF 7-3, 7-5
 - administering 13-1
 - defining 13-2
 - displaying 13-2
 - modifying 13-3
 - network ranges
 - adding 13-4
 - modifying 13-4
 - removing 13-5
 - removing 13-3
 - ARP (Address Resolution Protocol)
 - cache 4-7, 10-11
 - displaying 10-11
 - flushing 10-12
 - removing an entry from 10-12
 - defined 4-7, 10-11
 - location in OSI Reference Model 4-1
 - reply 4-8
 - request 4-8
 - ASP (AppleTalk Session Layer Protocol) 8-9
 - ATM ARP cache
 - displaying 10-14
 - flushing 10-15
 - removing an entry 10-15
 - ATM ARP servers
 - about 4-10
 - defining
 - server 10-13
 - nodes that can function as an 4-10
 - ATP (AppleTalk Transaction Protocol) 8-8
 - autonomous system routers, OSPF 7-8
 - autonomous systems, OSPF 7-3
-
- B**
- backbone areas, OSPF 7-4
 - backbone routers, OSPF 7-7
 - backup designated routers, OSPF 7-8, 7-12
 - BOOTP
 - hop count limit 10-17
 - relay threshold 10-17
 - bridge menu 1-3, 1-4
 - bridging/routing
 - CoreBuilder 2500 model 3-4
 - traditional model 3-4
 - bulletin board service A-2
-
- C**
- cache
 - AARP 14-6 to 14-8
 - ARP 4-7, 10-11
 - displaying the IP multicast 11-8
 - checksum, configuring AppleTalk 14-10
 - Chooser, Macintosh 8-2
 - conventions
 - notice icons, About This Guide 2
 - text, About This Guide 3
 - CoreBuilder
 - ports and IP interfaces 10-4
 - CoreBuilder 2500
 - bridging/routing model 3-6
 - intranetwork router, as an 3-2
 - subnetting with 3-2
 - cost
 - of OSPF interface 7-10, 13-12
 - of RIP interface 10-20, 10-21
 - See also metric
-
- D**
- data link layer 4-1
 - database description packets, OSPF 7-7
 - datagram delivery protocol 8-5
 - DDP statistics 14-11

default route metric, OSPF
 defining 13-6
 displaying 13-5
 removing 13-6
 setting 13-5

default route, IP
 defined 4-6, 10-8
 removing 10-11
 setting 10-10

default VLAN 2-5
 defining 10-5

defining an ATM ARP 10-13
 ATM ARP server 10-5

designated routers, OSPF 7-8, 7-13

direct, route status 10-8

documentation
 user comments, About This Guide 5

DVMRP (Distance Vector Multicast Routing Protocol)
 about 5-2
 enabling 11-1
 metric value 5-5, 11-3

dynamic routes 4-6
 See also RIP *and* OSPF *and* SAP

dynamic routes, IPX 6-9, 6-13

E

enhanced RIP mode, IPX
 setting 12-9

event group, RMON 15-8

extended network numbers 8-2

extended switching, overview 1-1

F

fax service (3ComFacts) A-3

flooding, IP multicast routing 5-3

flushing
 ARP cache 10-12
 learned routes
 IP 10-10
 IPX 12-6

forwarding
 configuring AppleTalk 14-10
 setting IPX 12-8

G

gateway
 IP address 10-8
 routing table, and the 4-5

H

hello packets, OSPF 7-6, 7-7, 7-12

host group, RMON 15-8

hostTopN group, RMON 15-8

hysteresis mechanism, RMON 15-7

I

ICMP (Internet Control Message Protocol)
 defined 4-8
 echo (request and reply) 10-23
 echo reply 4-8
 echo request 4-8
 location in OSI Reference Model 4-1
 ping and 10-23
 redirect 4-8
 time exceeded 4-8

ICMP Router Discovery, enabling 10-19

ICMP statistics 10-27

IGMP (Internet Group Management Protocol)
 about 5-2
 enabling 11-2

interface
 defining an IP 10-4

interface, AppleTalk
 defining 14-3
 displaying 14-2
 parts of 14-2
 removing 14-4

interface, IP
 defining a LIS 10-5
 defining a VLAN 10-4
 displaying 10-3
 parts of 10-2
 removing definition 10-6

interface, IP multicast
 administering 11-3
 characteristics 5-5
 defining 11-1
 disabling 11-4
 displaying 11-3
 enabling 11-4
 parts of 11-1

interface, IPX
 defining 12-3
 displaying 12-3
 modifying 12-4
 removing 12-4

interface, OSPF
 and OSPF areas 7-4
 area ID 7-10
 area ID, setting 13-11
 configuration, displaying 13-7
 configuring 13-6

- cost 7-10
- cost, setting 13-12
- dead interval 7-11
- dead interval, setting 13-14
- delay 7-10
- hello timer 7-10
- hello timer, setting 13-13
- mode 7-9
- mode, setting 13-10
- parts of 7-9
- password 7-11
- password, setting 13-14
- priority 7-9
- priority, setting 13-11
- retransmit timer 7-10
- retransmit timer, setting 13-13
- state 7-12
- statistics, displaying 13-8
- transmit delay, setting 13-12
- Interior Gateway Protocols (IGPs) 4-6, 6-9
- internal routers, OSPF 7-7
- Internet address. See IP address
- Internet Control Message Protocol. See ICMP
- Internet Protocol. See references with IP address
- intranetwork routing, diagram 3-2
- IP
 - statistics
 - ICMP 10-27
 - UDP 10-27
- IP (Internet Protocol)
 - address translation 10-11
 - ARP cache 10-11
 - enabling forwarding 10-18
 - interface 10-2
 - interface. See also IP interface
 - main menu 1-4
 - pinging a station 10-23
 - RIP interface configuration 10-19
 - RIP mode 10-20
 - RIP poison reverse mode 10-21
 - routes 10-8
 - statistics, displaying 10-24
- IP address
 - classes of 4-3
 - defined 4-2
 - derived from 4-2
 - division of network and host 4-2
 - example 4-4
 - for IP interface 10-2
 - network layer and the 4-1
 - OSPF, and 4-6
 - RIP, and 4-6
 - routing table, and the 4-5
 - subnet mask, and the 4-3
 - subnetwork part 4-3
 - IP forwarding, configuring 10-18
 - IP interface
 - address 10-2
 - cost 10-20, 10-21
 - defining 10-4
 - displaying 10-3
 - overlapped 10-18
 - removing definition 10-6
 - subnet mask 10-2
 - IP multicast
 - cache, displaying 11-8
 - routes, displaying 11-7
 - tunnels 11-5
 - defining 11-6
 - displaying 11-5
 - removing 11-6
 - IP multicast routing
 - about 5-1
 - algorithms 5-3
 - DVMRP metric value 5-5
 - flooding 5-3
 - interfaces 11-1
 - administering 11-3
 - defining 11-1
 - disabling 11-4
 - displaying 11-3
 - enabling 11-4
 - MBONE 5-2
 - pruning 5-4
 - rate limit 5-5, 11-3
 - reverse path forwarding 5-4
 - spanning tree 5-3
 - TTL threshold 5-5, 11-3
 - tunnels 5-5
 - IP route
 - default 10-8, 10-10
 - defining static 10-9
 - displaying table 10-9
 - gateway IP address 10-8
 - metric 10-8
 - removing from table 10-10
 - status 10-8
 - IP router, transmission process 4-2
 - IP routing
 - address classes 4-3
 - basic elements 4-2
 - ICMP 4-8
 - OSI reference model 4-1
 - references 4-11
 - router interface 4-4
 - routing table 4-5
 - transmission errors 4-8
 - IP routing over ATM 4-9

- IPX
 - enhanced RIP mode, setting 12-9
 - forwarding
 - setting 12-8
 - statistics 12-14
 - interface
 - defining 12-3
 - displaying 12-3
 - modifying 12-4
 - removing 12-4
 - main menu 1-7
 - RIP mode, setting 12-9
 - RIP statistics 12-12
 - RIP triggered updates, setting 12-10
 - route
 - administering 12-4
 - defining a static 12-5
 - removing 12-6
 - routing table
 - flushing learned routes 12-6
 - removing a route 12-6
 - SAP mode, setting 12-10
 - SAP statistics 12-13
 - SAP triggered updates, setting 12-11
 - server, flushing 12-8
 - server, removing 12-8
 - static server, defining 12-7
 - static server, displaying 12-7
 - summary statistics 12-11
 - IPX forwarding statistics, displaying 12-14
 - IPX RIP statistics, displaying 12-12
 - IPX routing
 - and RIP 6-9
 - and SAP 6-10
 - dynamic routes 6-9
 - packet format 6-5
 - router interface 6-7
 - routing table 6-8
 - server table 6-12
 - static routes 6-9
 - IPX SAP statistics, displaying 12-13
 - IPX summary statistics, displaying 12-11
-
- L**
 - layer 3 addressing 2-4
 - learned routes
 - flushing IP 10-10
 - flushing IPX 12-6
 - learned, IP route status 10-8
 - link state
 - acknowledge packets, OSPF 7-7
 - advertisements (LSAs), OSPF 7-3, 7-5, 7-7, 7-8, 7-10, 7-12, 7-13
 - database, OSPF
 - displaying 13-15
 - external network LSAs, displaying 13-20
 - network LSA summary, displaying 13-19
 - network LSAs, displaying 13-18
 - router LSAs, displaying 13-16
 - summary, displaying 13-15
 - protocol, OSPF 7-1
 - request packets, OSPF 7-7
 - update packets, OSPF 7-7
 - LIS
 - definition of 4-9
 - forwarding to nodes within an 4-10
 - LIS interfaces
 - characteristics of 10-3
 - defining 10-5
-
- M**
 - MAC address 3-3
 - ARP and 10-11
 - bridging in switching modules, and 3-6
 - compared to IP address 4-2
 - in ARP Request 4-7
 - located with ARP 4-7
 - use in IP routing 4-8
 - Macintosh, Chooser 8-2
 - management
 - IP interface 10-2
 - management console, RMON 15-1
 - matrix group, RMON 15-8
 - MBONE 5-2
 - memory partitions, OSPF
 - administering 13-24
 - displaying 13-24
 - modifying 13-24
 - menu
 - AppleTalk main 1-8
 - bridge 1-4
 - IP main 1-4
 - IPX main 1-7
 - metric
 - defined 4-5
 - DVMRP value 5-5, 11-3
 - in IP routing table 10-8
 - MIB (Management Information Base)
 - RMON 15-1, 15-2, 15-10
 - multicast routing, IP
 - about 5-1
-
- N**
 - name binding protocol (NBP) 8-8
 - statistics 14-16
 - named entities 8-2

neighbors, OSPF 7-6, 7-7, 7-10, 7-11, 7-12, 7-13
 adding 13-22
 administering 13-21
 displaying 13-21
 removing 13-22

NetWare
 defined 6-1
 OSI Reference Model, and the 6-2
 protocols 6-1 to 6-3

network address 3-3

network layer
 and IP address 4-1
 AppleTalk 8-5

network numbers
 extended 8-2
 nonextended 8-2

network supplier support A-3

nodes, AppleTalk 8-2

nonextended network numbers 8-2

O

online technical services A-1

OSI Reference Model
 AppleTalk routing and 8-4
 IP routing and 4-1
 IPX routing and 6-2

OSPF
 agencies 7-12
 administering. *See* administering OSPF
 area ID 7-10
 area parameters 7-5
 areas 7-3
 autonomous systems 7-3
 backbone areas 7-4
 backbone routers 7-7
 cost 7-10
 database description packets 7-7
 dead interval 7-11
 defined 4-6
 delay 7-10
 hello packets 7-6, 7-7, 7-12
 hello timer 7-10
 illustration 7-1
 interface state 7-12
 interface, parts of 7-9
 interfaces 7-4
 link state
 acknowledge packets 7-7
 advertisements (LSAs) 7-3, 7-5, 7-7, 7-8,
 7-10, 7-12, 7-13
 protocol 7-1
 request packets 7-7
 update packets 7-7

location in OSI Reference Model 4-1
 mode 7-9

neighbors 7-6, 7-7, 7-10, 7-11, 7-12, 7-13
 password 7-11
 path trees, shortest 7-13
 priority 7-9
 protocol packets 7-7
 retransmit timer 7-10
 router databases 7-5
 router IDs 7-8
 routers
 area border 7-4, 7-8
 autonomous system 7-8
 backbone 7-7
 backup designated 7-8, 7-12
 designated 7-8, 7-13
 internal 7-7

routing
 inter-area 7-14
 intra-area 7-14
 to different autonomous system 7-14
 to stub area 7-14
 shortest path trees 7-13
 stub areas 7-3
 stub default metric 7-11
 summary 7-1
 transit areas 7-3
 virtual links 7-4, 7-6, 7-7, 7-11, 7-14

overlapped IP interfaces 10-18

P

PAP (AppleTalk Printer Access Protocol) 8-9

physical layer, AppleTalk 8-5

pinging
 AppleTalk node 14-11
 IP station 10-23

poison reverse mode, RIP 10-21

printer access protocol (PAP) 8-9

probe, RMON 15-1, 15-2

protocol packets, OSPF 7-7

protocol, AppleTalk routing table maintenance 8-6

protocols, AppleTalk 8-6 to 8-9

pruning, IP multicast routing 5-4

PVC
 adding 10-7
 removing 10-7

R

rate limit, IP multicast 5-5, 11-3

references

Comer 4-11

IP 4-11

OSPF 4-11

Perlman 4-11

RIP 4-11

routing RFCs 4-11

Sterns 4-11

returning products for repair A-5

reverse path forwarding algorithm

IP multicast routing 5-4

RIP

broadcast address, and 10-21

RIP (Routing Information Protocol)

active mode 10-20

adding advertisement address 10-21

default mode 10-20

defined 4-6, 6-9

displaying interface configuration 10-19

displaying statistics 10-22

interface cost 10-20, 10-21

location in OSI Reference Model 4-1

off mode 10-20

passive mode 10-20

removing advertisement address 10-22

route configuration, and 4-6, 6-9

setting interface cost 10-20, 10-21

setting mode 10-20

setting poison reverse mode 10-21

using for dynamic routes 6-9

RIP mode, IPX

setting 12-9

RIP triggered updates, IPX

setting 12-10

RMON (Remote Monitoring)

agents 15-9

alarms 15-5, 15-6

benefits of 15-2

event group 15-8

host group 15-8

hostTopN group 15-8

hysteresis mechanism 15-7

management console 15-1

matrix group 15-8

MIB 15-1, 15-2, 15-10

probe 15-1, 15-2

statistics 15-4, 15-5

route, AppleTalk

administering 14-4

displaying 14-4

flushing 14-5

route, IP

default 10-8

defining static 10-9

gateway address 10-8

metric 10-8

removing default 10-11

removing from table 10-10

status 10-8

subnet mask 10-8

route, IPX

administering 12-4

removing 12-6

router databases, OSPF 7-5

router ID, OSPF 7-8

router interface, IP

described 4-4

diagram 4-4

router interface, IPX 6-7

routers

area border, OSPF 7-4, 7-8

autonomous system, OSPF 7-8

backbone, OSPF 7-7

backup designated, OSPF 7-8, 7-12

databases, OSPF 7-5, 13-15

designated, OSPF 7-8, 7-13

IDs, OSPF 7-8

internal, OSPF 7-7

seed 8-4

setting IDs, OSPF 13-23

routes, displaying IP multicast 11-7

routing

and bridging in switching modules 3-4

and bridging, traditional model 3-4

AppleTalk 8-1

CoreBuilder 2500 system, and the 3-1 to 3-7

inter-area, OSPF 7-14

intra-area, OSPF 7-14

to different autonomous system, OSPF 7-14

to stub area, OSPF 7-14

See also IP routing, IPX routing, and AppleTalk

routing

Routing Information Protocol. See RIP

routing table, AppleTalk 8-7

routing table, IP

contents 4-5, 10-8

default route 4-6

default route, setting 10-10

described 4-5

display routes 10-9

dynamic routes 4-6

example 4-5

flushing learned routes 10-10

metric 4-5

removing default route 10-11

removing route 10-10

static routes 4-6

- routing table, IPX
 - contents 6-8
 - described 6-8
 - dynamic routes 6-9
 - example 6-9
 - flushing learned routes 12-6
 - removing a route 12-6
 - static routes 6-9
- RTMP (AppleTalk Routing Table Maintenance Protocol)
 - description of 8-6
 - statistics 14-13

S

- SAP (Service Advertising Protocol)
 - aging mechanism 6-14
 - packet structure 6-11
 - request handling 6-14
 - using for dynamic routes 6-13
- SAP mode, IPX
 - setting 12-10
- SAP statistics, displaying 12-13
- SAP triggered updates, IPX
 - setting 12-11
- seed routers 8-4
- segmentation, increasing 3-3
- server 10-5, 10-13
 - defining a static IPX 12-7
 - displaying a static IPX 12-7
 - flushing IPX 12-8
 - removing IPX 12-8
- server table
 - contents 6-12
 - described 6-12
- server table, IPX
 - displaying 12-7
- server, IPX
 - flushing 12-8
 - removing 12-8
- Service Advertisement Protocol. *See* SAP
- session layer protocols, AppleTalk 8-8
- session protocol, AppleTalk (ASP) 8-9
- shortest path trees, OSPF 7-13
- spanning tree algorithm, IP multicast routing 5-3
- static route, IP 4-6
 - defining 10-9
 - status of 10-8
- static route, IPX 6-9, 6-13
 - defining 12-5
- static server, IPX
 - defining 12-7
- statistics
 - AppleTalk 14-11
 - AppleTalk DDP 14-11
 - AppleTalk NBP 14-16

- AppleTalk RTMP 14-13
- AppleTalk ZIP 14-14
- ICMP 10-27
- IP 10-24
- IPX forwarding 12-14
- IPX RIP 12-12
- IPX SAP 12-13
- IPX summary 12-11
- OSPF, displaying 13-28
- RIP 10-22
- RMON 15-4, 15-5
- UDP 10-27
- stub areas, OSPF 7-3
- stub default metrics, OSPF 7-11
 - administering 13-25
 - defining 13-25
 - displaying 13-25
 - removing 13-25
- subnet mask
 - defined 4-3
 - diagram 4-3
 - example 4-4
 - for IP address 10-2
 - in IP routing table 10-8
 - in routing table 4-5
- subnetting
 - defined 4-3
 - Ethernet switching and 3-2
 - subnet mask, and the 4-3
 - with the CoreBuilder 2500 3-2

T

- technical support
 - 3Com URL A-1
 - bulletin board service A-2
 - fax service A-3
 - network suppliers A-3
 - product repair A-5
- timing out, IP route status 10-8
- transit areas, OSPF 7-3
- transmission errors
 - ICMP echo reply 4-9
 - ICMP echo request 4-9
 - ICMP redirect 4-9
 - ICMP time exceeded 4-9
 - IP routing 4-8
 - reasons for 4-8
- TTL threshold, IP multicast 5-5, 11-3
- tunnels
 - defining IP multicast 11-6
 - displaying IP multicast 11-5
 - IP multicast 5-5, 11-5
 - removing IP multicast 11-6

U

- UDP Helper
 - and BOOTP 10-13
 - and DHCP 10-13
 - and routed networks 10-13
 - configuring overlapped IP interfaces 10-18
 - defining port and IP routing address 10-16
 - displaying configuration information 10-16
 - removing port and IP routing address 10-16
 - setting the BOOTP hop count limit 10-17
 - setting the BOOTP relay threshold 10-17
- UDP statistics 10-27
- URL A-1

V

- virtual links, OSPF 7-4, 7-6, 7-7, 7-11, 7-14
 - administering 13-26
 - defining 13-26
 - displaying 13-26
 - modifying 13-27
 - removing 13-27
- VLAN interfaces
 - about 10-2
 - characteristics of 10-2
 - defining 10-4
- VLANs
 - application-oriented 2-2
 - configuration examples 2-9
 - default 2-5
 - flooding within 2-5
 - information
 - defining 9-3
 - displaying 9-1
 - modifying 9-4
 - removing 9-5
 - layer 3 addressing 2-4
 - MAC address group 2-2
 - overlapped IP 2-7
 - port group 2-2
 - protocol-sensitive 2-3
 - routing between 2-8

W

- World Wide Web (WWW) A-1

Z

- zone information protocol (ZIP) 8-8
 - statistics 14-14
- zone information table (ZIT) 8-8
 - displaying the 14-8
- zone, AppleTalk
 - example of 8-2

