


Cisco Routers



© CAMPBELL LIND/SIS

address is 0010.> **Get productive with clear, accessible information.**< 7b3a.50b3 <bia 0010.7b3a.50b3> Internet add
us 130.10.64.1/19 MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, > **At a glance find key definitions.**< rely 255/255,
1/255 Encapsulation ARPA, ARP Timeout 04:00:00 Serial 0 is up, line protocol is up Hardware is HD64570 Processor B
ID 08867035, with hardware > **Configure routers with easy-to-follow instructions and diagrams.**< revision 00000000
software Version 2.0. NET2_BFE and GOSIP compliant 1 Ethernet/IEEE 802.3 interface(s) 8 Ethernet/IEEE 802.3-rend

que

Joe Habraken

PRACTICAL

Cisco Routers

Joe Habraken

Contents at a Glance

I Networking Overview

- 1 LAN Review 7
- 2 The OSI Model and Network Protocols 33
- 3 Wide Area Networking 53
- 4 Internetworking Basics 67
- 5 How a Router Works 77

II Router Design and Basic Configuration

- 6 Understanding Router Interfaces 99
- 7 Setting Up a New Router 111
- 8 Basic Router Configuration 123
- 9 Working with the Cisco IOS 141

III Routing LAN Protocols

- 10 TCP/IP Primer 167
- 11 Configuring IP Routing 195
- 12 Routing Novell IPX 211
- 13 Routing AppleTalk 227

IV Advanced Configuration and Configuration Tools

- 14 Filtering Router Traffic with Access Lists 243
- 15 Configuring WAN Protocols 259
- 16 Configuring the Router with Cisco ConfigMaker 271
- 17 Using a TFTP Server for Router Configuration Storage 289
- 18 Basic Router Troubleshooting 301

V Appendixes

- A Basic Router Command Summary 323
- B Selected Cisco Router Specifications 337
- Glossary 343

Index 359

que

A Division of Macmillan Computer Publishing, USA
201 W. 103rd Street
Indianapolis, Indiana 46290

Practical Cisco Routers

Copyright © 1999 by Que Corporation

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Neither is any liability assumed for damages resulting from the use of the information contained herein.

International Standard Book Number: 0-7897-2103-1

Library of Congress Catalog Card Number: 99-63284

Printed in the United States of America

First Printing: September 1999

01 00 99 4 3 2 1

Trademarks

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Que Corporation cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

Acquisitions Editor
Jenny Watson

Development Editor
Rick Kughen

Managing Editor
Lisa Wilson

Project Editor
Tonya Simpson

Copy Editor
Kate Givens

Indexer
Rebecca Salerno

Proofreader
Andy Beaster

Technical Editor
Ariel Silverstone

Interior Design
Anne Jones

Cover Design
Radar Design

Copy Writer
Eric Borgert

Layout Technicians
Stacey DeRome
Ayanna Lacey
Heather Miller

Contents

Introduction 1

About This Book 1

How This Book Is Organized 2

Who Should Use This Book 3

Conventions Used In This Book 3

I Networking Overview

1 LAN Review 7

The Advent of the PC 8

Networking PCs 8

Peer-to-Peer Networks 9

Server-Based Networks 10

Making the Connection 12

Network Interface Cards 13

Dealing with IRQs and I/O Ports 14

Network Cabling 17

Hubs, Repeaters, and MAUs 19

Understanding Network Topologies 20

Bus Network 21

Star Network 22

Ring Topology 23

Mesh Topology 25

Understanding Network Architectures 25

Ethernet 26

IBM Token Ring 28

FDDI 29

AppleTalk 30

2 The OSI Model and Network Protocols 33

OSI—The Theoretical Networking Protocol Stack 34

The OSI Layers 35

The Application Layer 38

The Presentation Layer 38

The Session Layer 38

The Transport Layer 40

The Network Layer 40

The Data-Link Layer 41

The Physical Layer 43

The Data-Link Sublayers 43

Real-World Network Protocols 44

NetBEUI 45

TCP/IP 45

IPX/SPX 48

AppleTalk 49

3 Wide Area Networking 53

Understanding Wide Area Connectivity 54

Getting Connected 54

Dial-Up Connections 55

Leased Lines 56

Switched Network Overview 59

Circuit Switching 60

Packet Switching 61

WAN Packet Switching Protocols 62

X.25 62

Frame Relay 64

Asynchronous Transfer Mode (ATM) 64

Other WAN Protocols 65

4 Internetworking Basics 67

What Is Internetworking? 68

Internetworking Devices 68

Repeater 70

Bridges 71

Switches 73

Routers 73

Gateways 74

Building a Campus Network 75

5 How a Router Works 77

Routing Basics 78

Path Determination 78

Logical and Hardware Addresses 80

Packet Switching 81

Routing Tables 82

Routable Protocols 85

Routing Protocols 85

Routing Protocol Basics 87

Routing Algorithms 87

Routing Metrics 89

Types of Routing Protocols 91

Interior Gateway Protocols 93

Exterior Gateway Protocols 95

II Router Design and Basic Configuration

6 Understanding Router Interfaces 99

Router Interfaces 100

LAN Interfaces 102

Serial Interfaces 104

Logical Interfaces 108

Loopback Interfaces 108

Null Interfaces 109

Tunnel Interfaces 109

7 Setting Up a New Router 111

Becoming Familiar with Your Router 112

Cisco Router Design 113

Router CPUs 113

Router Memory Components 113

Connecting the Console 115

Configuring the Router Console 117

Working with the Terminal Emulation Software 118

Connecting the Router to the Network 119

LAN Connections 119

Serial Connections 121

A Final Word on Physical Router Connections 122

8 Basic Router Configuration 123

Configuring a Router 124

Router Boot Sequence 126

Working with the System Configuration Dialog Box 128

Starting the Setup Dialog Box 129

Configuring Routed Protocols 131

Configuring Router Interfaces 132

Using the Different Router Modes 134

User (Unprivileged) Mode 135

Privileged Mode 136

Configuration Mode 137

Getting Around Lost Passwords 139

CONTENTS

9 Working with the Cisco IOS 141

Introducing the Internetworking Operating System 142

Command Structure 144

Exec Commands 144

Configuration Mode 145

The IOS Help System 147

Router Examination Commands 149

Using the Privileged Mode 153

Checking Router Memory 154

Checking Out the Internetwork Neighborhood 157

Working with CDP 157

Viewing CDP Neighbors 159

Using Ping 160

Creating a Router Banner 161

Creating the Network Subnet Mask 184

Calculating IP Subnet Ranges 186

Calculating Available Node Addresses 188

Creating Class B and Class C Subnets 188

Class B Subnetting 188

Class C Subnetting 190

Understanding Subnet 0 192

A Final Word on Subnetting 194

11 Configuring IP Routing 195

Configuring Router Interfaces 196

LAN Interfaces 198

WAN Interfaces 200

Configuring a Routing Protocol 201

Configuring RIP 202

Configuring IGRP 204

Dynamic Routing Versus Static Routing 207

Using Telnet 209

12 Routing Novell IPX 211

Introducing IPX/SPX 212

Routing-Related IPX/SPX Protocols 213

Understanding IPX Addressing 214

Understanding SAP 216

Configuring IPX Routing 217

Configuring Router Interfaces with IPX 219

LAN Interfaces 220

WAN Interfaces 222

Monitoring IPX Routing 223

13 Routing AppleTalk 227

Understanding AppleTalk 228

AppleTalk Addressing 229

AppleTalk Zones 232

III Routing LAN Protocols

10 TCP/IP Primer 167

The TCP/IP Protocol Stack 168

TCP/IP and the OSI Model 168

Application Layer 170

Host-to-Host Layer 171

Internet Layer 171

Network Access Layer 172

Working with IP Addresses 174

IP Classes 175

Binary Equivalents and First Octets 177

Basic Subnet Masks 178

Subnetting IP Addresses 180

Binary and Decimal Conversions 181

Creating Subnets on a Class A

Network 182

- Configuring AppleTalk Routing 232
 - Configuring LAN Interfaces* 235
 - Configuring WAN Interfaces* 236
- Monitoring AppleTalk Routing 237

IV Advanced Configuration and Configuration Tools

14 Filtering Router Traffic with Access List 243

- Understanding Access Lists 244
 - How Access Lists Work* 244
 - Building an Access List* 246

- Working with IP Access Lists 247
 - IP Wildcard Masks* 249
 - Creating the Access List* 252
 - Grouping the Access List to an Interface* 253

- Creating IPX Standard Access Lists 254

- Creating AppleTalk Standard Access Lists 256

15 Configuring WAN Protocols 259

- Understanding Serial and WAN Interfaces 260
- Configuring High-Level Data Link Control (HDLC) 261
- Configuring PPP 262
- Configuring X.25 263
- Configuring Frame Relay 265
- Configuring ISDN 268

16 Configuring the Router with Cisco ConfigMaker 271

- What Is Cisco ConfigMaker? 272
- Downloading ConfigMaker 272
- Installing ConfigMaker 273
- Designing Your Internetwork with ConfigMaker 274
 - Adding Devices* 276
 - Connecting LANs to Routers* 278
 - Connecting Routers to Routers* 281
- Delivering the Configuration to a Router 284

17 Using a TFTP Server for Router Configuration Storage 289

- What Is a TFTP Server? 290
 - Obtaining TFTP Software* 291
- Installing the Cisco TFTP Server Software 292
- Copying to the TFTP Server 294
- Copying from the TFTP Server 295
- Loading a New IOS from the TFTP Server 297

18 Basic Router Troubleshooting 301

- Troubleshooting Hardware Problems 302
 - Router Problems* 302
 - Other Hardware Problems* 305
 - Cabling Problems* 306
 - A Final Word on Hardware* 307
- Troubleshooting LAN Interfaces 307
 - Troubleshooting Ethernet with Show* 307
 - Troubleshooting Token Ring with Show* 309
- Troubleshooting WAN Interfaces 311

CONTENTS

Troubleshooting TCP/IP 313

Using ping 314

Using trace 315

Troubleshooting IPX 316

Troubleshooting AppleTalk 317

A Final Word on Troubleshooting 318

V Appendixes

A Basic Router Command Summary 323

Cisco IOS Command Summary 324

Router Examination Commands 324

Router Memory Commands 325

Password and Router Name Configuration
 Commands 326

Interface Configuration Commands 327

IP-Related Commands 328

IPX-Related Commands 330

AppleTalk-Related Commands 331

WAN-Related Commands 332

Troubleshooting Commands 334

Miscellaneous Commands 334

B Selected Cisco Router Specifications 337

Router Selection 338

Cisco 7500 Routers 338

Cisco 4500 Routers 339

Cisco 2500 Routers 340

Cisco 1000 Routers 341

A Final Note 342

Glossary 343

Index 359

About the Author

Joe Habraken is an information technology consultant and best-selling author whose publications include *The Complete Idiot's Guide to Microsoft Access 2000*, *Microsoft Office 2000 8-in-1*, *Easy Publisher 2000*, and *Sams Teach Yourself Microsoft Outlook 2000 in 10 Minutes*. Joe has a Masters degree from the American University in Washington, D.C. and over 12 years of experience as an educator, author, and consultant in the information technology field. Joe is a Microsoft Certified Professional and currently provides consulting services in the NT Server and internetworking arenas to companies and organizations. He also currently serves as the lead instructor for the Networking Technologies program at Globe College in St. Paul, Minnesota.

Dedication

To all the NSS students at Globe College.

Good luck with your careers, and thanks for staying awake in my Cisco class (even when I babbled excitedly about internetworking and routing technology).

Acknowledgments

Creating a book like this takes a real team effort, and this particular book was created by a team of incredibly dedicated professionals. I would like to thank Jenny Watson, our acquisitions editor, who worked very hard to assemble the team that made this book a reality and always made sure the right pieces ended up in the right places.

I would also like to thank Rick Kughen, who served as the development editor for this book and who came up with many great ideas for improving its content. He always asked the right questions and wasn't afraid to burn the midnight oil to get the job done.

Also a tip of the hat and a thanks to Ariel Silverstone, who as the technical editor for the project did a fantastic job making sure that everything was correct and suggested several additions that made the book even more technically sound. Finally, a great big thanks to our production editor, Tonya Simpson, who ran the last leg of the race and made sure the book made it to press on time—what a great team of professionals.

Tell Us What You Think!

As the reader of this book, *you* are our most important critic and commentator. We value your opinion and want to know what we're doing right, what we could do better, what areas you'd like to see us publish in, and any other words of wisdom you're willing to pass our way.

As an Associate Publisher for Que Corporation, I welcome your comments. You can fax, email, or write me directly to let me know what you did or didn't like about this book—as well as what we can do to make our books stronger.

Please note that I cannot help you with technical problems related to the topic of this book, and that due to the high volume of mail I receive, I might not be able to reply to every message.

When you write, please be sure to include this book's title and author as well as your name and phone or fax number. I will carefully review your comments and share them with the author and editors who worked on the book.

Fax: 317-581-4666

Email: hardware@mcp.com

Mail: Jim Minatel
Associate Publisher
Que Corporation
201 West 103rd Street
Indianapolis, IN 46290 USA

introduction

■ find it amazing how rapidly computer technology has changed over the last 10 years. Technology once considered too costly or too complex for small or medium-sized companies is now being embraced at breakneck speed. Internetworking devices, and routers in particular, are some of the former “big-company” technologies now being used by even the smallest companies.

Inexpensive, low-end routers provide the connection to service providers and the public switched telephone network for small companies (and even individuals) who are looking for more bandwidth as they increasingly use the Internet as a communication and marketing tool. And as companies grow, they also look for strategies to conserve the bandwidth on their company-owned LANs; LAN segmentation with routers has become a viable and cost-effective solution.

With this explosion of internetworking technology hitting the business world, there has been a growing need for professionals to configure, manage, and troubleshoot routers and other internetworking devices. And although several excellent books and training materials that relate to internetworking and Cisco products are available, most of these materials have been written for IT professionals with many years of experience or training already under their belts. A basic primer and entry-level book on the subject really hasn't been available—until now.

About This Book

When I sat down to write this book, I wanted to do two things: share my excitement about internetworking and Cisco router configuration and provide a book that someone new to this technology could use to explore the incredible possibilities this technology offers. I also wanted to create a solid learning tool and make the book useful as a reference for someone with little internetworking background, who suddenly found working with Cisco routers part of their job description. And although that sounds like somewhat of a tall order, I knew that I would have help.

Skilled designers and editors at Macmillan Publishing have worked very hard to create a book design that embraces fresh ideas and approaches that will provide an environment in which you can get the information you need quickly and efficiently. You will find that this book embraces a streamlined, conversational approach to the subject matter that will help you learn the concepts and become familiar with the hardware and software facts that you need to get the job done.

How This Book Is Organized

- Part I, "Networking Overview"—This section of the book helps you get up to speed or review several networking technologies. Information is provided on LANs, WANs, and internetworking. A chapter also provides information on the Open System Interconnection reference model and how it relates to real-world network protocols. The basics on how routers work is also included in this section.
- Part II, "Router Design and Basic Configuration"—This section walks you through the hardware components of a typical Cisco router. You are also introduced to the basic configuration of routers and learn an overview of the Cisco Internetwork Operating System.
- Part III, "Routing LAN Protocols"—This section provides information about popular LAN protocols, such as TCP/IP, IPX/SPX, and AppleTalk. You learn conceptual information on each of these protocol stacks. You also walk through the steps of configuring a Cisco router for each of these protocols.
- Part IV, "Advanced Configuration and Configuration Tools"—This section helps you become familiar with several WAN technologies available and how they are configured on a Cisco router. Restricting access to your routers and troubleshooting routers are also covered to give you a complete picture of working with internetworking devices. Information on using Cisco's ConfigMaker router configuration software is also included in this section. It provides someone who must get a router con-

nected and configured in a hurry, a step-by-step look at how to use the ConfigMaker software.

Who Should Use This Book

This book is for anyone who needs a primer on internetworking and the configuration of Cisco routers. And whether you work for a big company, small company, or are just beginning your education to become a network professional, this book is an excellent first step as you build your knowledge base.

Conventions Used In This Book

Commands, directions, and explanations in this book are presented in the clearest format possible. The following items are some of the features that will make this book easier for you to use:

- Commands that you must enter—Router commands that you'll need to type are easily identified by a monospace font. For example, if I direct you to get the encapsulation (the WAN protocol set) for a serial interface, I'll display the command like this: `show interface serial 0`. This tells you that you'll need to enter this command exactly as it is shown.
- Combination and shortcut keystrokes—Text that directs you to hold down several keys simultaneously is connected with a plus sign (+), such as Ctrl+P.
- Cross references—If there's a related topic that is prerequisite to the section or steps you are reading, or a topic that builds further on what you are reading, you'll find the cross reference to it at the end of the section, like this:

SEE ALSO

► *To see how to create newspaper columns, see page xxx.*

- Glossary terms—For all the terms that appear in the glossary, you'll find the first appearance of that term in the text in *italic* along with its definition.

- Sidenotes—Information related to the task at hand, or “inside” information from the author, is offset in sidebars that don’t interfere with the task at hand. This valuable information is also easier to find. Each of these sidebars has a short title to help you quickly identify the information you’ll find there. You’ll find the same kind of information in these that you might find in notes, tips, or warnings in other books but here, the titles should be more informative.

part

NETWORKING OVERVIEW

LAN Review	7	1
The OSI Model and Network Protocols	33	2
Wide Area Networking	53	3
Internetworking Basics	67	4
How a Router Works	77	5

chapter

1

LAN Review

The Advent of the PC

Networking PCs

Making the Connection

Understanding Network Topologies

Understanding Network Architectures



The Advent of the PC

How and where people use computer technology has changed dramatically over the past 30 years. In the 1960s, computing revolved around large mainframe computers. In the early days, users typically interfaced with this highly centralized computer through an intermediary: an IS administrator or programmer. As computer technology evolved further, mainframe users were able to directly communicate with the computer using a dumb terminal (basically, a monitor and a keyboard hard-wired to the mainframe). In the 1970s, the miniframe gained dominance in the computing world, making computer technology accessible to a larger number of companies and organizations (even though these companies paid a premium for their ability to compute). All storage and computing power was still centralized, however, much the same as in the mainframe environment.

In the 1980s the personal computer (particularly the IBM Personal Computer) revolutionized the way you compute. Computing power was brought to the individual desktop. Not only was this new type of computer relatively easy to use (when compared to mainframes and miniframes) but also it was very affordable. The only flaw in this computing renaissance was the inability of users to collaborate and share resources. The individuality of the PC isolated its users.

Networking PCs

To overcome this decentralized computing model offered by the PC, software and hardware were developed in the 1980s and 1990s to connect PCs into networks that could share resources (such as printers and files). Networked PCs made it easy to design a collaborative computing environment for any business situation. Networked computers can share a variety of resources, including hardware (printers, modems), software (application software), and user-created files.

Different networking models arose to fit different types of networking needs. In situations where a few computers needed to share a particular hardware device, such as a printer, but did not require centralized file storage, the *peer-to-peer network* evolved. The only time individual users interfaced with this type of network was when they

printed. The alternative to the peer-to-peer network was a network with more centralized control of resources and better security. This type of network—a *server-based network*—uses a server computer (the central controller of the network) to authenticate users on the network and provide central file storage (as well as access to a number of different hardware and software resources). How these two networking models differ deserves some additional discussion.

Peer-to-Peer Networks

Peer-to-peer networks provide an easy way to share resources, such as files and printers, without the need for an actual server computer. Peer computers act as both *clients* (the users of resources) and *servers* (the providers of resources). The only real requirements for building a peer-to-peer network are installing an operating system on the PCs that supports peer-to-peer networking and physically connecting the PCs.

Several operating systems, such as Microsoft Windows 3.11, Microsoft Windows 95/98, and Microsoft Windows NT Workstation, have peer-to-peer networking capabilities built in. Local drives, folders, and printers can be shared with others on the peer-to-peer network (see Figure 1.1).

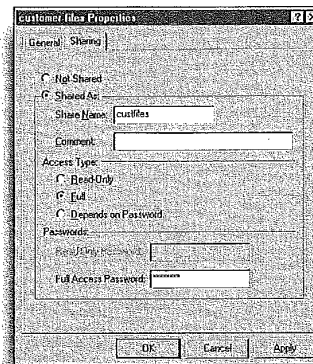


FIGURE 1.1
Operating systems such as Windows 98 make it easy for you to share resources on a peer-to-peer network.

When security is not the issue

If you are setting up a peer-to-peer network where security isn't an issue and all the users on the network are known to each other (and trust each other), you can choose not to assign a password to your shares—folders or drives set up for sharing on the network—or assign the same password to all of them. This takes some of the inconvenience out of sharing separate resources, but leaves resources wide open for use by anyone physically attached to the network.

Each resource that is shared (such as a drive or printer) potentially will have a different share password. This is one of the downsides of peer-to-peer networking—every resource is capable of having a separate password. If many resources are shared across the network, you will have to remember the password for each resource. This type of security is referred to as *share-level security*.

Peer-to-peer networks also don't require a great deal of additional administration because each user can manage resources on his own peer computer. Peer networks, however, do have their downsides:

- Increased performance hit on computers because of resource sharing
- No centralized location of shared files makes it difficult to back up data
- Security must be handled on a resource-by-resource level
- Decentralization of resources makes it difficult for users to locate particular resources
- Users might have to keep track of numerous passwords

Although peer-to-peer networking may seem like a fast and cheap way to connect a group of computers, the biggest drawback in using this type of networking is that only a small number of users can be accommodated. Peer networking isn't scalable (meaning expandable, because most peer networks are limited to 10 peer computers) and so is certainly not the appropriate choice for a growing company.

It is pretty much a consensus among IS managers that peer-to-peer networking works ideally with five or fewer peer machines.

SEE ALSO

➤ *For more information on the physical connections, see page 12.*

Server-Based Networks

Server-based networks provide greater centralized control of resources and expandability if required. A server computer is basically a special-purpose machine that logs in users and “serves” up resources to them. Because the server verifies users, this type of network makes it easier to manage your resources by providing different

access levels to the various users in your user pool. A username and one password puts users onto the network and gives them access to any resource for which they have the appropriate permissions.

A server-based network typically employs a more powerful (in terms of processor speed, RAM, and hard-drive capacity) computer to act as the server. In addition to hardware that can handle a large number of user requests for services, the server computer must run special software—a network operating system (NOS). Two commonly used network operating systems are Microsoft Windows NT Server and Novell NetWare.

Server-based networks, as mentioned before, are scalable. This means that the network can grow along with your company. Servers can be added to your network that take on specialized duties. For example, one server may handle user login and verification (a primary domain controller on a Windows NT network would be an example), while another server on the network may handle the email system (a communications server). Table 1.1 lists some of the specialized servers you might use on a local area network.

Table 1.1 LAN Server Types

Server Type	Use
File server	Stores shared user files and provides home directory space for users (such as a Novell NetWare server)
Communication server	Provides communication services such as email (such as an NT Server running Microsoft Exchange Server)
Application server	Provides access to a database or other application (such as an SQL server database)
Print server	Provides the print queue and other services related to a network printer

A server-based network of computers that is limited to a fairly small geographical area such as a particular building is described as a local area network (LAN). LANs are found in small, medium, and large companies. When several LANs are connected, you are dealing with an *internetwork*, which is a network of networks (this type of network can also be referred to as a *campus*). When you start connected campuses and create networks that span large geographical areas, you are working in the realm of the Wide Area Network (WAN).

Server-based networks are really the standard for even small local area networks; these types of networks do have their downside, however. Much of the downside, at least for the small company wanting to set up a PC network, is cost—the cost of at least one server PC and the cost of the network operating system. Server-based networks also typically require the hiring of a full-time administrator to maintain and manage the network (and whereas management sees this as an additional cost, the network administrator sees it as money well spent).

Other negatives associated with the server-based network revolve around server failures, *broadcast storms* (tons of broadcast traffic from devices on the network), and other hardware- and software-related disasters that are too numerous to mention in this book. Networks are by nature challenging, and that is why a good network administrator is worth his or her weight in gold.

SEE ALSO

➤ For more information on internetworking, see page 67.

SEE ALSO

➤ For more information on wide area networking see page 53.

Making the Connection

To create a computer network, you must use some type of connective medium that allows the transfer of your data. This medium can range from copper cable to microwave transmissions to a beam of infrared light (our discussion of network media will be restricted to copper and fiber-optic cables, with the understanding that there are a lot of possibilities for moving data from one point to another).

After you choose a connective medium, such as copper cable, you also need a device that can prepare the data on the computer so that it can travel along your network cabling. This data restructuring is handled by a network interface card (NIC). A NIC is typically placed in one of the computer's bus expansion slots and then the network cable is attached to a port on the NIC. Understanding how the NIC works, and your options as far as copper and fiber-optic cabling, will go a long way when you have to sit down and design even the smallest networks.

Network Interface Cards

The network interface card (NIC) provides the connection between the PC and the network's physical medium (such as copper or fiber-optic cable). Data travels in parallel on the PC's bus system; the network medium demands a serial transmission. The transceiver (a transmitter and receiver) on the NIC card is able to move data from parallel to serial and vice versa.

Network interface cards each have a unique address that is burned onto a ROM chip on each NIC. This addressing system is used to move data from one physical connection to another (and you will find that resolving logical addresses such as IP addresses to NIC hardware addresses is really what networking is all about).

NICs are available for a number of bus types (Figure 1.2 shows a PCI Ethernet NIC), so make it a point to open up the PC or PCs that you are going to network and check to see what type of bus slots are available. Newer PCs will typically have PCI slots available. Older computers mean that you will have to deal with ISA and possibly EISA slots. Obviously, purchasing the appropriate card is extremely important in making the computer network-ready. The remainder of the battle is installing the network card and the appropriate software drivers for the NIC and getting the computer to recognize both.




FIGURE 1.2
Network interface cards provide the physical connection between a computer and the network.

Match the NIC to the network architecture

If you are putting together an IBM Token Ring network, you need to purchase Token Ring network cards. Although this may be one of those things that goes without saying, acquiring the hardware (NICs and cabling) that is appropriate to the type of network you are building (say Ethernet versus Token Ring) is a complete and utter necessity.

Make sure you have the CD or disk set for the operating system running on the computer (such as Windows 98) and that you have any disks or CDs that came with the network card. Implement the following steps to get the PC up and running on the network:

Setting up the PC on the network

1. Open the case on the computer and install the NIC in an open expansion slot.
2. Close the case and attach the network medium (typically twisted-pair cabling).
3. Boot up the computer. If you purchased a plug-and-play network card and are using Windows 95/98, the card will be detected and the appropriate software drivers installed. You may be prompted to provide the drivers during this process (these drivers are on a disk or CD that came with the network card).
4. If you are using an operating system that doesn't detect new hardware devices, you will have to manually install the NIC. If the card came with installation software, use that software to install the necessary drivers.
5. Some operating systems will require that you select an IRQ and I/O port for the new NIC (this is the case with Windows NT 4—both the server and workstation OS; select an open IRQ and I/O port and then complete the installation of the card as required by your operating system).

After you physically install the card and add the appropriate driver to your software operating system, you should be up and running on the network (you might have to reboot the machine after installing any drivers for the NIC). Problems associated with NICs usually revolve around improper installation (press the card firmly into the expansion slot) and IRQ conflicts. The latter is discussed in the next section.

Dealing with IRQs and I/O Ports

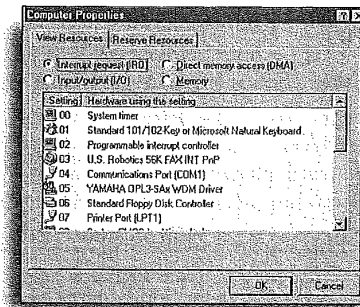
One of the most common pitfalls when installing any new device into one of the expansion slots on a PC is an IRQ conflict. IRQ

stands for *Interrupt ReQuest*. Each device in your computer, such as the mouse, keyboard, and NIC, are assigned an Interrupt Request line that the device uses to alert the microprocessor (CPU) that the device needs data processed. Each device must be assigned a unique IRQ or you have (yes, you guessed it) an IRQ conflict. Neither device will probably operate correctly if two devices are vying for the same IRQ. Knowing which IRQs are already spoken for on your system will make it easier for you to assign an IRQ to a new device such as a NIC.

Finding the available IRQs isn't that difficult, and each operating system (both PC operating systems and network operating systems) provides you with a tool to view both the used and available IRQs on a system.

For DOS clients, you can use the executable file MSD.EXE, which runs the Microsoft System Diagnostics program. This program is also available for Windows 3.11 clients.

For Windows 95 and 98, open the Control Panel (double-click My Computer and then double-click the Control Panel icon). In the Control Panel, double-click the System icon. On the System Properties dialog box, click the Computer icon, and then click Properties. A list of the IRQs on the system will appear (see Figure 1.3).



The latest operating systems make it easier to install NICs

Windows NT 2000 Server and Windows NT 2000 Professional both embrace Microsoft's Plug-and-Play scheme for plug-and-play hardware devices. This means that both of these operating systems in most cases will identify and install the appropriate drivers for a number of the network interface cards available on the market. And although you can't call what they do "plug and play," Novell NetWare 4.2 and Novell NetWare 5 both do a pretty good job of helping you set up the appropriate network card in your network server when you install either of these Novell network operating systems.

FIGURE 1.3 Operating systems like Windows 95 typically provide a tool that you can use to determine the available IRQs on a system.

In Windows NT Workstation 4.0 and Windows NT Server 4.0, you can check the available IRQs by clicking the Start menu, and then pointing at Programs. Point at Administrative Tools (Common), and then click Windows NT Diagnostics. On the Windows NT Diagnostics dialog box, click the Resources tab to view the IRQ assignments on the system.

Table 1.2 shows the standard IRQ settings for a PC. As you can see, several IRQs are reserved for particular system devices.

Table 1.2 IRQ Settings

IRQ	Use
0	System Timer
1	Keyboard
2	Cascade to secondary IRQ controller
3	COM Port 2 and 4 (serial port)
4	COM Port 1 and 3 (serial port)
5	LPT2 (printer port)
6	Floppy disk controller
7	LPT1 (printer port)
8	Real-time clock
9	Free
10	Primary SCSI adapter (or free)
11	Secondary SCSI adapter (or free)
12	PS/2 Mouse
13	Floating-point math coprocessor
14	Primary hard disk controller
15	Secondary hard disk controller (or free)

Obviously, in cases where the computer doesn't have a second COM port or an LTP2, these IRQs will be available. Each computer will vary, so use the tools mentioned earlier in this section to determine how your IRQs have been assigned.

Not only do devices need a unique IRQ to communicate with the processor, they also need a communication line that the micro-processor can use to route processed information to the device. The base I/O port for a device essentially serves as the address that the processor uses when sending and receiving data from that device. As with IRQs, each device needs a unique base I/O port. Typically, I/O ports 280h, 300h, 320h, and 360h are available for your NIC (I/O port addresses are written in hexadecimal, or base-16, format accounting for the h). The same tools for finding available IRQs on a system can also be used to determine the available base I/O ports.

Network Cabling

Copper cable is the most frequently employed network medium for local area networks. Fiber-optic cable is being increasingly employed because of its higher potential bandwidth and cable run. Fiber-optic cable is used in a number of high-speed networking implementations such as FDDI and SONET (Synchronous Optical Network, which delivers voice video and data over a high-speed fiber-optic network).

As already mentioned, copper cable is the most commonly used medium for LANs. And although copper cable comes in several different types, the most commonly used copper cable is now category 5 unshielded twisted pair (twisted-pair cable comes in 5 categories, with categories 3 to 5 being data grade cable).

Category 5 twisted pair allows Ethernet implementations of 10Mbps, 100Mbps (Fast Ethernet), and 1Gbps (Gigabit Ethernet). Unshielded twisted pair can also be used in IBM Token Ring networks. IBM has its own defining system for twisted-pair cable (both shielded and unshielded); Type 1 is the twisted-pair cable used most commonly in Token Ring installations. Twisted-pair cable typically uses an RJ-45 connector to hook to network cards, hubs, and other connectivity devices.

Although it's becoming less popular, installations of thicknet (RG-58 or RG-11 coaxial cable) can still be found in certain settings such as manufacturing companies. Thicknet is characterized by a cable backbone that is tied to servers and workstations on the network by vampire taps (the taps actually pierce the cable). The transceiver is

actually attached to the tap, and then the computer is connected to the transceiver/tap by a drop cable.

Thinnet (RG-58 coaxial cable) was the cable of choice at one time because of its relative ease of installation and lower cost. Thinnet LANs employ a bus topology where a T-connector is connected to each computer's network card. The computers are then chained together using appropriate lengths of cable. Thinnet installations require that each end of the network be terminated, and terminators are placed on the downside T-connector of the computers that reside on either end of the network.

Although copper wire is an inexpensive and easy-to-install network medium, it does have some inherent limitations. First, it can be highly susceptible to electromagnetic interference (EMI). Attenuation (the weakening of the signal over the length of the cable) also limits the length of copper cable that can be used. Copper wire can also be tapped, which may be an issue depending on the proprietary nature of the information that is being moved on the network.

Fiber-optic cable is a high-speed alternative to copper wire and is often employed as the backbone of larger corporate networks. Fiber-optic cable uses glass or plastic filaments to move data and provides greater bandwidth, longer cable runs, and is impervious to tapping. With the need for network speed seemingly on the rise, fiber installations are becoming commonplace.

Fiber-optic cable uses pulses of light as its data-transfer method. This means a light source is required and lasers and light emitting diodes (LEDs) are used. Fiber-optic cable is more expensive and more difficult to install than copper cable installations, but fiber's capability to move data faster and farther makes it an excellent alternative to copper.

Table 1.3 provides a quick summary of the various cable types. Figure 1.4 provides a look at each of the cable types listed in the table.

Choosing cable

When selecting cable for a network, a number of factors are important, including cost, cable bandwidth (the amount of information you can cram through the cable), the cable's susceptibility to EMI, attenuation (which affects the maximum cable length possible), and ease of installation. Choose the cable type that best suits your needs and budget.

Table 1.3 Network Cable Comparison

Cable Type	Bandwidth	Maximum Length	Cost
CAT 5 UTP	10Mbps to 100Mbps	100 meters	Inexpensive
Thinnet	10Mbps	185 meters	Inexpensive
Thicknet	10Mbps	500 meters	Expensive
Fiber optic	100Mbps to 2Gbps+	2 kilometers	Expensive




FIGURE 1.4
Thinnet, thicknet, twisted-pair, and fiber-optic cables are commonly used network media.

SEE ALSO

➤ For more information on the bus topology, see page 21.


Hubs, Repeaters, and MAUs

Depending on the type of cable you use and the topology of your network, you may need to use connectivity devices to connect the nodes or expand the number of nodes on your network. The type of connective device used will also depend on the type of network architecture you are using (Ethernet versus Token Ring), which is discussed later in this chapter.

Hubs are used in twisted-pair deployments and serve as the central connection point for the network. A basic hub contains no active electronics and so cannot be used to extend the network. It basically organizes your cables and relays signals to all the connective devices (see Figure 1.5).

FIGURE 1.5

A hub provides a central connection point for the network.



When is a hub no longer a hub?

Hub technology is evolving very quickly. Active hubs not only serve as the physical connection for your network nodes, but they can also serve as a repeater, allowing you to extend the size of a network. New hubs with switching capabilities are also available that can help you maximize the bandwidth on your network. Intelligent hubs are even available—they can actually help you troubleshoot connectivity problems with your network.

Physical medium equals OSI Physical layer

The actual physical medium such as the cable, hubs, and connectors operate at the Physical layer of the OSI networking model.

In cases where the network needs to be extended beyond the maximum length of the particular cable type that you are using, a repeater can be used. Repeaters take the signal that they receive and regenerate it.

In IBM Token Ring networks, the device that serves as the central connecting point is a multistation access unit, or MAU. These units actually contain active electronics and while serving as the physical connection for the devices on the network, they also provide the logical ring that is used to circulate network traffic. Multistation access units will be discussed further in the “IBM Token Ring” section of this chapter.

SEE ALSO

➤ For more information about the Physical layer, see page 43.

Understanding Network Topologies

A convenient way to discuss local area networks is by their physical layout, or *topology*. To a certain extent, the topology of a certain network will reflect the cable type used and the actual architecture of the network (such as Ethernet or IBM Token Ring). And although

the different types of topologies have been assigned particular characteristics (a bus topology, for instance, is considered to be a passive, contention-based network), the actual behavior of a particular network is better defined by the architecture used for the network. A short description of each basic network topology and a diagram of that topology type follow.

SEE ALSO

➤ For more information on network architectures, see page 25.

Bus Network

A *bus network* is characterized by a main trunk or backbone line with the networked computers attached at intervals along the trunk line (see Figure 1.6). Bus networks are considered a passive topology. Computers on the bus sit and listen. When they are ready to transmit, they make sure that no one else on the bus is transmitting, and then they send their packets of information. Passive, contention-based bus networks (contention-based because each computer must contend for transmission time) would typically employ the Ethernet network architecture.




FIGURE 1.6
A bus topology provides a passive network layout.

Bus networks typically use coaxial networking cable hooked to each computer using a T-connector. Each end of the network is terminated using a terminator specific to the cable type (if you use 50 Ohm cable, you use 50 Ohm terminators). Because the bus network is really just a collection of cable, connectors, and terminators, there is no amplification of the signal as it travels on the wire.

That bus has bounced!

When bus topology networks aren't terminated properly, the network will experience signal bounce; packets sent over the wire will actually bounce back up the line and cause collisions on the network and bring the network down. If you use the bus topology, always check the physical aspects of the network first when you are having problems. These types of networks are notorious for connector, cable, and termination problems.

Bus networks are easy to assemble and extend. They require a fairly limited amount of cabling when compared to other topologies. Bus networks are prone to cable breaks, loose connectors, and cable shorts that can be very difficult to troubleshoot. One physical problem on the network, such as a detached connector, can actually bring down the entire bus network.

SEE ALSO

➤ For more information on wide area networking, see page 25.

Star Network

In a *star topology*, the computers on the network connect to a centralized connectivity device called a *hub*. Each computer is connected with its own cable (typically twisted-pair cable) to a port on the hub (see Figure 1.7). Even though the star topology uses a hub (special hubs—multiport repeaters—can actually enhance the packet signals before passing them onto the network), this type of network still employs the passive, contention-based method of moving information on the wire that is embraced by the bus topology. Computers listen to the wire and then contend for transmission time.

Because the star topology uses a separate cable connection for each computer on the network, stars are easily expandable, with the main limiting factor being the number of ports available on the hub (although hubs can be daisy-chained together to increase the number of ports available). Expanding a star topology network is also very unobtrusive; adding a computer to the network is just a matter of running a wire between the computer and the hub. Users on the network will be pretty much unaware that the expansion is taking place.

Disadvantages of the star topology revolve around cabling needs and the hub itself. Because each computer on the network requires a separate cable, cable costs will be higher than a bus topology network (although twisted pair, the cable type used for stars, is the least expensive cable). Purchasing a hub or hubs for your network does add additional costs when you are building a network based on the star topology, but considering the benefits of this type of topology in terms of managing the physical aspects of your network, it is probably well worth it. (Hub prices have fallen to a point where even computer users with a small home network will probably want to use a hub to connect computers.)

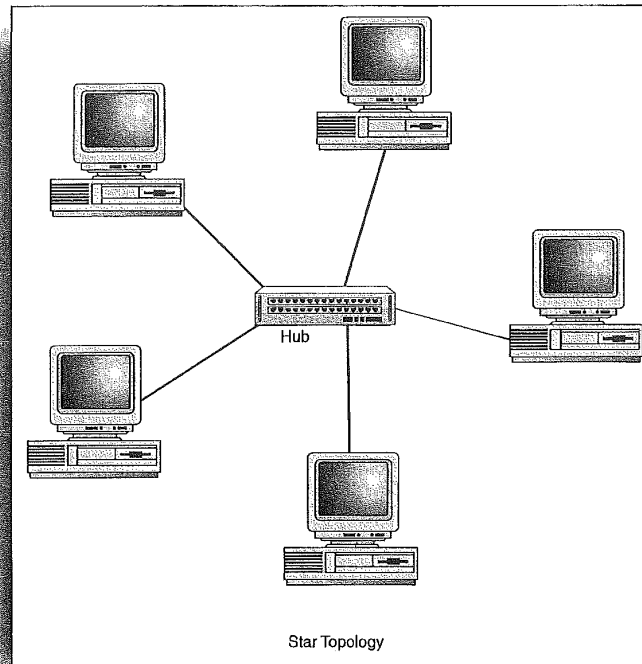



FIGURE 1.7
A star topology is easily expandable.

The most negative aspect of the star topology is related to the central hub. If the hub fails, so does the network. You will find that many network administrators who don't like crisis management keep an extra hub squirreled away just in case.

Ring Topology

A *ring topology* connects the networked computers one after the other on the wire in a physical circle (see Figure 1.8). The ring topology (an example of an architecture that uses a ring topology is Fiber Distributed Data Interface—*FDDI*) moves information on the wire in one direction and is considered an active topology. Computers on the network actually retransmit the packets they receive and then send them on to the next computer in the ring.

FIGURE 1.8
The ring topology uses a token-passing strategy to provide equal access to all the computers on the network.



Access to the network media is granted to a particular computer on the network by a token. The token circulates around the ring and when a computer wants to send data, it waits for the token to come around and then takes possession of it. The computer then sends its data onto the wire. After the computer that sent the data receives verification from the destination computer that the message was received, the sending computer creates a new token and passes it onto the next computer in the ring, beginning the token passing ritual again.

The fact that a computer must have the token to send data means that all the computers on the network have equal access to the network media. Token passing rings provide a more timely transmission of data (because of the level playing field provided by the token passing strategy) when compared to contention-based networks like the

bus or star. Token Rings actually degrade more gracefully (in terms of performance) during times of high traffic when compared to passive topologies, which can go down quickly in very high traffic situations due to increased packet collisions.

True ring topologies can be difficult to troubleshoot, and the failure of one computer on the ring can disrupt the data flow because data circulates around the ring in one direction. Adding or removing computers from this type of topology also can disrupt the operation of the network.

SEE ALSO

➤ For more information on FDDI see page 29.

Mesh Topology

The *mesh topology* uses redundant connections between computers on the network as a fault tolerance strategy. Each device on the network is connected to every other device. In short, this type of topology requires a lot of cable (see Figure 1.9). This type of topology also can weather a broken segment or two and still continue to operate as a network because of all the redundant lines.


Mesh networks, obviously, would be more difficult and expensive to install than other network topologies because of the large number of connections required. In most cases, networks that use this redundant connection strategy will actually be comprised of a hybrid mesh. In a hybrid mesh only highly important servers and mission-critical computers are configured with redundant connections. This protects the most important parts of the companywide network but doesn't require multiple lines to every computer.

Understanding Network Architectures

Network architectures provide different ways to solve a common problem—moving data quickly and efficiently on the network medium. The particular network architecture that you use, such as Ethernet, not only will define the topology for your network but also defines how the network media is accessed by the nodes on the network. There are several network architectures available, all with a different strategy for moving information on the network.

FIGURE 1.9

Each device on the network is connected to every other device on the network.

**Hybrid topologies**

As already mentioned, topologies are a convenient way to categorize the physical layout of a particular network and the strategy that it uses to move data on the wire. A number of hybrid topologies that combine the topologies discussed can exist. For example, you may chain a number of hubs together in a line, which would create a star bus topology. Or a ring network may use a connective device much like a hub that contains a logical ring (an example of a device that contains a logical ring is a Multistation Access Unit used as the central hub in an IBM Token Ring network). Computers are then connected in a star topology to this central device. This gives you a star ring configuration.

Ethernet

Ethernet is the most commonly deployed network architecture in the world. Ethernet provides access to the network using CSMA/CD (carrier sense multiple access with collision detection). This strategy of network access basically means that the nodes on the network listen (sense) to the network and wait until the line is clear. The computer then sends its packets out onto the line. If there is more than one computer transmitting, collisions result. Sensing the collisions, the computer stops transmitting and waits until the line is free. One of the computers will then transmit, gaining control of the line and completing the transmission of packets.

Ethernet is a passive, wait-and-listen architecture. Collisions are common on the network and computers are required to contend for transmission time. Ethernet networks typically will be found in a bus or star bus configuration depending on the type of network media used. One of the common implementations (on several different media types) of Ethernet runs at 10Mbps. This 10 Megabit Ethernet run over twisted pair would be designated as 10BaseT (the 10 stands for the Megabits per second, the Base means a baseband transmission (*baseband* simply means a single bit stream, or a digital flow of information), and the T stands for twisted pair). Table 1.4 lists some of the Ethernet implementations available.

Table 1.4 Ethernet Implementations

Ethernet Designation	Cable Type	Maximum Cable Length	Connector Types
10BaseT	CAT 5 UTP	100 meters	Hub
10Base2	Thinnet	185 meters	T connectors, barrel connectors, terminators
10Base5	Thicknet	500 meters	Vampire taps, transceiver drop cables, terminators
10BaseFL	Fiber optic	2 kilometers	Repeaters, terminators

When packets of information are prepared for transmission over the wire, their final form is called a frame. Ethernet actually embraces more than one frame type, which can cause problems on a network if you don't have all the nodes configured to use the same frame type. The various Ethernet frame types are as follows:

- Ethernet 802.3—Although this frame has the appropriate IEEE number, it is actually not completely in compliance with the specifications for Ethernet. This frame type is used by Novell NetWare 2.2 and 3.1 networks.
- Ethernet 802.2—This is the frame type that is in full compliance with the IEEE specifications. It is used by later versions of Novell NetWare, including NetWare 3.12, 4.x, and 5.x.
- Ethernet SNAP—This Ethernet frame type is used in AppleTalk networks.

The IEEE 802.3 specification

The specifications for running the Ethernet architecture have been defined by the Institute of Electrical and Electronic Engineers. Its designation is IEEE 802.3. Ethernet runs at the media access control sub-layer of the OSI model's Data-link layer. The OSI model and the various MAC specifications are discussed in Chapter 2, "The OSI Model and Network Protocols."

- Ethernet II—Networks running multiple protocols such as the Internet generate Ethernet II frames.

Although the 10 Megabit installations of Ethernet have been common, they are rapidly being replaced by Fast Ethernet (100 Mbps) and Gigabit Ethernet (1000Mbps or 1Gbps). Both of these versions of Ethernet require CAT 5 cabling and special network cards and hubs (Gigabit Ethernet in many cases uses CAT 6 twisted pair).

The main advantage of Ethernet is that it is one of the cheaper network architectures to implement. NICs, cabling, and hubs are fairly inexpensive when compared to the hardware required for other architectures such as Token Ring. A major disadvantage of Ethernet relates to the number of collisions on the network. The more collisions, the slower the network will run, and excessive collisions can even bring down the network.

SEE ALSO

- *Segmenting a network with a bridge or dividing a network into subnets with a router are two strategies for overcoming traffic problems on Ethernet networks. For more information, see page 67.*

IBM Token Ring

IBM Token Ring is characterized as a fast and reliable network that uses token passing as its media access strategy. Token Ring networks are wired in a star configuration with a *Multistation Access Unit* (MAU) providing the central connection for the nodes. The actual ring on which the token is circulated (the token moves in one direction as characterized by the ring topology) is a logical ring inside the MAU.

The token is passed around the ring until a computer wanting to send information out onto the network takes possession of the token. A computer that passes the token to the next computer on the logical ring would be called the nearest active upstream neighbor (*NAUN*). The computer being passed the token is the nearest active downstream neighbor (*NADN*).

After a computer takes possession of the token and transmits data, it then passes a new token to its *NADN* and the token makes its way around the ring until a node on the network takes possession to transmit.

The IEEE 802.5 specification

The specifications for running IBM Token Ring architecture have been defined by the Institute of Electrical and Electronic Engineers. Its designation is IEEE 802.5. Token Ring runs at the media access control sublayer of the OSI model's Data-link layer. The OSI model and the various MAC specifications will be discussed in Chapter 2, "The OSI Model and Network Protocols."

Token Ring is characterized by no collisions and equal access to the network media by all the nodes on the network. It is slower than some implementations of Ethernet (Token Ring can run at 4 and 16Mbps) but the network degrades more gracefully during times of high traffic. (A gigabit implementation of Token Ring will soon be a reality.)

Token Ring also provides some fault tolerance to the network with its error detection strategy, *beaconing*. When the computers on the network are first brought online, the first computer powered on is designated as the Active Monitor. The Active Monitor sends out a data packet every seven seconds that travels around the ring to help determine if any of the nodes on the network are done. For example, if a particular computer doesn't receive the packet from its NAUN, it creates a packet containing its address and the NAUN's address and sends the packet onto the network. This packet provides information that the Token Ring can actually use to automatically reconfigure the ring and maintain network traffic.

FDDI

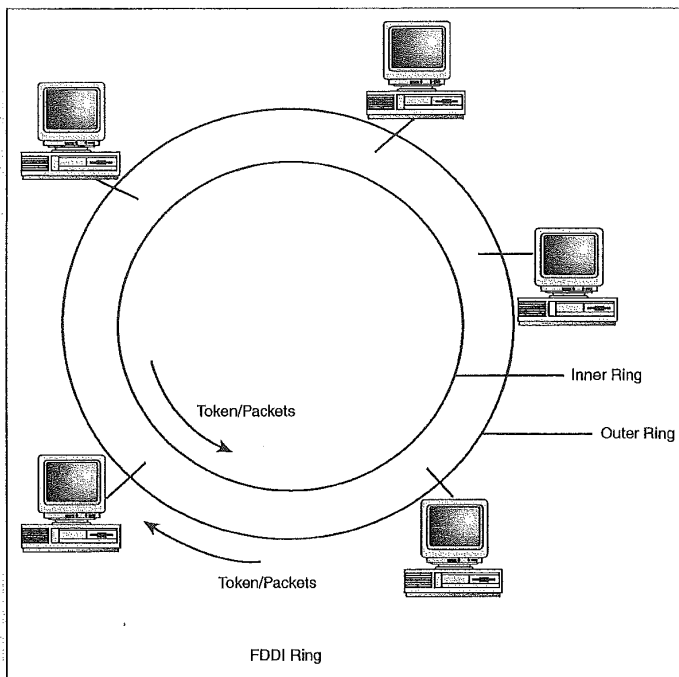
The Fiber Distributed Data Interface (*FDDI*) is an architecture that provides high-speed network backbones that can be used to connect a number of different network types. FDDI uses fiber-optic cable and is wired in a ring topology. FDDI uses token passing as its media access method and can operate at high speeds (most implementations are 100Mbps but faster data transfer rates are possible).

Because FDDI uses a token-passing media access strategy, it is reliable and provides equal access to all the nodes on the network. With FDDI you can set priority levels, however, servers on the network could be allowed to send more data frames onto the network than client computers.

Because FDDI uses a true ring topology, breaks in the cable system can be a problem. To build fault tolerance into the FDDI network, a secondary ring is used. When a computer cannot communicate with its downstream neighbor, it sends its data to the second ring (which circulates the data in the opposite direction from the one the primary ring uses).

Obviously, a special NIC is required to implement FDDI. Dual attachment stations (computers connected to both rings on the network) will use a special card that connects to both ring backbones. In place of hubs, concentrators are used on the FDDI network for the connection of LAN nodes. Because these computers don't sit directly on the FDDI ring, they only require a single attachment NIC for connection to the concentrator.

FIGURE 1.10
FDDI uses two true rings, which circulate data in opposite directions.



AppleTalk

AppleTalk is the networking architecture used by Apple Macintosh computers. The networking hardware required is already built into each Macintosh (although if you want to connect Macs to an Ethernet network, you need a Mac Ethernet NIC). The cabling system used to connect Macintosh computers is called *LocalTalk* and uses shielded twisted-pair cables with a special Macintosh adapter.

AppleTalk uses a special dynamic addressing system to determine the address of the nodes on the network. When a Macintosh is powered up on the network, the computer generates a random address and broadcasts it out onto the network. This random address becomes its network address (if another Macintosh isn't already using that address; if so, the newly powered on Mac will continue to generate random addresses until it finds one that is unused).

AppleTalk is similar to Ethernet in that it is a passive network architecture. AppleTalk uses Carrier Sense multiple access with collision detection—*CSMA/CA*. Basically the computers sit on the network and listen to determine whether the wire is clear. After making sure the network is clear, the computer will send a packet onto the network letting all the other computers know that it intends to transmit data. The computer then sends out its data.

The fact that a computer that intends to send data out onto the network notifies the other network nodes as to its intentions greatly reduces the number of collisions on a *CSMA/CA* network (especially when compared to Ethernet).

These announcement packets, however, do have a tendency to slow down the network and Macintosh networks only have a transmission speed of 230.4 Kbps. The fact that the hardware and software needed to network a group of Macintosh computers comes with each Macintosh (other than the LocalTalk cable) makes it an easy and inexpensive way to network several workstations to share a printer or files.

chapter

2

The OSI Model and Network Protocols

OSI—The Theoretical Networking
Protocol Stack

The OSI Layers

The Data Link Sublayers

Real-World Network Protocols

-
-
-
-

OSI—The Theoretical Networking Protocol Stack

Conceptual models are something that you run into no matter what discipline you tackle. Art embraces color and design theories; physics embraces nearly every theoretical model that Einstein scrawled on a napkin. Computer networking is no different and it also uses a conceptual model or framework that allows us to discuss a complex chain of events—data movement on a network.

In the late 1970s the International Standards Organization (ISO) began to develop a conceptual model for networking called the *Open Systems Interconnection Reference Model*. Networking folk more commonly refer to it as the OSI model (and I'm sure a number of them have forgotten what the OSI stands for). In 1984, the model became the international standard for network communications, providing a conceptual framework that helps explain how data gets from one place to another on a network.

The OSI model describes network communication as a series of seven layers that operate in a stack; each layer is responsible for a different part of the overall process of moving data. This framework of a layered stack, while conceptual, can then be used to discuss and understand actual protocol stacks that we see used for networking. For example, TCP/IP and AppleTalk are two real-world network protocol stacks; protocols that actually serve as layers in a protocol suite like TCP/IP can then be discussed in terms of how they relate to and serve at various levels of the OSI model's stack.

SEE ALSO

- To learn more about several of the commonly used network protocol suites, see page 44.

The OSI model provides the model for a number of important events that take place during network communication. It provides basic rules of thumb for a number of different networking processes:

- How data is translated into a format appropriate for your network architecture. When you send an email or a file to another computer, you are working with a certain application such as an email client or an FTP client. The data you transmit using this application must be placed in a more generic format if it is going to move out onto the network and to the intended recipient.

ISO seems to ring a bell

The International Standards Organization (ISO) is involved in developing sets of rules and models for everything from technical standards for networking to how companies do business in the new global market. You've probably seen banners on businesses announcing that they are ISO 9002 certified. This means that they are in compliance with the set of rules and protocols that have been developed by the ISO for doing business in the world marketplace. Another common ISO certification—ISO 9660—defines file systems for CD-ROMs.

- How PCs or other devices on the network establish communications. When you send data from your PC, there must be some mechanism that supplies a communication channel between sender and receiver. It's not unlike picking up a telephone and making a call.
- How data is sent between devices and how sequencing and error checking is handled. After a communications session has been established between computers, there must be a set of rules that controls how the data passes between them.
- How logical addressing of packets is converted to the actual physical addressing provided by the network. Computer networks use logical addressing schemes such as IP addresses. There must be a conversion of these logical addresses to the actual hardware addresses found on the NICs in the computers.

The OSI model provides the mechanisms and rules that make the handling of the issues discussed in the bulleted list possible. Understanding the various layers of the OSI model not only provides insight into actual network protocol suites, but it also provides you with a conceptual framework that can be used to better understand complex networking devices like switches, bridges and routers. (Much of this book is devoted to a discussion of routers and routing.)

The OSI Layers

The layers of the OSI model explain the process of moving data on a network. As a computer user, the only two layers of the model that you actually interface with are the first layer—the Physical layer—and the last layer—the Applications layer.

- The *Physical layer* constitutes the physical aspects of the network (the network cabling, hubs, and so on). You've probably interfaced with the physical layer at least once, when you tripped over a poorly situated cable.
- The *Application layer* provides the interface that you use on your computer to send email or place a file on the network.

Obviously, this would be a very short chapter if we only discussed these two layers, but you will find each and every layer of the OSI model plays an important part in the networking of information.

So, what's a protocol stack?

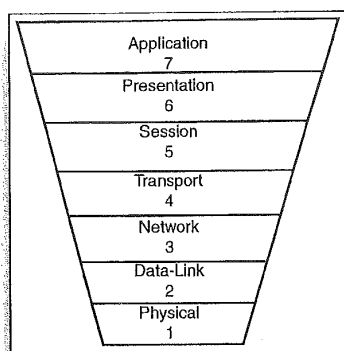
Protocol stacks or suites (or layers) are a group of small protocols that work together to accomplish the movement of data from one node on a network to another. Protocol stacks are not unlike relay-race runners, although packets of data rather than a baton are handed off to each subsequent protocol until the packets of data are in a form (a single bit stream) that can be placed on the network medium.

The ISO/OSI protocol stack exists!

While network protocol stacks like NetWare's IPX/SPX and TCP/IP are something with which most network administrators are quite familiar, there is actually a real protocol suite based on the OSI model; it's called the OSI protocol stack. Unfortunately, it is not embraced by any of the network operating systems (such as Novell NetWare or Windows NT) with which you will actually work.

Figure 2.1 provides a list of the OSI model layers from the top of the stack to the bottom. An upside-down pyramid is also an apt representation of the model because data is taken in a fairly complex form and eventually converted to a simple bit stream that can be placed on the network wire. You will notice that the layers are numbered, however, from top to bottom. For instance, in a discussion of the Network layer, you may hear the layer described as Layer 3. Whether you use the name or number is unimportant; you just need to make sure that you understand the role of each layer in the overall process of data communications.

FIGURE 2.1
The OSI model provides a conceptual basis for how data moves from a sending computer to a receiving computer.



A good way to remember the network layers from bottom to top is the following mnemonic: **Please Do Not Throw Sausage Pizza Away**. And (unfortunately, you may be thinking), you really do need to remember the OSI model; it is important to any discussion of networking technology from the very simple to the very complex. Every book or article you pick up on networking will make some reference to the model.

Before we discuss each of the layers in the stack, it makes sense to get a general idea of what takes place when data moves through the OSI model. Let's say that a user decides to send an email message to another user on a network. The user sending the email will take advantage of an email client or program (such as Outlook or Eudora) that serves as the interface tool where the message is composed and then sent. This user activity takes place at the Application layer.

After the data leaves the Application layer (the layer will affix an Application layer header to the data packet) it moves down through the other layers of the OSI stack. Each layer in turn does its part by providing specific services related to the communication link that must be established, or by formatting the data a particular way.

No matter what the function of a particular layer is, it adds header information (the headers are represented as small boxes on Figure 2.2) to the data. (The Physical layer is hardware—a cable, for instance—so it doesn't add a header to the data.)

The data eventually reaches the Physical layer (the actual network medium such as twisted pair cable and the hubs connecting the computer) of the email sender's computer and moves out onto the network media and to its final destination—the intended recipient of the email.

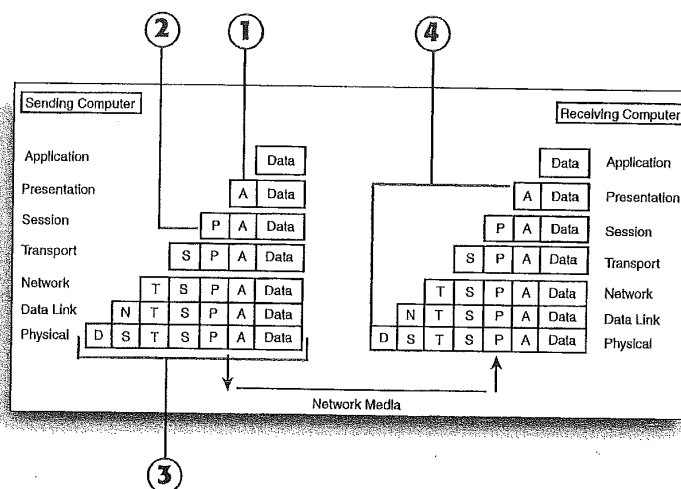


FIGURE 2.2

Data moves down through the OSI stack of the sending computer and moves up through the OSI stack on the receiving computer.

- ① Application layer header
- ② Presentation layer header
- ③ Packet with full complement of OSI layer headers
- ④ Headers are removed as the data moves up the OSI stack

The data is received at the Physical layer of the recipient's computer and moves back up through the OSI stack. As the data moves through each layer, the appropriate header is stripped from the data. When the data finally reaches the Application layer, the recipient can use his or her email client to read the received message.

The following discussion of the OSI layers will discuss the layers in the stack from top to bottom (Application layer to Physical layer).

Communications take place between peer layers

While data moves down through the protocol stack on the sender's computer (such as an email message) and eventually out onto the wire and then up the protocol stack on the receiving computer, communications do take place between complementary layers on each computer. For example, there is virtual communication between two computers sending and receiving data at the Session layer. Which makes sense because this is the layer that controls the communication between the two computers over the network media (which could be twisted-pair wire, fiber optic wire, or other connective media).

The Application Layer

The Application layer provides the interface and services that support user applications. It is also responsible for general access to the network.

This layer provides the tools that the user actually sees. It also provides network services related to these user applications such as message handling, file transfer, and database queries. Each of these services are supplied by the Application layer to the various applications available to the user. Examples of information exchange services handled by the Application layer would include the World Wide Web, email services (such as the Simple Mail Transfer Protocol—more commonly referred to as SMTP—found in TCP/IP), and special client/server database applications.

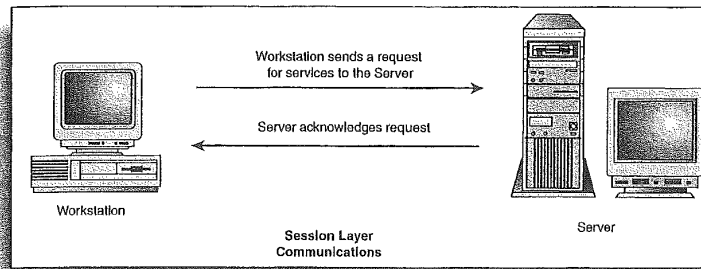
The Presentation Layer

The Presentation layer can be considered the translator of the OSI model. This layer takes the packets (packet creation for the movement of the data to the network actually begins in the Application layer) from the Application layer and converts it into a generic format that can be read by all computers. For instance, data represented by ASCII characters will be translated to an even more basic, generic format.

The Presentation layer is also responsible for data encryption (if required by the application used in the Application layer) and data compression that will reduce the size of the data. The packet created by the Presentation layer is pretty much the final form that the data will take as it travels down through the rest of the OSI stack (although there will be some additions to the packets by subsequent layers and data may be broken into smaller packet sizes).

The Session Layer

The Session layer is responsible for setting up the communication link or *session* between the sending and receiving computers. This layer also manages the session that is set up between these nodes (see Figure 2.3).

**FIGURE 2.3**

The Session layer provides the communication link between the two communicating computers.

After the session is set up between the participating nodes, the Session layer is also responsible for placing checkpoints in the data stream. This provides some fault tolerance to the communication session. If a session fails and communication is lost between the nodes, once the session is reestablished only the data after the most recently received checkpoint will need to be resent. This negates the need to tie up the network by resending all the packets involved in the session.

Actual protocols that operate at the Session layer can provide two different types of approaches to getting the data from sender to receiver: connection-oriented communication and connectionless communication.

Connection-oriented protocols that operate at the Session layer provide a session environment where communicating computers agree upon parameters related to the creation of checkpoints in the data, maintain a dialogue during data transfer, and then simultaneously end the transfer session.

Connection-oriented protocols operate much like a telephone call: You establish a session with the person you are calling. A direct connection is maintained between you and the party on the other end of the line. And when the discussion concludes both parties typically agree to end the session.

Connectionless protocols operate more like the regular mail system. They provide appropriate addressing for the packets that must be sent and then the packets are sent off much like a letter dropped in the mailbox. It is assumed that the addressing on the letter will get it to its final destination, but no acknowledgment is required from the computer that is the intended destination.

Users must run the same protocol stack to communicate

In the previous example of an email message being sent and received, it was assumed that both the sender and receiver of the data involved were running the same protocol stack (the theoretical OSI stack) on their client computers. Very different computers running very different operating systems can still communicate if they embrace a common network protocol stack. This is why a UNIX machine, an Apple Macintosh, or a PC running Windows all use TCP/IP to communicate on the Internet. A case where two computers could not communicate would be where a computer running TCP/IP is trying to communicate with a computer that is only running IPX/SPX. Both of these real-world protocols use different rules and data formats, making communication impossible.

Application layer services make user applications work over the network

When a user working in a particular application (Excel, for example) decides to save a worksheet file to his or her home directory on the network file server, the Application layer of the OSI model provides the appropriate service that allows the file to be moved from the client machine to the appropriate network volume. This transaction is transparent to the user.

Each layer performs functions on outgoing and incoming data

Remember that each layer in the OSI model (or in an actual network protocol stack such as IPX/SPX or TCP/IP) have responsibilities related to outgoing and incoming information. When data is moving down the stack on a sending computer, the Presentation layer converts information from a particular application to a generic format. On the receiving computer the Presentation layer would take generic information moving up the OSI stack and convert it into a format usable by the appropriate Application layer program on the receiving computer.

The Transport Layer

The Transport layer is responsible for the flow control of data between the communicating nodes; data must not only be delivered error-free but also in the proper sequence. The Transport layer is also responsible for sizing the packets so that they are in a size required by the lower layers of the protocol stack. This packet size is dictated by the network architecture.

SEE ALSO

» For more about network architectures such as Ethernet and Token Ring, see page 25.

Communication also takes place between peer computers (the sender and receiver); acknowledgements are received from the destination node when an agreed upon number of data packets have been sent by the sending node. For example, the sending node may send three bursts of packets to the receiving node and then receive an acknowledgement from the receiver. The sender can then send another three bursts of data.

This communication at the Transport layer is also useful in cases where the sending computer may flood the receiving computer with data. The receiving node will take as much data as it can hold and then send a “not ready” signal if additional data is sent. After the receiving computer has processed the data and is able to receive additional packets, it will supply the sending computer with a “go-ahead” message.

The Network Layer

The Network layer addresses packets for delivery and is also responsible for their delivery. Route determination takes place at this layer, as does the actual switching of packets onto that route. Layer 3 is where logical addresses (such as the IP address of a network computer) are translated to physical addresses (the hardware address of the NIC—Network Interface Card—on that particular computer).

Routers operate at the Network layer and use Layer 3 routing protocols to determine the path for data packets.

How routes are determined and how routers convert logical addresses to physical addresses are subjects that we will look at in much more detail throughout this book.

SEE ALSO

- Our discussion of the Network layer will be greatly expanded in later chapters. To begin an exploration of how routers operate at the Network layer see page 77.

The Data-Link Layer

When the data packets reach the Data-Link layer, they are placed in data frames defined by the network architecture embraced by your network (such as Ethernet, Token Ring, and so on). The Data-Link layer is responsible for data movement across the actual physical link to the receiving node and so uniquely identifies each computer on the network based on its hardware address that is encoded into the NIC (Network Interface Card). Figure 2.4 shows the hardware address for the network interface card used in a networked computer running Windows 98.




FIGURE 2.4

Each node on the network will have a unique physical address.

Real-world protocols use a combination of connection-oriented and connectionless communication

You will find that in network protocol stacks—such as TCP/IP and IPX/SPX—both connection-oriented and connectionless communication strategies are used to move data on the network. Typically, more than one protocol will operate at the Session layer to handle these different communication strategies.

Header information is added to each frame containing the sending address and the destination address. The Data Link layer is also responsible for making sure that the frames sent over the physical link are received error-free. So, protocols operating at this layer will add a *Cyclical Redundancy check* (CRC) as a trailer on each frame. The CRC is basically a mathematical calculation that takes place on the sending computer and then on the receiving computer. If the two CRCs match up, the frame was received in total and its integrity was maintained during transfer.

Again, as mentioned earlier, the frame type produced by the Data Link layer will depend on the network architecture that your network embraces, such as Ethernet, IBM Token Ring, or FDDI. Figure 2.5 shows an Ethernet 802.2 frame. Table 2.2 lists and describes each of the frame components. While you may not fully understand all the parts of the frame shown, note that the makeup of the frame is basically header information that describes the frame, the actual data in the frame, and then Data-link layer information (such as Destination Service Access Points and Service Access Points) that not only define the Frame type (in this case Ethernet) but also serve to help get the frame to the receiving computer. (For more about the IEEE 802 specifications, see the "Ethernet Frame Trivia" sidebar.)

FIGURE 2.5
The Ethernet frame is created at the Data Link layer of the OSI model.




Table 2.2 Ethernet Frame Segments

Segment	Purpose
Preamble	Alternating bits (1s and 0s) that announces that a frame has been sent
Destination	The destination address
Source	The source address
Length	Specifies the number of bytes of data in the frame
DSAP	Destination Service Access Point—this tells the receiving network card where to place the frame in buffer memory
SSAP	Provides the Service Access Point information for the frame (Service Access points are discussed in the "Data-Link section later in this chapter)
CTRL	A Logical Link control field (Logical Link control is discussed in the "Data-Link Sublayers" section later in this chapter).
Data	This part of the frame holds the actual data being sent
FCS	Frame Check Sequence field contains the CRC value for the frame

The Data Link layer also controls how computers access the physical network connections. This aspect of Layer 2 will be discussed more fully in the “Data Link Sublayers” section that follows this discussion of the OSI layers.

The Physical Layer

At the Physical layer the frames passed down from the Data Link layer are converted into a single bit stream that can then be sent out onto the network media. The Physical layer also defines the actual physical aspects of how the cabling is hooked to the computer’s NIC. On a computer that is receiving data, the Physical layer receives the bit stream (information consisting of 1s and 0s).

SEE ALSO

» To learn more about the commonly used network media and cable types, see page 17.

The Data-Link Sublayers

Before we end our discussion of the OSI networking model, we need to back track a little and discuss additional specifications that were developed for the Data Link layer of the OSI model by the IEEE. The IEEE 802 specifications divided the Data Link layer into two sublayers: Logical Link Control (LLC) and Media Access Control (MAC).

The Logical Link Control sublayer establishes and maintains the link between the sending and receiving computer as data moves across the network’s physical media. The LLC sublayer also provides Service Access Points (SAPs), which are reference points that other computers sending information can refer to and use to communicate with the upper layers of the OSI stack on a particular receiving node. The IEEE specification that defines the LLC layer is 802.2 (see IEEE specifications sidebar for more information on the categories).

Finding MAC addresses on Windows computers

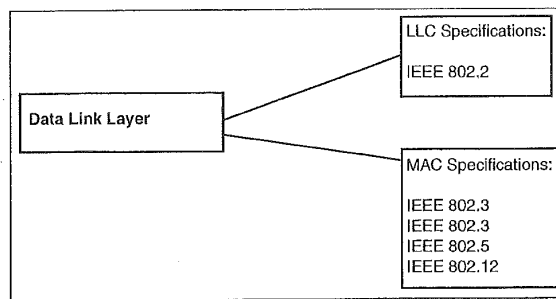
To find the address of a network card running on a Windows 95/98 computer, click the **Start** menu, and then click **Run**. In the Run dialog box, type `winnicfg`, and then click **OK**. The IP Configuration dialog box will appear for the computer and provide the address for the Network card. On a Windows NT computer, right-click on the Network Neighborhood icon and then select the **Adapters** tab on the Network dialog box. Select your network adapter and then click the **Properties** button. The MAC address of the NIC should be provided.

The Media Access Control sublayer determines how computers communicate on the network and how and when a computer can actually access the network media and send data. The 802 specifications actually break the MAC sublayer down into a list of categories (ways of accessing the network media) that directly relate to specific network architectures such as Ethernet and Token Ring (see Figure 2.6).

SEE ALSO

- » For more information on some of the common network architectures like Ethernet and Token Ring, see page 25.

FIGURE 2.6
The Data Link Layer consists of two sublayers: the LLC and the MAC.



Real-World Network Protocols

Now that we've taken a look at the theoretical model for how data moves from one computer to another on a network, as seen in the different layers of the OSI model, we can take a look at some of the most commonly used network protocol stacks and map their different layers to the OSI model. This will provide you with a good understanding of how these real-world protocol stacks operate and provide data transport on the network.

You will also see which protocols in a particular protocol stack are involved at the Network layer of the OSI model. These protocols will become important as we discuss the routing of packets on an Internetwork (something that we will do for much of the book).

NetBEUI

NetBEUI (NetBIOS Extended User Interface) is a simple and fast network protocol that was designed to be used with Microsoft's and IBM's NetBIOS (Network Basic Input Output System) protocol in small networks. NetBEUI operates at the Transport and Network layers of the OSI model.

Because NetBEUI provides only the services needed at the Transport and Network layers of the OSI stack, it needs NetBIOS, which operates at the Session layer of the OSI stack, and is responsible for setting up the communication session between two computers on the network. Two other networking components found in Microsoft networks are the Redirector and the Server Message Block. The Redirector operates at the Application layer and makes a client computer perceive all the network resources as if they were local. Server Message Block (SMB) provides peer-to-peer communication between the Redirectors on client and network server machines. The Server Message Block operates at the Presentation layer of the OSI model.

While an excellent transport protocol with very low overhead, NetBEUI is not a routable protocol, so it cannot be used on Internetworks where routing takes place. This means that while you should remember NetBEUI as a network protocol possibility for small, simple networks, it is not an option for larger networks that make use of routers (and so this is the last time you will hear about NetBEUI in this book).

TCP/IP

Often referred to as the "protocol of low bid" (see the TCP/IP Trivia sidebar for more information on TCP/IP's interesting genesis), TCP/IP has become the de-facto standard for enterprise networking. TCP/IP networks are highly scalable, so TCP/IP can be used for small or large networks.

A word about hardware addresses

NIC hardware addresses are also called *MAC Addresses*. MAC stands for Media Access Control and it is one of the sublayers of the Data-Link layer (the MAC sublayer will be discussed in the "Data-Link Sublayers" section later in this chapter). Hardware addresses are burned onto ROM chips on network interface cards, giving each of them a unique address. The addressing scheme was developed by the Institute for Electrical and Electronic Engineers (IEEE). The actual address takes the form of a 48-bit address that is written in hexadecimal format. An example of a MAC address is 00-00-B3-B3-B3-3F.

Ethernet frame trivia

The Ethernet frame used by early versions of Novell NetWare (NetWare 2.x and 3.x) was created before the IEEE specifications were completed. This means that the Ethernet 802.3 frame type is actually not to specifications as outlined by the IEEE. New versions of NetWare and other Ethernet network operating systems now use the 802.2 Ethernet frame, which is completely compliant with the IEEE specifications (the IEEE specifications are listed later in this chapter).

TCP/IP is a routable protocol stack that can be run on a number of different software platforms (Windows, UNIX, and so on) and it is embraced by most network operating systems as the default network protocol. TCP/IP contains a number of “member” protocols that make up the actual TCP/IP stack. And because the TCP/IP protocol stack was developed before the completion of the OSI reference model, these protocols do not map perfectly to the various layers of the model. Figure 2.7 shows the TCP/IP stack mapped to the OSI layers (the figure provides a general overview of TCP/IP and is not an exhaustive list of all the protocols in the stack). Table 2.3 describes the protocols listed in the figure. More information will be provided on all the protocols in the TCP/IP stack in Chapter 10, “TCP/IP Primer.”

FIGURE 2.7
TCP/IP is a large protocol stack using a number of member protocols at various layers of the OSI model.

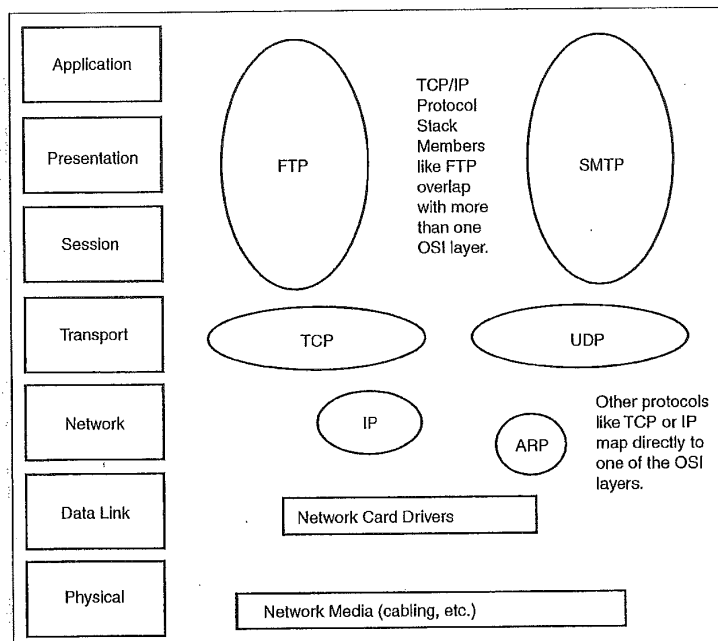


Table 2.3 TCP/IP Protocol Stack Members

Protocol	Role
FTP	File Transfer Protocol provides an interface and services for file transfer on the network.
SMTP	The Simple Mail Transport Protocol provides email services on the Internet and IP networks.
TCP	The Transport Control Protocol is a connection-oriented transport protocol. TCP handles a connection between sending and receiving computers much like a phone conversation.
UDP	User Datagram Protocol is a connectionless transport protocol that provides transport services in conjunction with TCP.
IP	The Internet Protocol is the basis for all addressing on TCP/IP networks and it provides a connectionless oriented Network layer protocol. Works much like an addressed letter that is dropped in a mail box and then delivered to the intended destination.
ARP	Address Resolution Protocol maps IP addresses to MAC hardware addresses. ARP will be discussed in greater detail in Chapter 10.

TCP/IP not only provides a very rich set of network-related features (which means that TCP/IP requires a fair amount of overhead to run) but also provides a unique logical addressing system. Anyone connected to the Internet is familiar with the 32-bit IP address, which is commonly written as 4 octets (an *octet* being 8 bits of information). The typical IP address is written in the format 129.30.20.4, where each of the four dotted decimal values actually represent 8 bits of binary information. Much more information concerning IP addressing will be discussed in Chapter 10.

Because of TCP/IP's importance in Internetworks and the complexities related to routing TCP/IP networks, an entire chapter of this book has been provided reviewing all the aspects of TCP/IP addressing. A great deal of information will also be provided on the commands related to routing TCP/IP on a campus or enterprise network.

SEE ALSO

- The best place to start in on TCP/IP and routing is Chapter 10, "TCP/IP Primer," beginning on page 167.

The IEEE 802 specifications

The IEEE 802 specifications provide categories that define the Logical Link Layer and the different network architectures that can be embraced by the MAC layer. A complete list of the 802 categories is provided:

- 802.1 Internetworking
- 802.2 Logical Link Control
- 802.3 Ethernet(CSMA/CD) LAN
- 802.4 Token Bus LAN
- 802.5 Token Ring LAN
- 802.6 Metropolitan Area Network
- 802.7 Broadband Technical Advisory Group
- 802.8 Fiber Optic Technical Advisory Group
- 802.9 Integrated Voice and Data Networks
- 802.10 Network Security
- 802.11 Wireless Networks
- 802.12 Demand Priority LAN

IPX/SPX

IPX/SPX (Internetwork Packet Exchange/Sequenced Packet Exchange) is a network protocol stack developed by Novell for use in the Novell NetWare network operating system. IPX/SPX is a leaner stack than TCP/IP and does not require the overhead needed by TCP/IP. IPX/SPX is suitable for small and large networks and is a routable network protocol suite.

Figure 2.8 maps protocols in the IPX/SPX stack to the OSI Layers. Table 2.4 gives a brief description of each of the protocols.

FIGURE 2.8

IPX/SPX is an efficient network protocol stack used on large and small networks.

TCP/IP trivia

TCP/IP was developed by Defense Advanced Research Projects Agency (DARPA). The Department of Defense needed a protocol stack that could communicate across unlike networks. The unlike networks existed because the government uses a bidding system and suddenly found itself with different computer systems at various branches of the Defense Department: the Army, Navy, and so on. So, TCP/IP is jokingly called the protocol of low bid because it was in part developed to fix a problem that arose because of the way the government takes bids for procuring technology and other goods.

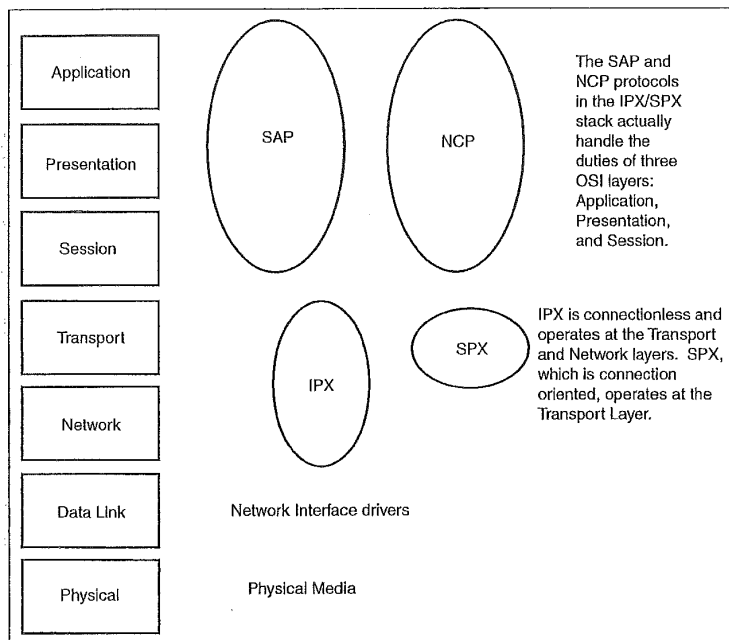


Table 2.4 IPX/SPX Protocol Stack Members

Protocol	Role
SAP	The Service Advertising Protocol is used by NetWare File Servers and Print Servers to announce the address of the server.
NCP	The NetWare Core Protocol handles network functions at the Application, Presentation, and Session layers. It handles packet creation and is responsible for providing connection services between clients and servers.
SPX	Sequenced Packet Exchange Protocol is a connection-oriented transport protocol
IPX	Internetwork Packet Exchange Protocol is a connectionless transport protocol that handles addressing and routing on the network.

Our major concern with IPX/SPX is routing this protocol suite on an Internetwork. More information on routing IPX/SPX and how the IPX/SPX stack moves data on the network is provided later in this book.

SEE ALSO

- *Routing IPX/SPX is discussed in Chapter 12, "Routing Novell IPX," which begins on page 211.*

AppleTalk

While many network administrators would not consider AppleTalk an Internetworking or enterprise network protocol, AppleTalk is routable. And with the appropriate type of NIC (Apple Macintoshes can participate on an Ethernet network if they are outfitted with EtherTalk cards or other adapters) it can support Ethernet, Token Ring, and FDDI architectures. It is not uncommon to have Macintosh computers in the Enterprise to support graphic manipulation and other multimedia duties and so it makes sense to include AppleTalk as another key routable protocol stack on the corporate network.

Earlier, in Chapter 1, we discussed AppleTalk as architecture, but it is also a network protocol stack. Figure 2.9 maps the protocols in the AppleTalk stack to the layers of the OSI model. Table 2.5 gives a brief description of each protocol.

Figure alert!

Figures 2.7 through 2.9 map real-world protocols to the OSI model. To understand these figures, think back to how the OSI model describes in seven layers how data moves from one computer to another and the transformation that it must undergo. Real-world stacks like TCP/IP perform all the tasks described in the OSI model; they just do it with fewer protocols. Rather than having seven protocols (one for each of the OSI layers) TCP/IP has certain protocols that handle the duties of more than one OSI layer. For example, FTP handles Application, Presentation and Session layer duties. The circle around FTP spans all three of the layers on the OSI model (the layers are the boxes).

FIGURE 2.9
AppleTalk is a routable protocol stack for Macintosh networks that can communicate with Ethernet, Token Ring, and FDDI networks.

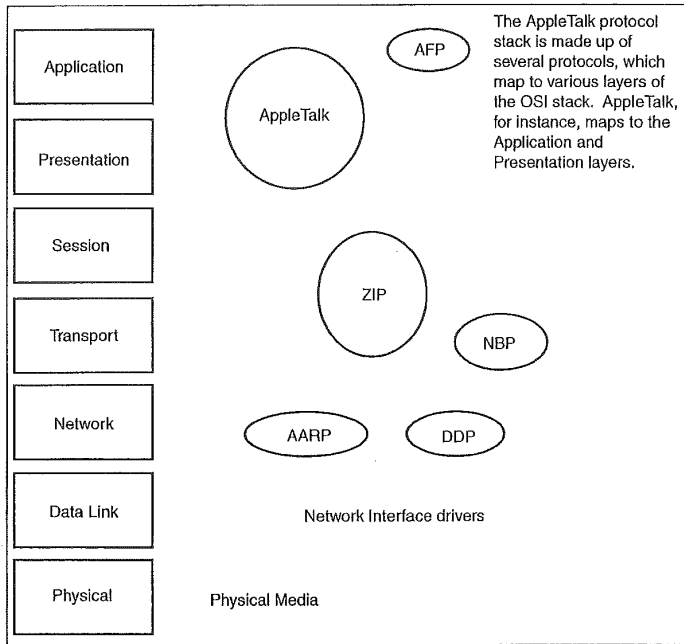
Terminology alert!

Before we go too much farther, we should sort out some terms that you will find throughout this book.

Internetwork: a network of networks. Local Area Networks connected by an Internetwork device such as a bridge or router. Internetworking is discussed in detail in Chapter 4, "Internetworking Basics."

Internet: The global network of networks. TCP/IP is the de-facto standard for this global collection of heterogeneous computers.

Intranet: A corporate network that is internal to the enterprise (not connected to the global Internet) but uses Internet protocols such as Simple Mail Transport Protocol and Hypertext Transport Protocol (the protocol used by Web Browsers) to share information among corporate users. An *extranet* is an intranet that provides corporate network access to specified users outside the company.



The AppleTalk protocol stack is made up of several protocols, which map to various layers of the OSI stack. AppleTalk, for instance, maps to the Application and Presentation layers.

Table 2.5 AppleTalk Protocol Stack Members

Protocol	Role
AppleShare	AppleShare provides services at the Application layer
AFP	The AppleTalk Filing Protocol provides and managing file sharing among nodes on the network
ATP	The AppleTalk Transaction protocol provides the Transport layer connection between computers
NBP	The Name Binding Protocol maps computer hostnames to Network layer addresses
ZIP	The Zone Information Protocol controls AppleTalk zones and maps zone names to network addresses
AARP	The AppleTalk Address Resolution Protocol maps logical Network layer addresses to Data Link hardware addresses

Protocol	Role
DDP	The Datagram Delivery Protocol provides the addressing system for the AppleTalk network and provides connectionless transport of datagrams between computers


As with IPX/SPX, our interest in the AppleTalk network protocol stack relates to routing AppleTalk. More information on how AppleTalk networks are configured and how AppleTalk is routed on a Cisco router is provided later in this book (see Chapter 13, "Routing Apple Talk.")

SEE ALSO

- *More information on how AppleTalk networks are configured and how AppleTalk is routed on a Cisco router is provided on page 227.*

Where are the routing protocols?

You might have noticed that the diagrams that map various protocol stacks to the OSI model did not include routing protocols. Obviously, each protocol stack has a default routing protocol; for example, RIP is the default routing protocol for TCP/IP and the Routing Table Maintenance Protocol is the routing protocol for the AppleTalk stack. These protocols will be discussed in greater detail when routing of these protocol stacks is discussed later in this book.



chapter

3

Wide Area Networking

Understanding Wide Area Connectivity

Getting Connected

Switched Network Overview

Circuit Switching

Packet Switching Protocols

Other WAN Protocols

-
-
-
-
-
-

Understanding Wide Area Connectivity

As the PC local area network became more and more important to businesses, corporations, and institutions, the need to expand and then connect LANs became a necessity. Expanding or connecting LANs locally (in a fairly limited geographic area) was taken care of by internetworking devices such as repeaters, bridges, switches, and routers. However, when connecting LANs over large distances, other technology must come into play.

A need for technology that provided network administrators with the ability to connect LANs over greater geographic areas became extremely important as networking the enterprise (the enterprise is the entire corporation—which in many cases can be a worldwide operation) became an imperative.

Expanding a network across great distances can be accomplished by taking advantage of several different wide area networking technologies. Networks can be connected with services provided by the public switched telephone network (PSTN) or private carrier companies. Extremely large companies can invest in their own WAN infrastructure and invest in microwave and satellite transmission equipment.

WAN technology can be used to connect networks between two cities, across the country, or around the world. As with LANs and internetworks, after the Physical layer aspects have been taken care of, various protocols are used to move the data on the WAN. On a LAN, the cable and the hubs provide the Physical layer, while on a WAN, the Physical layer can be a T1 leased line or a satellite dish.

SEE ALSO

➤ For more information on internetworking, see page 67.

WANs take advantage of wireless technologies

Although LANs typically use some sort of physical wiring (copper or fiber-optic cable), WANs can take advantage of several wireless technologies, including microwave transmissions, satellite links, infrared light, and network communications via radio signals (both single-frequency and spread-spectrum radio transmissions).

Getting Connected

While the actual physical infrastructure (the cabling and networking devices such as hubs, repeaters, and so on) of a LAN will be owned by a company, most businesses and institutions find it too costly to own the physical WAN connections that they use.


Three types of WAN connections are available: a connection over the Public Switched Telephone System via a modem, a dedicated connection such as a full-time leased line, or a switched connection that enables multiple users to take advantage of the same line.

Each of these WAN connection possibilities offers its own set of pros and cons and each embraces different hardware needs. The following sections discuss these three WAN connection alternatives.

Dial-Up Connections

The simplest and least expensive type of dial-up connection uses a modem to connect two computers over a regular analog voice-grade telephone line. The modem converts the digital information on the computer to an analog signal (modulation) and vice versa (demodulation). That's how the modem got its name. This conversion process allows computer data to be sent over the analog line. Modems are now available that have potential transmission speeds of up to 56Kbps, but line noise can limit the speed at which a connection over an analog line can run.

Routers can be outfitted with a modem connection (and then are often referred to access servers). This means that two LANs could be connected via a dial-up connection and packets routed (although this would provide a very slow connection between the networks). Figure 3.1 shows two LANs connected via a dial-up connection with routers serving as the connection point for the asynchronous connection over modems.



Keep the packets on the company-owned wire

One of the secrets of being a highly successful WAN or internetwork administrator is designing large networks and setting up the potential routing of packets so that traffic circulates on the company-owned networking infrastructure as much as possible. Running a cost-effective internetwork or WAN is the ultimate challenge when involving leased lines and connections for which you basically pay for the bandwidth.

FIGURE 3.1
LANs can be connected using dial-up connections via modems.

56K modems don't give you 56Kbps

Although 56K modems greatly increase the speed of downloads and uploads from a home or office PC (when compared with the previously available 33.6 modems), you probably have found that you never get more than 53Kbps as your throughput. This is because it takes increased power to move data over regular phone lines at higher speeds (such as those supported by the 56K modem technologies) and the FCC has placed a limit on the power available. This is because increasing the power on the phone lines increases the amount of interference (or crosstalk) between the wires present in the twisted-pair copper wiring used for the phone system. Your connection speed will also be limited by the age and amount of interference on the lines. You might find that you never get more than 49Kbps at times, and this is directly related to the "dirty" phone lines that you are forced to use as your communication medium.

Leased Lines

Dedicated leased lines provide a full-time connection between two networks through the PSTN or another service provider. Leased lines are typically digital lines. They provide much more bandwidth than analog lines and are less susceptible to the line noise and interference found on voice-grade connections.

Digital lines commonly used for data transfer are DDS (digital data service) lines and the T-carrier system, which provides a range of line types that provide different data rates.

DDS Lines

DDS lines, which are typically available from your local phone service provider, can provide bandwidth of up to 64Kbps and supply your network with a permanent, full-duplex connection (data can be sent and received at the same time). Because DDS lines are digital, they require a Channel Service Unit/Data Service Unit (CSU/DSU) as the connecting point between your LAN and the DDS line. The CSU/DSU converts the LAN data transmission into a digital signal on the DDS line.

The DSU side of the CSU/DSU is connected to your LAN and the leased digital line is connected to the CSU port of the device. Some sort of internetworking device such as a bridge or router will typically sit between your network and the CSU/DSU (on both ends of the DDS connection). Figure 3.2 shows two LANs connected by a DDS line.

T-Carrier Lines

The T-carrier system takes advantage of technology that allows several transmissions, which can include voice communication and data (divided onto several different channels), to be combined and then transmitted on a single high-speed line.

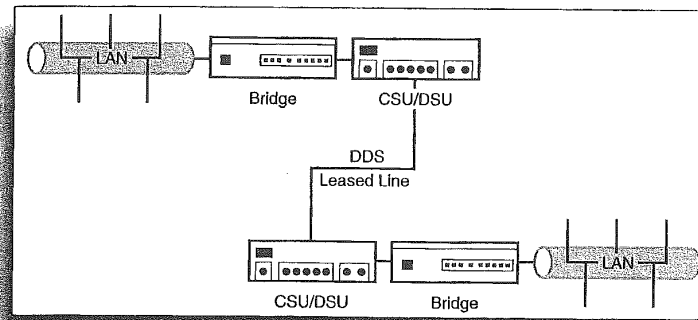


FIGURE 3.2
DDS lines provide a constant connection between local area networks.

The device that combines signals carried on these separate channels (when data must be sent over the digital line as a single data stream) and also has the capability to split a received data stream into the appropriate channels is called a multiplexor, or MUX. Figure 3.3 shows two different LANs connected by a T-carrier. Multiplexors are used at either end of the digital connection to assemble and disassemble different data channels, including data from any attached networks and the company's voice channel (used for their telephone system).

The T-1 line is the basic unit of the T-carrier system. It provides 24 channels that can be combined to provide a total transmission bandwidth of 1.544Mbps. Several other T-carrier classes exist, which can provide a larger number of channels and extremely high data rates. A greater number of channels and a higher data throughput, however, relates directly to the cost for the carrier line. Table 3.1 provides a listing of the T-carriers.

Table 3.1 The T-Carrier Systems

Carrier Line	Channels	Total Data Rate
T1	24	1.544Mbps
T2	96	6.312Mbps
T3	672	44.736Mbps
T4	4032	274.760Mbps

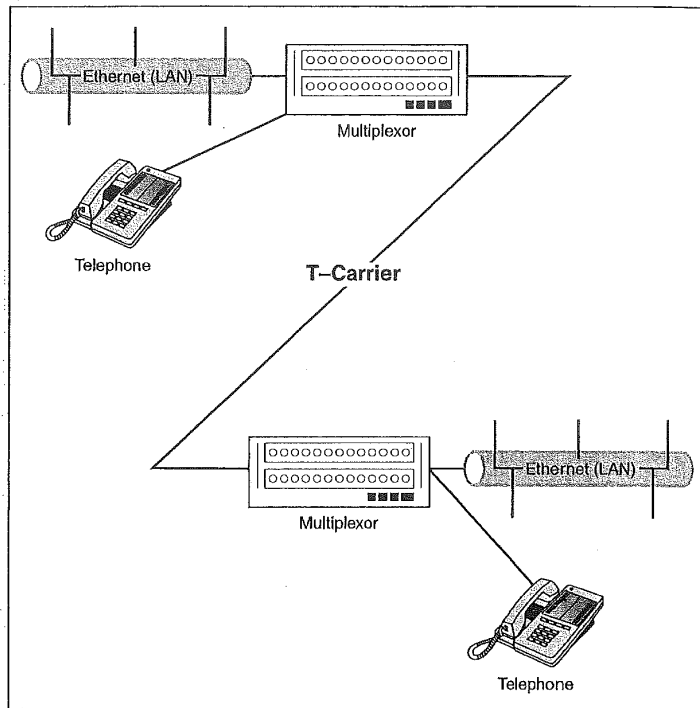
Telephone trivia

The Public Switched Telephone Network (PSTN) is also often referred to as POTS. This stands for Plain Old Telephone System (you might also hear it referred by other names, especially when your leased lines are down).

DDS lines are not ISDN lines

DDS lines are digital lease lines. They are special lines to which the phone company provides access. ISDN is a digital technology designed to use digital technology over the existing phone lines.

FIGURE 3.3
Multiplexors can combine channels for transmission or disassemble a data stream into its channels when connected to T-carrier lines.



DDS line connections are being replaced by other WAN technologies

DDS line connections are becoming a thing of the past and are fast being replaced by other WAN technologies, specifically some of the packet-switching alternatives such as Frame-Relay. A drop in the price of T-carrier lines and the use of Fractional-T connections (a portion of a T1 line is leased) are giving network administrators more bandwidth for their buck in the WAN arena.

The T-1 line is the most affordable T-carrier (and the most leased of the T-carrier classes) and can be deployed on copper wire. And as mentioned earlier, smaller companies can lease a portion of a T-1 line, selecting to use only a certain number of channels. T-2 lines aren't available to the public and are used internally by the phone company.

T-carriers offering greater bandwidth such as T-3 and T-4 lines are employed by only very large corporations or government entities (in a large part because of cost). T-3 lines and T-4 lines also require fiber optic cabling.

When attaching a company to a T-carrier line, the equipment needed is the same as for a connection to any dedicated digital line. A CSU/DSU sits between the line connection and an internetworking device such as a bridge, switch, or router, which is attached to your computer network. And as mentioned before, if channels on the carrier line are going to be split between voice and data communications, a multiplexor is needed to combine and split the signals as needed (outgoing versus incoming in relation to the T-carrier line).

Switched Network Overview

The third alternative for WAN connections is the use of switched networks. Switched networks allow multiple users to take advantage of the same line. Switched networks offer a cheaper alternative to the cost of leasing dedicated lines.

Basically, your network is connected to the wide area network via a service provider or the phone company itself. Data leaving your network through the WAN connection then enters the switched network (which is often represented on diagrams as a cloud because the path of your switched data can potentially be different each time as it makes its way to the designated destination). See Figure 3.4.

The connection between your network and the switched network, or *PDN* (public data network), will be in the form of a digital terminal device (*DTE*) such as a router. (A *DCE*—data circuit terminating equipment—such as a DSU/CSU may sit between your router and the PDN; with the DCE providing the bandwidth and timing settings for the transfer of your data.) The PDN provides the lines and switching equipment that will move your data through the switched network cloud.

Two types of switched network possibilities exist for WAN connections: circuit switching and packet switching.

Cable television holds the key

The concept of what a multiplexor does (at least on the receiving end of a single bit stream from a T-carrier connection) isn't really new to anyone who has cable television in his or her home. A single data signal comes into the house, and your cable-ready television or VCR contains a multiplexor that breaks down the data feed into the 100+ television channels that you constantly are surfing through. That's what broadband transmission is all about—multiple channels on a single feed.

The T-carriers are multiples of T-1s

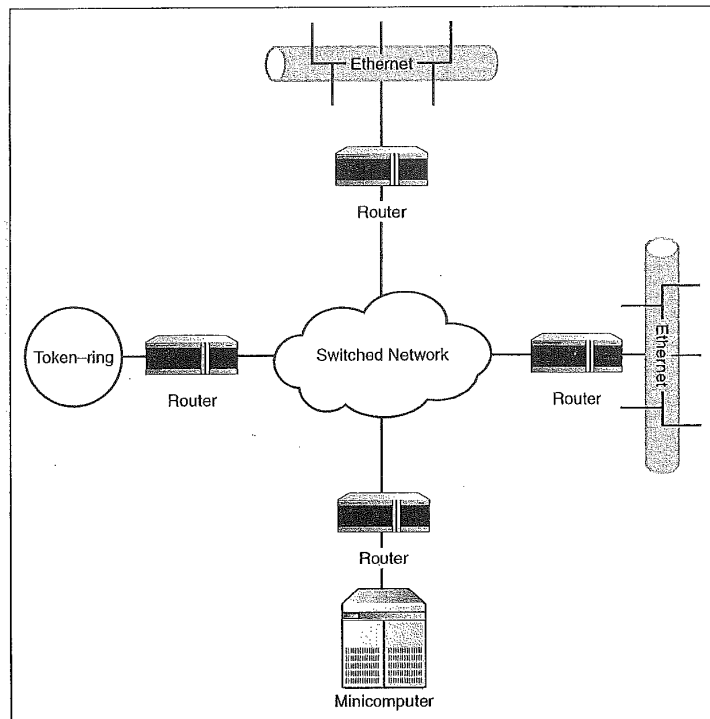
The T-1 line is the basic unit of the T-carrier system. All the other T-carriers available can actually be thought of as simply a particular multiple of T-1 lines. The T-2 line consists of 4 T-1s; T-3 consists of 28 T-1s; and T-4 consists of 168 T-1 lines.

FIGURE 3.4

Switched networks enable you to connect LANs at different sites so that they can share data over greater distances.

Public switched network versus private data networks

When you work with WAN technology, you basically have two types of networks that you can use as the carrier for your data. The public switched network—also known as the public data network and the plain old telephone system (POTS)—is one avenue for the movement of your data. You can also use private data networks as your carrier. These networks are owned by companies like GE, Sprint, and MCI and provide another avenue for WAN-switched technologies and leased lines.



Circuit Switching

Circuit switching establishes a dedicated connection between the sender and receiver on the PDN. Data moves from the source to the destination along the circuit (the lines) that has been established for the particular session. When the data transfer is complete, the connection between the sender and receiver ends and the circuit is terminated.

An example of a circuit-switching WAN technology is *ISDN* (Integrated Services Digital Network). ISDN is available from local phone providers and takes advantage of digital phone switching systems. The cost of an ISDN connection will be dictated by how often the line is used for data transfer. Your usage charge is determined by

the connection charge (and there is also often a recurring monthly charge for being connected to the service). ISDN comes in two flavors: basic rate ISDN (BRI) and primary rate ISDN (PRI).

Basic rate ISDN provides three channels: two B channels that each provide 64Kbps of bandwidth for data transfer and a D channel operating at 16Kbps that is used exclusively for setup and control information. BRI can be used for both voice and data communications by dedicating a B channel for each. Typically, however, the two B channels are combined in BRI to provide a data transfer speed of 128Kbps.

Primary rate ISDN—designed for larger businesses that require greater bandwidth—uses a T1 line and provides 23 B channels (each operating at 64Kbps). One D channel is also necessary (as with BRI) to handle setup and control of the connection).

SEE ALSO

➤ *Configuring ISDN on a router is discussed on page 268.*

Packet Switching

In WAN connections that use *packet switching*, the data is divided into packets. The small packet size used in packet-switching WANs provides fast and efficient delivery of data.

Each packet has its own control information and is switched through the network independently. This means that data packets can follow different routes through the WAN cloud and reach the destination out of sequence. However, sequencing information in the packet header can be used by the receiving device to reassemble the data in the appropriate order.

Packet switching networks can take advantage of virtual circuits when transferring data. A *virtual circuit* establishes a defined route across the WAN cloud so that all the data packets move to the destination along the same route (remember that this route is shared by packets from many other users because switched networks use shared lines). The use of virtual circuits in packet switching networks can improve the overall performance of your data transfers.

Basic rate ISDN a thing of the past?

Basic rate ISDN was designed for the small business and home users who required faster connections of very small networks or single workstations to WANs, specifically the Internet. A newer technology that provides faster connection speeds is digital subscriber line service (DSL). DSL offers voice and data communication over the digital line with speeds of up to 7Mbps. Again, handled by your local phone service provider, DSL costs can actually be cheaper than the sum of a home phone bill and the monthly cost for a 56K analog connection to an Internet service provider. This makes DSL a much cheaper alternative to BRI (DSL charges are typically flat rate, while BRI charges are based on the amount of time the connection is used).

Several packet switching technologies exist such as X.25, frame relay, and ATM. The next section of this chapter discusses these packet switching protocols.

WAN Packet Switching Protocols

Packet switching networks have been available since the late 1970s when X.25 became available. The lower cost of packet switching networks (when compared to dedicated leased lines) led to a fairly rapid evolution of packet switching protocols that now makes the movement of data over packet switching networks very fast and extremely efficient. The following sections discuss some of the popular packet switching protocols.

X.25

X.25 was designed for use over the PDNs that were operated by companies such as AT&T and General Electric. The X.25 protocol stack provides point-to-point communications between local area networks on the WAN using DTEs and DCEs (with the DCE providing the connection from a DTE, such as a router, to the actual WAN connection).

Because the purpose of any WAN is to connect geographically separated LANs, X.25 sessions consist of communications between two DTEs. For example, you might have a LAN in Chicago that is connected to a router that then provides a connection to a PDN. Likewise, you have a LAN in Minneapolis that is connected to a PDN via a router. The X.25 protocol can then handle a connection between these two DTE devices on the WAN, so that they can exchange data (see Figure 3.5).

The X.25 protocol stack consists of protocols that operate at the Network, Data Link, and Physical layers of the OSI model. These protocols are as follows:

- **Packet Layer Protocol (PLP)**—Operating at the Network layer, this protocol manages the exchange of packets between the connected LANs (such as the routers in Chicago and Minneapolis discussed earlier). PLP establishes the virtual circuit

between the DTE devices and is also responsible for segmenting and reassembling the packets as they move from the sending to the receiving device. PLP also closes the virtual circuit when data transfer is complete.

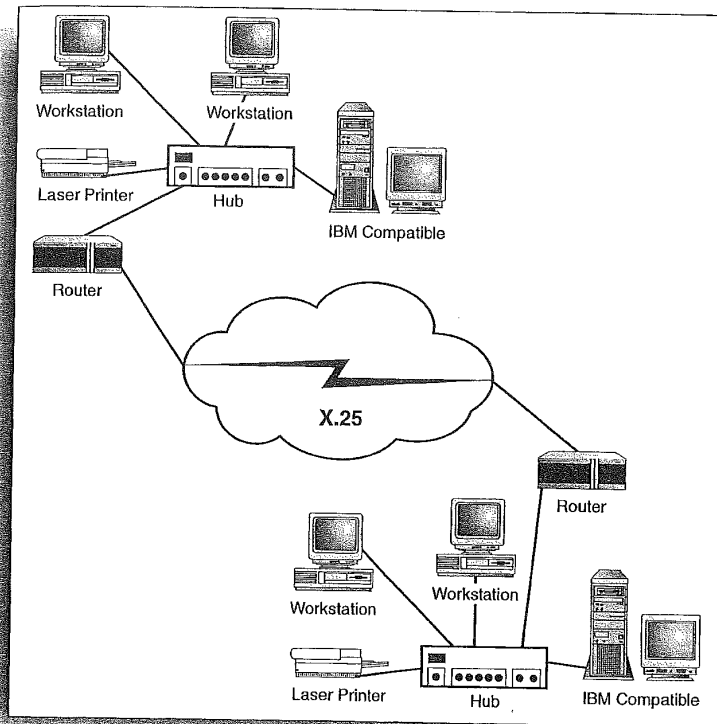


FIGURE 3.5
X.25 establishes a virtual circuit between two DTEs (routers) on the WAN.

- **Link Access Procedure/Balanced Protocol (LAP/B)**— Operates at the Data Link layer and makes sure the frames are delivered error free and in the proper sequence.
- **X.21bis**—This Physical layer protocol provides the activation and deactivation of the physical media that connects the DTE and DCE devices.

Two types of virtual circuits

Two different types of virtual circuits can be established using X.25 packet switching. Switched virtual circuits (SVC) are set up for a particular communication session and then torn down at its conclusion. A PVC (permanent virtual circuit) is established and used for recurring communication between two points, providing an active session for communication between networks.

Because X.25 was created for use on the public switched telephone network that typically consisted of noisy analog lines, X.25 is bogged down with a great deal of error-checking capabilities. Although still in use, X.25 is fast being replaced by speedier packet switching protocols such as frame relay and ATM.

SEE ALSO

- *Configuring X.25 on a router is discussed on page xxx.*

Frame Relay

Frame relay is the successor to the X.25 protocol. It is a layer 2 WAN protocol that provides high-speed connections between DTE devices (such as bridges and routers) that typically operate over fiber optic cable. DCE devices on frame relay networks consist of the carrier-owned switches. Frame relay is faster than X.25 because it has shed some of the control and error-checking functions that slowed the packet-switching capabilities of X.25.

Frame relay uses permanent virtual circuits for communication sessions between points on the WAN. These virtual circuits are identified by a DLCI (Data Link connection identifier)—a value provided by the frame relay service provider. Because several virtual circuits can exist on a frame relay interface, the DLCI for a particular virtual circuit (the one you are using to move your packets) can be used as a reference or pointing device that makes sure the packets end up at the proper destination. This is done by mapping the logical addresses (IP addresses, for example) of the sending and receiving DTEs to the DLCI of the virtual circuit that they use to communicate.

SEE ALSO

- *Configuring frame relay on a router is discussed on page 265.*

Asynchronous Transfer Mode (ATM)

Another packet switching WAN technology is asynchronous transfer mode (ATM). ATM is an advanced packet-switching protocol that uses fixed packet sizes (53 bytes) called *cells*. Using a fixed packet size (which X.25 and frame relay do not) actually increases the

throughput speed of the data because the switching and routing equipment can move the consistently sized cells faster. ATM can move data at a theoretical speed of up to 2.4Gbps. Typically ATM WAN speeds fall between 45 and 622Mbps. The 622Mbps is achieved on the fastest WAN network medium available—ONET (synchronous optical network, a fiber optic network developed by Bell Communications Research that provides voice, data, and video at high speeds).

ATM is similar to frame relay in that it assumes that the lines it uses are noise-free; therefore, it doesn't require a great deal of overhead for error checking (which, as you remember, slowed X.25 packet switching down). ATM can be used over FDDI backbones on metropolitan area networks (at speeds of 100Mbps) and T-3 leased lines (at speeds of 45Mbps).

While they aren't one of the WAN technologies that we will look at in terms of router configurations, ATM networks take advantage of fast ATM switches that quickly move data from port to port as it travels from sending station to receiving station.

Other WAN Protocols

When you work with routers, two other WAN protocols become important: High-Level Data-Link Control (HDLC) and Point-to-Point Protocol (PPP). Each is commonly configured as the protocol for router serial interfaces.

- **HDLC**—HDLC is the default WAN protocol for Cisco Router serial interfaces and is used for synchronous serial connections (digital connections such as ISDN). Cisco's version of this Data Link layer WAN protocol is proprietary and unfortunately will not communicate with other HDLC implementations.
- **PPP**—PPP is widely used as the protocol for connecting dial-up connections to TCP/IP networks such as the Internet. PPP can be used over asynchronous (dial-up) or synchronous lines. PPP supports data compression and provides authentication using either the Password Authentication Protocol (PAP) or the Challenge Handshake Authentication Protocol (CHAP).

SEE ALSO

- *Configuring HDLC on a router is discussed on page 261.*

SEE ALSO

- *Configuring PPP on a router is discussed on page 262.*

chapter

4

Internetworking Basics

What Is Internetworking? ●

Internetworking Devices ●

Building a Campus Network ●

What Is Internetworking?

In its strictest sense, internetworking is the connecting of two or more LANs where the LANs still function as separate entities. In a broader sense, internetworking is a strategy for expanding, segmenting, and connecting LANs so that *bandwidth* on and between the networks is maximized. (Bandwidth in this case would be the potential throughput of the medium that you are using, for example a 10baseT LAN runs at a bandwidth of 10Mbps.)

Internetworking embraces both LAN and WAN technologies to move information between the networks. The great thing about internetworking is that it could be used as a strategy for connecting networks that embrace the same network architecture (such as two Ethernet LANs) or as a strategy for networks that use different network architectures (such as an Ethernet network and Token Ring network). An excellent example of a real-world internetwork is the Internet.

Figure 4.1 shows an internetwork that employs some of the internetworking strategies and devices that are typically used.

Internetworking Devices

As your company's network grows, you will need to deal with issues related to extending the network, conserving bandwidth on your network, and connecting your network across greater geographical distances (using WAN technology). Several different internetworking devices exist to fill network expansion needs. These devices are as follows:

- Repeaters
- Bridges
- Switches
- Routers
- Gateways

Network growth means expansion and segmentation

As LANs grow, you will find that you must expand the number of workstations locally. In cases where the network expands beyond your local geography (from a building on one side of town to a building on the other side of town, for example) your expansion might require the use of WAN technology to connect the two networks. As localized networks grow (LANs in a large corporate office location), you will find that the more workstations and servers you add, the more burden you put on the network in terms of maintaining the bandwidth available. In this case you use segmentation as a strategy to break the larger network into segments that operate as separate units when communicating locally on the segment. This helps preserve bandwidth. This chapter discusses both of these situations in terms of internetworking strategies.

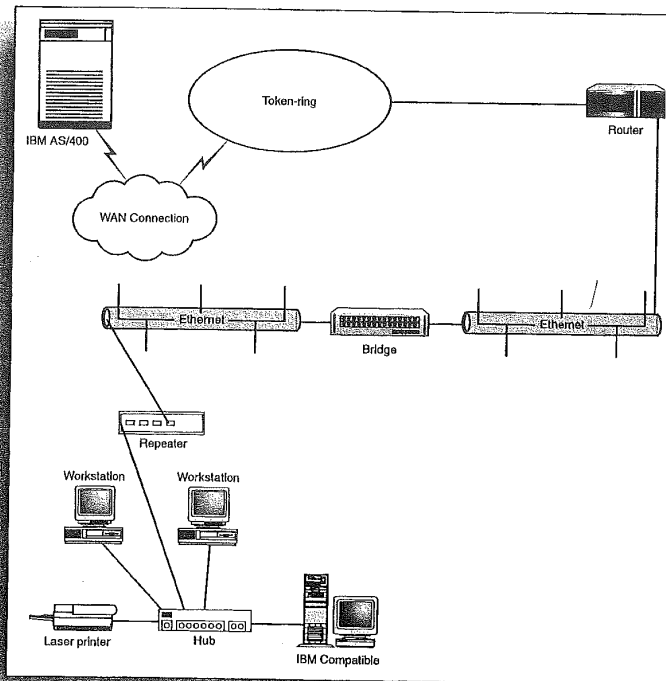


FIGURE 4.1
Internetworks embrace LAN and WAN technologies to expand and connect networks.

So, internetworking is...?

For purposes of discussion in this chapter, internetworking embraces the expansion, segmentation, and connection of LANs using LAN, internetworking, and WAN devices and protocols. This means that repeaters, bridges, switches, routers, and gateways will all be included in the list of internetworking devices (whereas in the very strictest definition of internetworking, only routers and gateways would qualify as internetworking devices).

LAN and WAN protocol reviews

Because internetworking uses both LAN and WAN protocols, you might want to review Chapters 1, "LAN Review" and 3, "Wide Area Networking" if you find this chapter difficult to follow.

The overall capabilities and duties of the device will be related to where the device operates in the OSI model. For example, repeaters work at the Physical layer boosting the data signal over a greater distance on your network (allowing you to beat *attenuation* on long cable runs—attenuation is the degradation of the data signal over the run of the cable). On the other end of the internetworking spectrum, gateways operate at the upper layers of the OSI model (such as the Application and Presentation layers) and provide a way to connect computer systems using unlike network protocols (such as connecting an Ethernet LAN to an IBM AS400 miniframe).

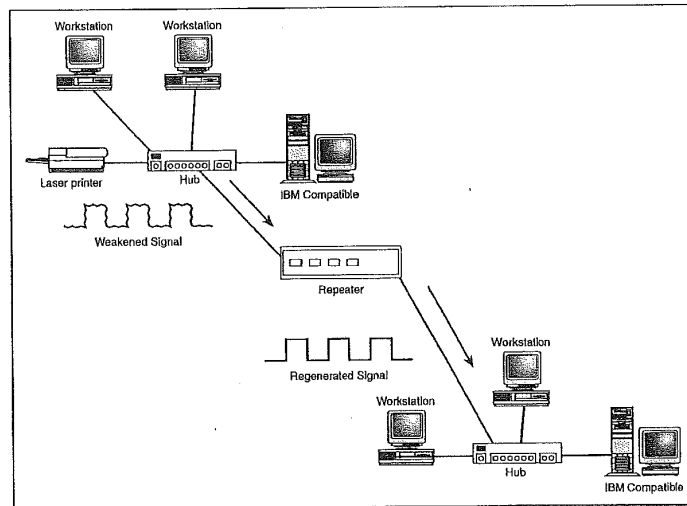
Repeaters are simple devices, whereas gateways require both hardware and software to accomplish the task of allowing very different kinds of networks to communicate. The other internetworking devices that will be discussed, bridges and routers, fall in complexity in between repeaters and gateways.

Repeaters

Repeaters take the signal that they receive from network devices and regenerate the signal so that it maintains its integrity along a longer media run than is normally possible. Because all media types (copper cable, fiber optic cable, and wireless media) must deal with attenuation limiting the possible distance between network nodes, repeaters are a great way to physically enlarge the network.

Because repeaters are Physical layer devices, they don't examine the data packets that they receive, nor are they aware of any of the logical or physical addressing relating to those packets. This means that placing a repeater on a network doesn't slow down the flow of information on the network to any great degree. The repeater just sits on the network boosting the data signals received on one particular segment and passing it back out to another segment on the network as the data makes its way to its final destination (see Figure 4.2).

FIGURE 4.2
Repeaters boost the data signal from one network segment and pass it on to another network segment, extending the size of the network.



Bridges

Bridges are internetworking devices that operate at the Data Link layer of the OSI model. This means that they have greater capabilities (networking-wise) than Layer 1 devices like repeaters and hubs. Bridges are used to segment networks that have grown to a point where the amount of data traffic on the network media is slowing the overall transfer of information.

Bridges (which consist of the bridge hardware and some type of bridge operating system software) have the capability to examine the MAC address (also known as the hardware address; remember it's burned onto the NIC in each computer on the network) on each data packet that is circulating on the network segments that are connected to the bridge. By learning which MAC addresses are residents of the various segments on the overall network, the bridge can help keep data traffic that is local to a particular segment from spreading to the other network segments that are serviced by the bridge.

So basically bridges provide a segmentation strategy for recouping and preserving bandwidth on a larger homogenous network (homogenous meaning that the entire network consists of a particular architecture such as Ethernet). For example, you may segment a larger network using a bridge into three different segments as shown in Figure 4.3.

Let's say that a computer on segment A transmits data that is intended for another computer on segment A. The bridge will examine these data packets (checking out their source and destination MAC addresses), determine that they stay on segment A, and discard the packets. (It doesn't clear the packets from the network; remember that Ethernet is a passive architecture where all the nodes on the network sense the data on the carrier line.) The fact that the bridge doesn't forward the packets to the other segments on the network preserves the bandwidth on those segments (their lines aren't cluttered up by data that isn't intended for the computers on that particular segment).


Internetworking with an Ethernet bent

You will find that as the various internetworking devices and internetworking itself are discussed in this chapter, much of the information relates more directly to Ethernet networks than other architectures such as Token Ring and FDDI. The reason for this is simple: Ethernet is the most commonly employed network architecture, and many internetworking devices were devised because of connectivity issues with Ethernet networks. For a wealth of information on Token Ring and other LAN technologies (related to IBM hardware such as Token Ring and FDDI NICs), check out the white papers offered by IBM on its support Web site at <http://www.networking.ibm.com/nethard.html>. These white papers come in HTML and PDF formats (for Adobe Acrobat Reader) and are a great free resource for network administrators.

A good tutorial on the basics of FDDI can be found at http://www.data.com/tutorials/boring_facts_about_fddi.html. Another good source of networking articles can be found at www.cmpnet.com/, which has links to a large number of sites that provide information on LAN and WAN technologies.

FIGURE 4.3

Bridges segment larger networks to keep segment data traffic localized.



Repeaters, concentrators, and active hubs

Repeaters are also referred to as concentrators. Hubs that have the same signal boosting capabilities as repeaters are referred to as active hubs or multiport repeaters. All these devices (no matter what you call them) operate at the Physical layer of the OSI model.

In another scenario, a computer on segment A transmits data that is intended for a computer on segment C. Again, the bridge will examine the MAC addresses of these packets and in this situation it will forward the packets from segment A to segment C. The bridge is very specific about where it forwards the packets. No packets will be forwarded to segment B.

Although bridging might sound like the ultimate answer to maximizing network throughput, it actually does have some downsides. Bridges forward broadcast packets from the various nodes on the network to all the segments (such as NETBIOS and other broadcasts). Also, in cases in which the bridge is unable to resolve a MAC address to a particular segment on the network, it forwards the packets to all the connected segments.

Switches

Switches are another Layer 2 internetworking device that can be used to preserve the bandwidth on your network using segmentation. Switches are used to forward packets to a particular segment using MAC hardware addressing (the same as bridges). Because switches are hardware-based, they can actually switch packets faster than a bridge.

Switches can also be categorized by how they forward the packets to the appropriate segment. There are store-and-forward switches and cut-through switches.

Switches that employ store-and-forward switching completely process the packet including the CRC check and the determination of the packet addressing. This requires the packet to be stored temporarily before it is forwarded to the appropriate segment. This type of switching cuts down on the number of damaged data packets that are forwarded to the network.

Cut-through switches are faster than store-and-forward switches because they forward the packet as soon as the destination MAC address is read.

Routers

Routers are internetworking devices that operate at the Network layer (Layer 3) of the OSI model. Using a combination of hardware and software (Cisco Routers use the Cisco IOS—Internetwork Operating System), routers are used to connect networks. These networks can be Ethernet, Token Ring, or FDDI—all that is needed to connect these different network architectures is the appropriate interface on the router.

Because routers are Layer 3 devices, they take advantage of logical addressing to move packets between the various networks on the Internetwork. Routers divide the enterprisewide network into logical subnets, which keep local traffic on each specific subnet. And because routers don't forward broadcast packets from a particular subnet to all the subnets on the network, they can prevent broadcast storms from crippling the entire network.

Transparent bridges build a bridging table

Transparent bridges are employed on Ethernet networks; they forward packets (or drop packets that are part of local segment traffic) on the network based on a bridging table. The bridge builds the table by sampling the packets received on its various ports until it has a complete list of the MAC addresses on the network and the particular network segment that they are present on.

Source-routing bridges

Source-routing bridges on Token Ring networks don't work as hard as transparent bridges on Ethernet networks. Source-routing bridges are provided the path for a particular set of packets it receives within the packets themselves. The bridge only needs to follow the directions contained in the packets to forward them to the appropriate segment.

The horror of broadcast storms

Because bridges forward broadcast packets, which can really flood a network with data, bridges don't protect you against broadcast storms. Malfunctioning NICs and other devices can generate a large amount of broadcast packets, resulting in a broadcast storm that can cripple an entire network.

Email gateways

Another common use of gateways is as translators between different email standards. For example, a gateway is used to translate between Lotus Notes Mail server and a Microsoft Exchange Server (an email server).

Because this book is about routers and routing (specifically Cisco Routers and the Cisco IOS), the ins and outs of how routers work and the routing protocols that they use to move packets between subnets are discussed in more detail in Chapter 5, "How a Router Works."

Gateways

Gateways are used to connect networks that don't embrace the same network protocol and so protocol translation is necessary between the two disparate networks. For example, a gateway can be used as the connection between an IBM AS400 miniframe and a PC-based LAN.

Gateways function at the upper layers of the OSI model—the Transport, Session, Presentation, and Application (4, 5, 6, and 7) layers. Gateways typically consist of an actual computer that runs software which provides the appropriate gating software that converts the data between the two unlike computing environments. In our example of the gateway between the IBM AS400 and the PC LAN, the gateway computer might be running Windows NT Server with a special translation software package installed.

Gateways typically are situated on high-speed backbones such as FDDI networks, where they connect a mainframe or miniframe to LANs that are connected to the FDDI backbone via routers (see Figure 4.4). Although gateways are certainly necessary to connect networks where data conversion is necessary, they can slow traffic on the network (especially the data traffic moved between the two connected networks). And because gateways typically connect very different systems, their configuration can be relatively more complex than other internetworking devices (*relatively* is the key word; don't ever try to tell someone who configures routers that setting up a gateway is a more difficult task).




FIGURE 4.4
Gateways provide the connecting point between high-speed backbones and mainframe and miniframe computers.

Building a Campus Network

Before leaving the subject of internetworking, a few words should be said about network scale. A Campus network is defined as a portion of the enterprise network that serves an entire corporation or institution. Network campuses usually are limited to a building or group of buildings and primarily use LAN technologies, such as Ethernet, Token Ring, and FDDI.

Building and maintaining a campus-sized network is really a study in connecting different LAN architectures (using routers) and taking advantage of internetworking devices that help relieve congestion on the network (such as switches and bridges).

Networking the enterprise—connecting the various campus networks—requires the use of WAN technologies, which also employ internetworking devices, particularly routers with the appropriate WAN interfaces.

The next chapter discusses how a router works. This should help you take the puzzle pieces that were provided to you in Chapters 1, 3, and 4 and allow you to better understand how LANs can become WANs and how networking the enterprise isn't an insurmountable task (at least in theory).

I thought routers were gateways

When you configure a particular computer on a network (particularly on a TCP/IP network), you must configure the default gateway for the node. The default gateway is typically the logical address of the router port that the node (and the rest of its subnet) connects to. Don't confuse router interfaces (when they are referred to as gateways) with actual gateways that translate data between two different computer systems.

chapter

5

How a Router Works

Routing Basics

Routable Protocols

Routing Protocols

Routing Protocol Basics

Types of Routing Protocols

-
-
-
-
-

Routing Basics

In cases where information needs to be moved between two networks, an internetworking device, called a *router* (you learned a little bit about routers in Chapter 4, "Internetworking Basics"), is responsible for the movement of this data. Routing data on an internetwork requires that a couple of different events take place: an appropriate path for the packets must be determined, and then the packets must be moved toward their final destination.

Both path determination and routing of packets (or *switching* as it is also referred to—packets are switched from an incoming interface to an outgoing interface on the router) take place at layer 3 (Network layer) of the OSI model. Another important layer 3 event is the resolution of logical addresses (such as IP numbers when TCP/IP is the routed protocol) to actual hardware addresses. Additional discussion related to these three layer 3 events will give you a better idea of the overall routing process.

SEE ALSO

- *To review the OSI model before continuing with this chapter, see page 35.*

Path Determination

As discussed in Chapter 4, routers enable you to divide a large network into logical subnets; doing so keeps network traffic local on each subnet, enabling you to take better advantage of the bandwidth available. It's then the job of the router to move data packets between these subnets when necessary. Routers can also serve as the connective device between your network (all your subnets are viewed by other enterprise networks as a single network even though you've divided them into logical parts). Routers also can serve as the connective device to other networks to which your network may be attached. The best example of many different networks connected for communication purposes is the Internet.

For the purpose of discussion, let's create a network that contains subnets that are connected by a router. You will also create a logical addressing system.

Understanding subnets

Creating subnets is an extremely important part of implementing routing on a network. For now, understand that subnets are logical divisions of a larger corporate network. Creating subnets in a TCP/IP environment will be discussed in great detail in Chapter 10, "TCP/IP Primer."

Figure 5.1 shows a network that has been divided into two subnets using a router. The type of connections between the subnets (Ethernet, Token Ring, and so on) and the router aren't important at this point in our discussion, so just suppose that the appropriate protocols and interface connections would be used to connect these subnets to the router.

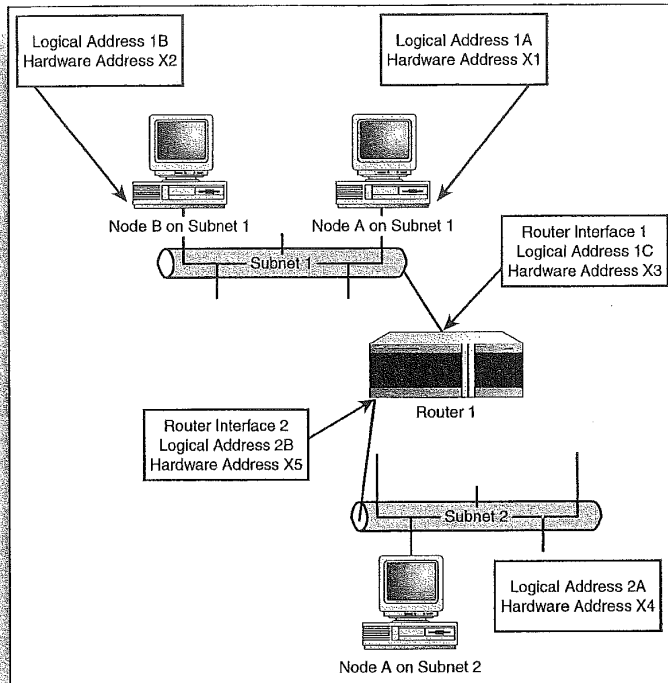


FIGURE 5.1
A network divided into two logical subnets.

In this example, the router has two network interfaces, Interface 1 and Interface 2, which are connected to Subnet 1 and Subnet 2, respectively. The logical addressing system that is used to address the various nodes on the network (logical addresses must be assigned to each interface on the router as well) is the subnet number followed by a letter designation. So, Node A on Subnet 1 is assigned the logical address 1A (subnet designation then node designation).

Don't try this at home

Be advised that the logical addresses that you assign to your nodes and router interfaces are for our discussion of how the router determines when and when not to forward frames to a network. These aren't real logical addresses. Real logical addresses such as IP addresses would be used on a real-world network.

Real-world addresses

To give you an idea of what the addresses for these various router interfaces and nodes would be in a real IP network, each node and interface is listed below with a Class B IP address:

Subnet 1: 130.10.16.0

Node A: 130.10.16.2

Node B: 130.10.16.3

Router Interface 1:
130.10.16.1

Subnet 2: 130.10.32.0

Node A: 130.10.32.2

Router Interface 2:
130.10.32.1

Notice that subnetting has taken place on the network and the Subnet 1 nodes and router interface have the third octet value of 16 and the Subnet 2 nodes and router interface have a third octet value of 32; these different numbers identify the different subnets used. You will learn all about this in Chapter 10, "TCP/IP Primer."

Each node on the network will also have a hardware address (remember that a hardware address is actually assigned to each NIC when they are built at the factory; router interfaces are also assigned a burned-in hardware address when they are manufactured). For ease of discussion, the hardware addresses for each of the nodes is an X followed by a number. For example, the hardware address for Node A on Subnet 2 is X4 (remember all hardware addresses are different, that's how the cards are manufactured).

Now that you have a small internetwork, let's take a look at what happens when one of the computers attempts to send packets to another computer on the network.

Logical and Hardware Addresses

When you connect networks using a router, you end up with two different types of data traffic. You end up with local data traffic, where nodes on the same subnet communicate with each other. You also have network traffic where nodes on different subnets are communicating with each other. This type of traffic must pass through the router. The next two sections explain how communication within a subnet and communication between subnets take place.

Communication on the Same Subnet

First, let's look at a situation in which two computers on the same subnet communicate. Node A on Subnet 1 must send data to Node B on Subnet 1. Node A knows that the packets must go to the logical address 1B and Node A knows that 1B resides on the same subnet (so in this case the router will not actively be involved in the movement of packets). However, the logical address 1B must be resolved to an actual hardware address.

Now, Node A might already know that logical address 1B actually refers to the hardware address X2. Computers actually maintain small memory caches where they keep this type of logical-to-hardware address-resolution information. If Node A has no idea what the hardware address of logical address 1B is, it will send a message out to the network asking for the logical address 1B to be resolved to a hardware address. When it receives the information, it will send the packets to Node B, which accepts the packets because they are tagged with its hardware address—X2.

As you can see, node-to-node communication on the same subnet is pretty straightforward.

Communication Between Different Subnets

Now let's look at a scenario where a computer wants to send data to a computer on another subnet.

Node A on Subnet 1 wants to send data to Node A on Subnet 2. So, Node A on Subnet 1 wants to send the data to logical address 2A. Node A on Subnet 1 knows that address 2A isn't on the local subnet, so it will send the packets to its *default gateway*, which is the router interface that is connected to Subnet 1. In this case the logical address of the Node A (on Subnet 1) gateway is 1C. However, again this logical address must be resolved to a hardware address—the actual hardware address of Router Interface 1.

Again, using broadcast messages, Node 1 on Subnet 1 receives the hardware address information related to logical address 1C (the hardware address is X3) and sends the packets on to Router 1 via Router Interface 1. Now that the router has the packets, it must determine how to forward the packets so that they end up at the destination node. It will take a look at its routing table and then switch the packets to the interface that is connected to the destination subnet.

Packet Switching

After the router has the packets, packet switching comes into play. This means that the router will move the packets from the router interface that they came in on and switch them over to the router interface connected to the subnet they must go out on. However, in some cases the packets might have to pass through more than one router to reach the final destination. In our example, only one router is involved. Router 1 knows that the logical address 2A is on Subnet 2. So the packets will be switched from Router Interface 1 to Router Interface 2.

Again, broadcast messages are used to resolve logical address 2A to the actual hardware address X4. The packets are addressed appropriately and then forwarded by the router to Subnet 2. When Node A on Subnet 2 sees the packets with the Hardware Address X4, it grabs the packets.

Nodes collect addressing information
Computers use broadcast messages and tables of information (that they build from broadcast information placed out on the network by other computers) to determine which addresses are local and which addresses are remote on an internetwork.

So, you can see that routing involves both the use of logical addressing and hardware addressing to get packets from a sending computer to a destination computer. Each routable protocol (TCP/IP versus IPX/SPX) uses a slightly different scheme to resolve logical addresses to hardware addresses, but the overall theory is pretty much the same as outlined here (TCP/IP addressing was used as the model for our discussion).

Routing Tables

Before I finish this basic discussion of routing, we should discuss how the router determines which router port it switches the packets to (this information will be reviewed when IP routing is discussed in Chapter 11, "Configuring IP Routing"). Routers use software to create routing tables. These routing tables contain information on which the hardware interface on the router is the beginning route (for the router) that will eventually get the packets to the destination address.

Routers, however, aren't concerned with individual node addresses when they build their routing tables; they are only concerned with getting a particular set of packets to the appropriate network. For example, using your logical addressing system from Figure 5.1, a router's routing table would appear as shown in Table 5.1. Notice that each router interface is mapped to a particular subnet. That way the router knows that when it examines the logical address of a packet, it can determine which subnet to forward the packets to.

Where do routing tables come from?

Routing tables actually have two sources. In *static routing*, the network administrator actually types in the different routes that are available between segments on the internetwork. These network administrator-created routing tables use a series of router commands to build a table that looks somewhat like Figure 5.1. Routing tables can also be built dynamically by routing protocols such as RIP and IGRP (which are discussed later in this chapter). Dynamic routing tables also end up looking like a table (again somewhat like Figure 5.1).

Table 5.1 A Basic Routing Table for Router 1

Subnet Logical Designation	Router Interface
1	1
2	2

Basically, this routing table means that packets that are destined for any node on Subnet 1 would be routed to the Router 1 Interface on the router. Any packets destined for Subnet 2 would be switched to the Router Interface 2 (just as I discussed earlier). Obviously, the logical designation for a subnet on a real-world network would consist

of something like a network IP address, such as 129.10.1.0, which designates a class B IP subnetwork. And the router interface would be designated by the type of network architecture it supports, such as E0 for the primary Ethernet interface, or S0 for the primary serial interface on the router.

When multiple routers are involved—on larger networks—the routing tables become populated with more information. For example, let's expand your one router, two-subnet network into five subnets that employ two routers. Figure 5.2 shows this network.

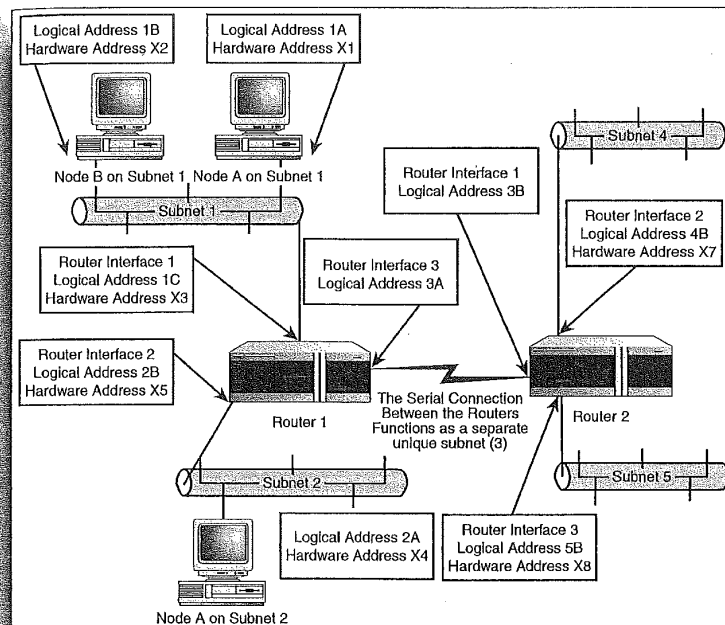


FIGURE 5.2
A network divided into five logical subnets that use two routers.

Now, you might be thinking that you see only four subnets. Actually, any serial connection between two routers is, in effect, a separate subnet and must be provided with unique logical addresses.

With the size of the network expanded and the number of subnets increased, Router 1 will have a decidedly different routing table. It now must potentially pass on packets that go to nodes on Subnets 4

and 5. However, as I stated earlier, a router doesn't worry about getting the packets to the actual recipient nodes; it only forwards the packets so that they get to the correct subnet.

Table 5.2 shows a routing table for Router 1 using your (fictional) logical addressing system for your subnets. Notice that Router 1 forwards packets for Subnets 4 and 5 through the same interface—its Interface 3. So, Router 1 is content with forwarding packets for Subnets 4 and 5 (sent from Subnets 1 or 2) to Router 2. Router 2 is then responsible for switching the packets to the correct interface that is connected to the appropriate subnet.

Table 5.2 An Expanded Routing Table for Router 1

Subnet Logical Designation	Router Interface
1	1
2	2
4	3
5	3

Router 2 would have a similar routing table that would designate that all packets for Subnets 1 and 2 be routed out of its Interface 1 to Router 1. Router 1 would then handle the routing of the packets to the appropriate subnet.

All these routing decisions made by the routers will involve software. Software that is responsible for network transport (network, or *routable*, protocols such as TCP/IP, IPX/SPX, and AppleTalk) and software that helps the router determine the best path for a set of packets to the next step in their journey to a final node destination. This type of software is called a *routing protocol*. Routable protocols (network protocols that can be routed) and routing protocols will be discussed in the next two sections.

SEE ALSO

➤ For more information on IP routing and routing tables, see page 195.

Routable Protocols

Before you take a look at the protocols that determine the path for packets routed through the router (and also maintain the routing table used by the router to forward the packets), a few words should be said about routable or *routed protocols*. Chapter 2, "The OSI Model and Network Protocols," discussed commonly used network protocols: TCP/IP, IPX/SPX, AppleTalk, and NetBEUI. Of these four protocols only TCP/IP, IPX/SPX, and AppleTalk are routable. This is because these three protocols all provide enough information in the Network layer header of their packets for the data to be sent from sending node to destination node even when the packets must be forwarded across different networks (by a device such as a router).

SEE ALSO

- To review network protocols such as TCP/IP, see page 44.

Routing Protocols

Whereas routable protocols provide the logical addressing system that makes routing possible, *routing protocols* provide the mechanisms for maintaining router routing tables. Routing protocols allow routers to communicate, which allows them to share route information that is used to build and maintain routing tables.

Several different routing protocols exist, such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Enhanced Interior Gateway Protocol (EIGRP). And while these different routing protocols use different methods for determining the best path for packets routed from one network to another, each basically serves the same purpose. They help accumulate routing information related to a specific routed protocol such as TCP/IP (IP is the routed portion of the TCP/IP stack).

It's not uncommon in LANs and WANs to find host and server machines running more than one network protocol to communicate. For example, an NT server in a *NT Domain* (an NT Domain is a network managed by an NT server called the Primary Domain Controller) may use TCP/IP to communicate with its member

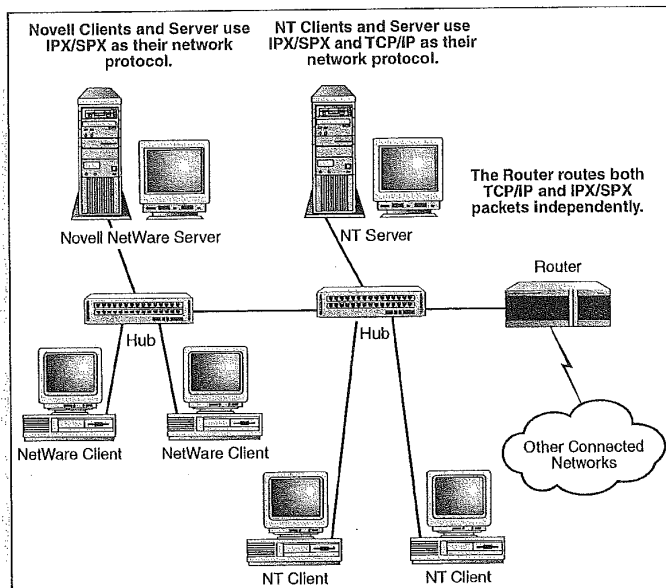
Why isn't NetBEUI routable?

NetBEUI does provide a logical naming system to deliver packets to computers; it uses NetBIOS names, (the name you give your computer when you set it up), which are then resolved to MAC addresses on computers using a series of NetBIOS broadcasts. Unfortunately, the NetBIOS naming system doesn't have a Network layer logical addressing system that can be used to direct packets across a router on an internetwork. NetBIOS names just don't provide enough information (no network information at all) for the packets to be moved between the various networks connected by a router. Plus the NetBEUI/NetBIOS network stack doesn't contain a routing protocol. So, in NetBEUI's case it has two strikes and no route.

clients. But it may also serve as a gateway to various printers and file servers that use the Novell NetWare operating system; meaning that the NT server will also embrace IPX/SPX as a network protocol. These protocols basically operate in their own tracks simultaneously and do not interfere with each other (see Figure 5.3).

This same concept of simultaneously but independently running protocols is also embraced by routing protocols. Multiple independent routing protocols can run on the same router, building and updating routing tables for several different routed protocols. This means that the same network media can actually support different types of networking.

FIGURE 5.3
Networks can embrace multiple network protocols, and routers can simultaneously route multiple network protocols using multiple routing protocols.



SEE ALSO

- For a quick look at two theoretical routing tables, see page 82.
- For more information on the types of routing protocols and specific routing protocols, see page 91.

Routing Protocol Basics

Routing protocols must not only provide information for router routing tables (and be able to adequately update routers when routing paths change), they are also responsible for determining the best route through an internetwork for data packets as they move from the sending computer to the destination computer. Routing protocols are designed to optimize routes on an internetwork and also to be stable and flexible.

Routing protocols are also designed to use little processing overhead as they determine and provide route information. This means that the router itself doesn't have to be a mega computer with several processors to handle the routing of packets. The next section discusses the mechanism that routing protocols use to determine paths.

Routing Algorithms

An algorithm is a mathematical process that is used to arrive at a particular solution. In terms of routing protocols, you can think of the algorithm as the set of rules or process that the routing protocol uses to determine the desirability of paths on the internetwork for the movement of packets. The routing algorithm is used to build the routing table used by the router as it forwards packets.

Routing algorithms come in two basic flavors: *static* and *dynamic* algorithms. Static algorithms aren't really a process at all, but consist of internetwork mapping information that a network administrator enters into the router's routing table. This table would dictate how packets are moved from one point to another on the network. All routes on the network would be static—meaning unchanging.

The problem with static algorithms (other than it's a real pain to have to manually enter this information on several routers) is that the router cannot adapt to changes in the network topology. If a particular route becomes disabled or a portion of the internetwork goes down, there is no way for the routers on the network to adapt to these changes and update their routing tables so that data packets continue to move toward their final destinations.

Routed protocols and routing protocols are configured on the router

Although this chapter delves into the theoretical aspects of how a router works and discusses the relationship between routed and routing protocols, keep in mind that these are all issues that you deal with on the router when you actually configure it. The Cisco IOS provides the commands and functions that enable you to set the routed and routing protocols used by a specific router. More on the Cisco IOS is discussed in Chapter 9, "Working with the Cisco IOS."

Convergence is the key for dynamic routing protocols

When an internetwork experiences a downed link or some other network problem, it's very important for all the routers on the network to update their routing tables accordingly. *Convergence* is the time it takes for all the routers on the network to be up-to-date in terms of the changes that have taken place in the network topology (such as the unavailability of a certain route because of a downed line). The longer it takes for all the routers on the internetwork to converge, the greater the possibility that packets will be routed to routes that are no longer available on the network. This type of problem is certainly not unheard of on the Internet either, and this is why email can end up traveling a road to nowhere and never get to its destination.

Dynamic algorithms are built and maintained by routing update messages. Messages that provide information on changes in the network prompt the routing software to recalculate its algorithm and update the router's routing table appropriately.

Routing algorithms (and the routing protocols that employ a certain algorithm) can also be further classified based on how they provide update information to the various routers on the internetwork.

Distance-vector routing algorithms send out update messages at a prescribed time (such as every 30 seconds—an example is the Routing Information Protocol—RIP). Routers using distance-vector algorithms pass their entire routing table to their nearest router neighbors (routers that they are directly connected to). This basically sets up an update system that reacts to a change in the network like a line of dominos falling. Each router in turn informs its nearest router neighbors that a change has occurred in the network.

For example, in Figure 5.4, Router 1 realizes that the connection to Network A has gone down. In its update message (sent at 30-second intervals), it sends a revised routing table to Router 2 letting its neighbor know that the path to Network A is no longer available. At its next update message, Router 2 sends a revised routing table to Router 3, letting Router 3 know that Router 2 no longer serves as a path to Network A. This updating strategy continues until all the routers on the network know that the Network A line is no longer a valid path to the computers on that particular part of the entire internetwork.

The downside of distance-vector routing is that routers are basically using hearsay information to build their routing tables; they aren't privy to an actual view of a particular router's interface connections. They must rely on information from a particular router as to the status of its connections.

Another strategy for updating routing tables on an internetwork is the link-state routing algorithm. Link-state routing protocols not only identify their nearest neighbor routers, but they also exchange link-state packets that inform all the routers on the internetwork about the status of their various interfaces. This means that only information on a router's direct connections is sent, not the entire routing table as in distance-vector routing.

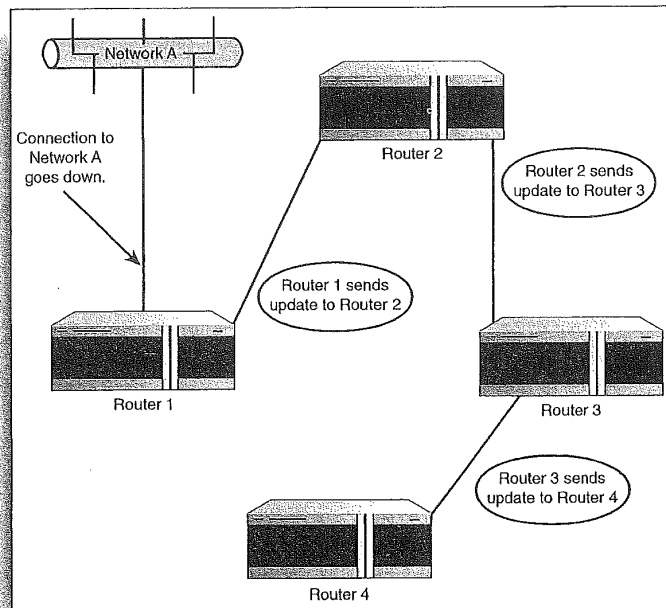


FIGURE 5.4
In distance-vector routing, nearest neighbors provide updated routing tables.

This also means that link-state routers are able to build a more comprehensive picture of the entire internetwork and make more intelligent decisions when choosing paths for the routing of packets. Convergence also takes place more rapidly on a link-state routing system than it does when distance-vector routing is used.

Routing Metrics

Now that you have learned the different types of routing algorithms (static versus dynamic) and the two ways that they update their router tables (distance vector versus link state), you should take a look at how routing protocols actually determine the best route between a sending computer and a destination computer when more than one route is available.

Static versus dynamic routing

Although you might get the impression that dynamic routing is a much better way to manage the demands of internetworking (when compared to static routing), dynamic routing does require more overhead (in terms of bandwidth and processing power) from internetworking devices such as routers because of all the broadcast messages and editing of the routing tables. Dynamic routing is, obviously, a much more "fun" process to monitor. However, in some cases, setting up static routing tables can provide an overall faster throughput on the network as packets are routed.