SIXTH EDITION

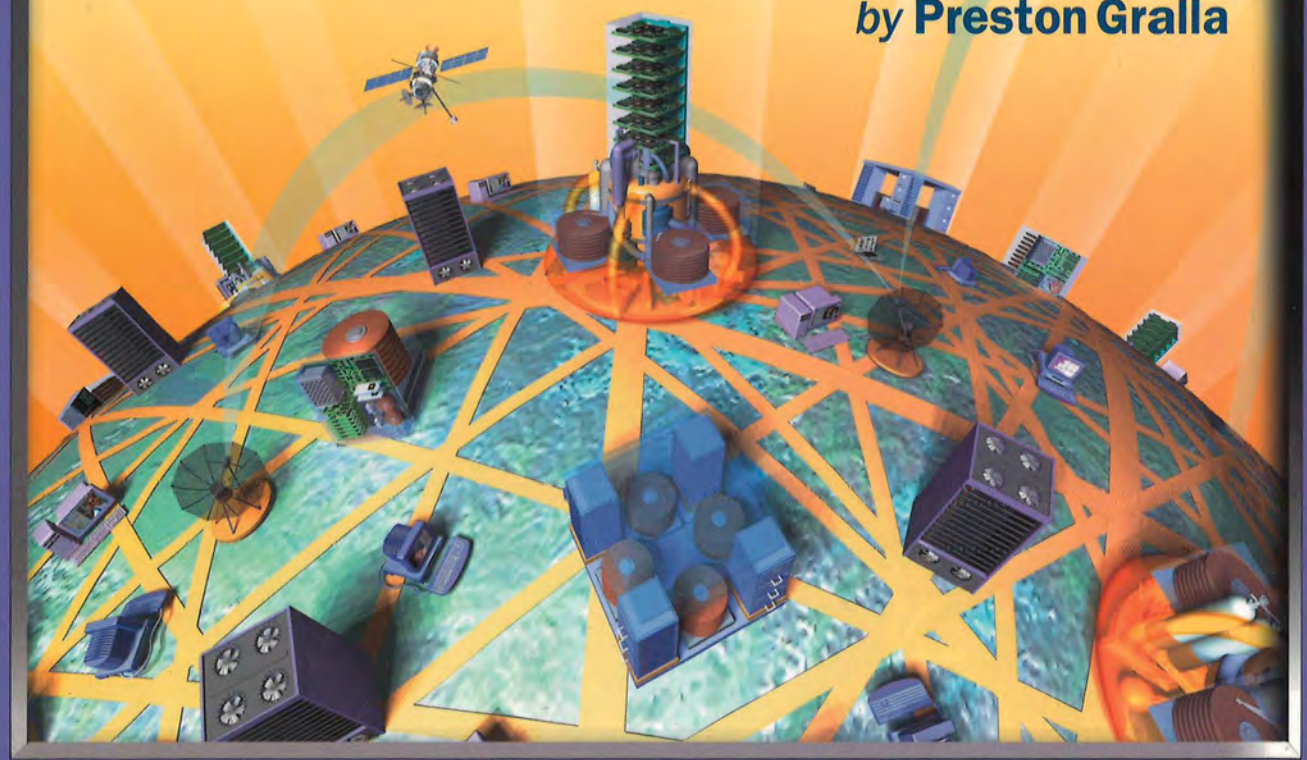
NATIONAL BESTSELLER
MORE THAN 2,000,000 How It Works SOLD

HOW THE INTERNET WORKS

"*How the Internet Works* is a richly-illustrated guide to the magic behind this wonderful, world-transforming technology—a crystal-clear labor of love from one of the Web's true editorial pioneers."

Sam Vincent Meddis
Technology Editor
USATODAY.com

by **Preston Gralla**



que

Illustrated by **MICHAEL TROLLER**

APPLE 1005

ABOUT THE AUTHOR



Preston Gralla is the award-winning author of 20 books, including *How Wireless Works*, *How To Expand and Upgrade PCs*, and *The Complete Idiot's Guide to Protecting Yourself Online*. He is an executive editor and columnist for CNet and ZDNet; is a technology columnist for the *Dallas Morning News*; and has written about technology for many magazines and newspapers, including *USA Today*, *PC Magazine*, the *Los Angeles Times*, *Boston Magazine*, *PC/Computing*, *Computerworld*, and *FamilyPC* among many others. Gralla has won several writing and editing awards, including one from the Computer Press Association for the best feature article in a computer magazine.

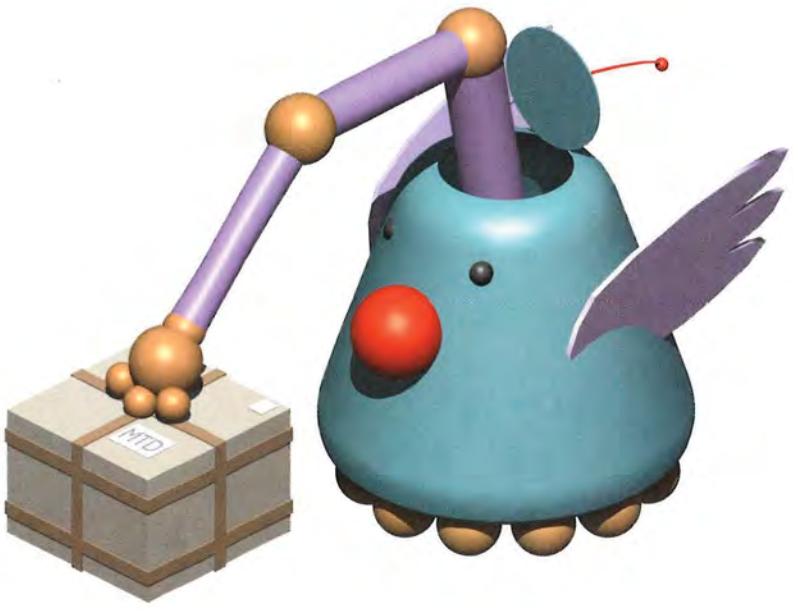
As a well-known expert on computers and the Internet, he has appeared frequently on numerous TV and radio shows and networks, including the *CBS Early Show*, CNN, National Public Radio's *All Things Considered*, MSNBC, CNBC, TechTV, and CNet Radio.

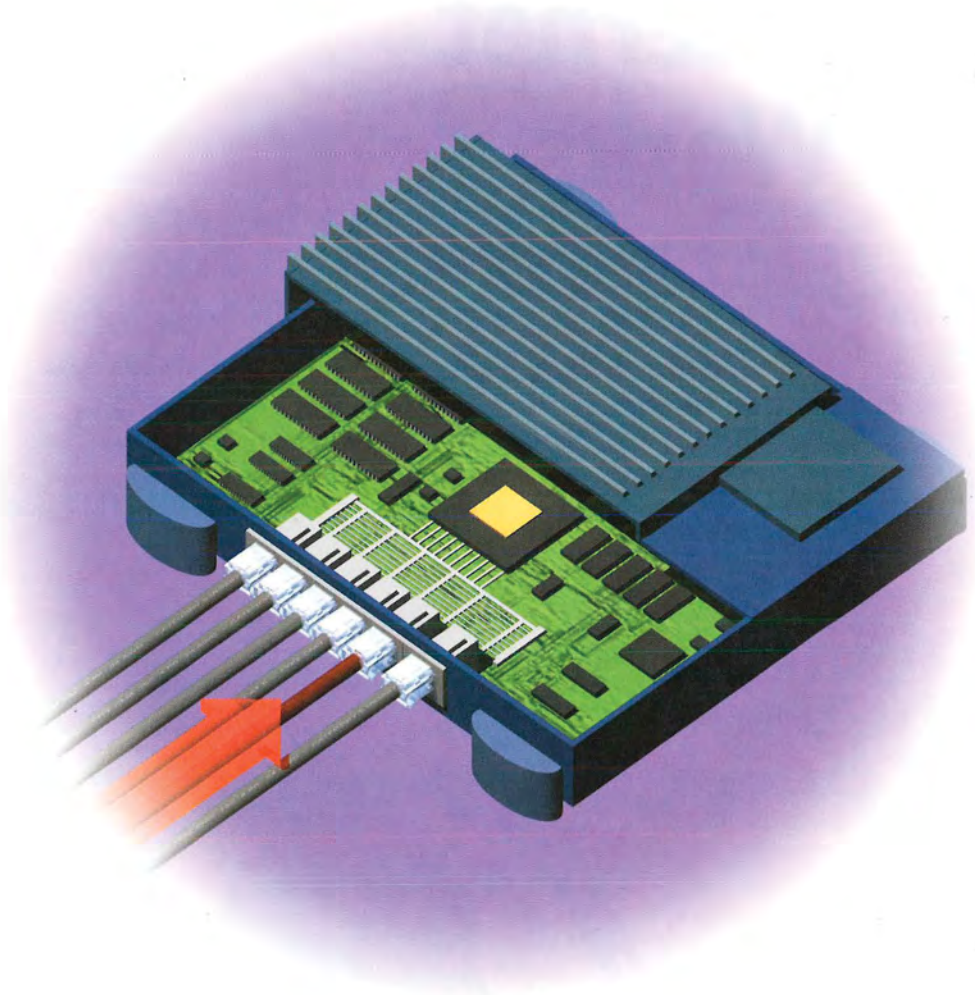
He was the founding managing editor of the well-known newspaper *PC Week* and a founding editor of *PC/Computing*. Under his editorship, *PC/Computing* was a finalist for General Excellence from the National Magazine Awards.

Gralla lives in Cambridge, Massachusetts, with his wife Lydia, children Gabriel and Mia, and a rabbit named Polichinelle. He also writes the free *Gralla's Internet Insider* e-mail newsletter. To subscribe to it for free, send an e-mail to `preston@gralla.com` with the words SUBSCRIBE NETINSIDER on the subject line.

How the Internet Works

Sixth Edition





How the Internet Works

Sixth Edition

Preston Gralla

QUE®
201 W. 103rd Street
Indianapolis, IN 46290

How the Internet Works, Sixth Edition
Copyright © 2002 by Que

Associate Publisher	Greg Wiegand
Acquisitions Editor	Stephanie J. McComb
Development Editor	Nicholas J. Goetz
Managing Editor	Thomas F. Hayes
Project Editor	Karen S. Shields
Production Editors	Megan Wade Amy Jay
Indexer	Erika Millen
Technical Editor	Chris Faust
Illustrator	Michael Troller
Team Coordinator	Sharry Lee Gregory
Page Layout	Gloria Schurick

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. Nor is any liability assumed for damages resulting from the use of the information contained herein.

International Standard Book Number: 0-7897-2582-7

Library of Congress Catalog Card Number: 20-01087727

Printed in the United States of America

First Printing: September, 2001

04 03 02 01 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 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.

Introduction

Part 1: Understanding the Internet's Underlying Architecture 1

Chapter 1
What Is the Internet? 4

Chapter 2
How Computer Networks Send Data Across the Internet 8

Chapter 3
How TCP/IP Works 12

Chapter 4
Understanding the Internet's Software Structure 16

Chapter 5
How Internet Addresses and Domains Work 20

Chapter 6
How Routers Work 28

Chapter 7
How Internet File Types Work 32

Part 2: Connecting to the Internet 36

Chapter 8
How Computers Connect to the Internet 40

Chapter 9
How a Modem Works 44

Chapter 10
How Internet/Television Connections Work 48

Chapter 11
How the Digital Subscriber Line (DSL) Works 56

Chapter 12
How ISDN Works 60

Chapter 13
How You Connect to the Internet from Online Services 64

Chapter 14
How Wireless and Palmtop Connections Work 68

Chapter 15
How Home Networks Work 78

Part 3: Communicating on the Internet 84

Chapter 16
How E-mail Works 88

Chapter 17
How E-mail Spam Works 100

Chapter 18
How Newsgroups Work 106

Chapter 18
How Internet Chat and Instant Messaging Work 110

Chapter 20
How Internet Phone Calls Work 116

Part 4: Using the World Wide Web 122

Chapter 21
How Web Pages Work 126

Chapter 22
How Web Browsers Work 134

Chapter 23
How Markup Languages Work 140

Chapter 24
How Hypertext Works 148

Chapter 25
How URLs Work 152

Chapter 26
How Image Maps and Interactive Forms Work 156

Chapter 27
How Web Host Servers Work 160

Chapter 28
How Web Sites Work with Databases 164

Chapter 29
How Internet-Based Software Works 168

Part 5: Using Common Internet Tools 172

Chapter 30
How Tenet Works 176

Chapter 31
How FTP Downloading Works 180

Chapter 32
How Internet Searching Works 186

Chapter 33
How Agents Work 192

Chapter 34
How Java, ActiveX, and JavaScript Work 196

Chapter 35
How CGI Scripting Works 204

Part 6: Enjoying Entertainment and Multimedia on the Internet 208

Chapter 36
How Music and Audio Work on the Internet 212

Chapter37
How Napster and File Sharing Work 220

Chapter 38
How Video on the Internet Works 224

Chapter 39
How Multicast IP and the MBone Work 232

Chapter 40
How Virtual Reality Works 236

Chapter 41
Animation on the Web 240

Part 7: Shopping and Doing Business on the Internet 248

Chapter 42
How Intranets Work 252

Chapter 43
Shopping on the Internet 258

Part 8: Protecting Yourself on the Internet 268

Chapter 44
How Firewalls Work 272

Chapter 45
How Hackers Can Cripple the Internet and Attack Your PC 280

Chapter 46
How Viruses Work 286

Chapter 47
How Internet Sites Can Invade Your Privacy 292

Chapter 48
Cryptography, Privacy, and Digital Certificates 302

Chapter 49
How the FBI's "Carnivore" Program Works 308

Chapter 50
Parental Controls on the Internet 312

Glossary 316

Index 326

Acknowledgments

THIS book, like the Internet, is a collaborative work. My name might be on the cover, but I am far from the only person involved in its creation.

Development editor Nick Goetz managed to pull together all the elements of a complex project, all the while keeping a cool head, improving the illustrations, and offering editorial insight and help. And acquisitions editor Stephanie McComb was instrumental in getting the book off the ground, and as always, was a pleasure to work with.

The illustrator for the book, Michael Troller, deserves much thanks. Not only did he create new illustrations from scratch specifically for this book, he redid all the ones from past books and made this edition a more coherent, richer, more visually pleasing experience.

Thanks also go to illustrators of past editions: Mina Reimer, Sarah Ishida, Shelly Norris, and Stephen Adams. And many thanks to the entire team at Que who produced the book, including Karen Shields, Megan Wade, Amy Jay, Gloria Schurick, and Erika Millen.

Thanks also have to go to the many, many people I interviewed for this book. People from Quarterdeck Corporation, Chaco Communications, Progressive Networks, White Pine Software, Microsoft, Netscape, Headspace, SurfWatch Software, WebTV, Accrue, VDONet Corporation, America Online, Yahoo!, Hilgraeve, Fairmarket, eWallet, and Nuborn Technologies are only a few of the folks who gave their time to help me understand the nitty-gritty of how various Internet technologies work. Tim Smith from ZDNet offered me vital help as well.

I gleaned much information from the many FAQs and similar documents widely available on the Internet. I'd like to thank the anonymous authors of those documents, whoever they are.

Many thanks go to Chris Faust, the technical editor for the book, who did a superb job of ensuring that I always got the information right.

Finally, big thanks have to go to my wife, Lydia. She put up with those occasional glassy-eyed looks that were replies to simple questions like, "Did you leave your keys in the refrigerator again?" She also endured my extreme absent-mindedness while I was figuring out ways to explain how firewalls, ISDN, or Web robots work when I should have been concentrating on more immediate matters.

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, I welcome your comments. You can fax, e-mail, 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

E-mail: `feedback@quepublishing.com`

Mail: Greg Wiegand
Associate Publisher
Que
201 West 103rd Street
Indianapolis, IN 46290 USA

Introduction

IN the course of cruising the World Wide Web and clicking a link, have you ever wondered, "How does that work?" Or perhaps this question popped into your mind while you were transferring a file to your computer via FTP, or reading a newsgroup message, or when you first heard about technologies such as spam, cookies, and firewalls. Maybe you've wondered how a message sent from your computer travels through the vastness of cyberspace and ends up in the right e-mail box halfway across the world. Have you ever wanted to know how search tools find the exact piece of information you want out of the millions of pieces of information on the whole Internet? How can you listen to music and view animations while surfing the Web?

This book is designed for everyone interested in the Internet. Its guiding principle is this: No matter how much of a cyberpro you are—or how much of a novice—there's a lot you don't understand about the Internet. Here's just one small example. I have a friend who has made his living with companies involved with the Internet for many years. He's a complete cyberpro who lives and breathes the Internet. One day, he almost whispered to me, "I don't like to admit this, but I don't know what a proxy server is. How does it work, anyway?"

He's not alone. The Internet changes so quickly and the technology advances so rapidly that it can seem almost impossible to keep up with all of it. If you're like just about everyone else involved in the Internet, your questions are similar to those of my friend. You'll find your answers here.

In Part 1, "Understanding the Internet's Underlying Architecture," I explain the underlying basics of the Internet: who runs it, how TCP/IP works, how to understand Internet addresses and domains, and similar topics. Here's where you'll find out about things such as routers and how the client/server architecture underpins virtually every aspect of the Internet.

Part 2, "Connecting to the Internet," depicts the various ways you can connect your computer to the Internet. Here's where to turn if you're interested in any of the following topics: how a cable modem, a digital subscriber line (DSL), or ISDN works; how online services connect to the Internet; how you can connect by satellite to the Internet; how "palmtop" computers, such as 3COM's Palm, can browse the Web; how a modem works; how you can access the Internet from your own home network; and a host of similar subjects.

Part 3, "Communicating on the Internet," covers every aspect of Internet communications. It shows how e-mail and newsgroups work, how IRC chat works, what e-mail "spam" is and what you can do to prevent it, how instant messaging works, and how you can use the Internet to make telephone calls anywhere in the world.

Part 4, "Using the World Wide Web," covers what has become by far the most popular part of the Internet—the World Wide Web. You'll learn virtually every aspect of how the Web works. It delves into how browsers work, how Web server software works, and how Hypertext Markup Language (HTML) works. This section also covers the ways in which the Web is becoming integrated directly into your computer, how Web pages are published and organized on a site, and every other aspect of the Web that is likely to be of interest to you.

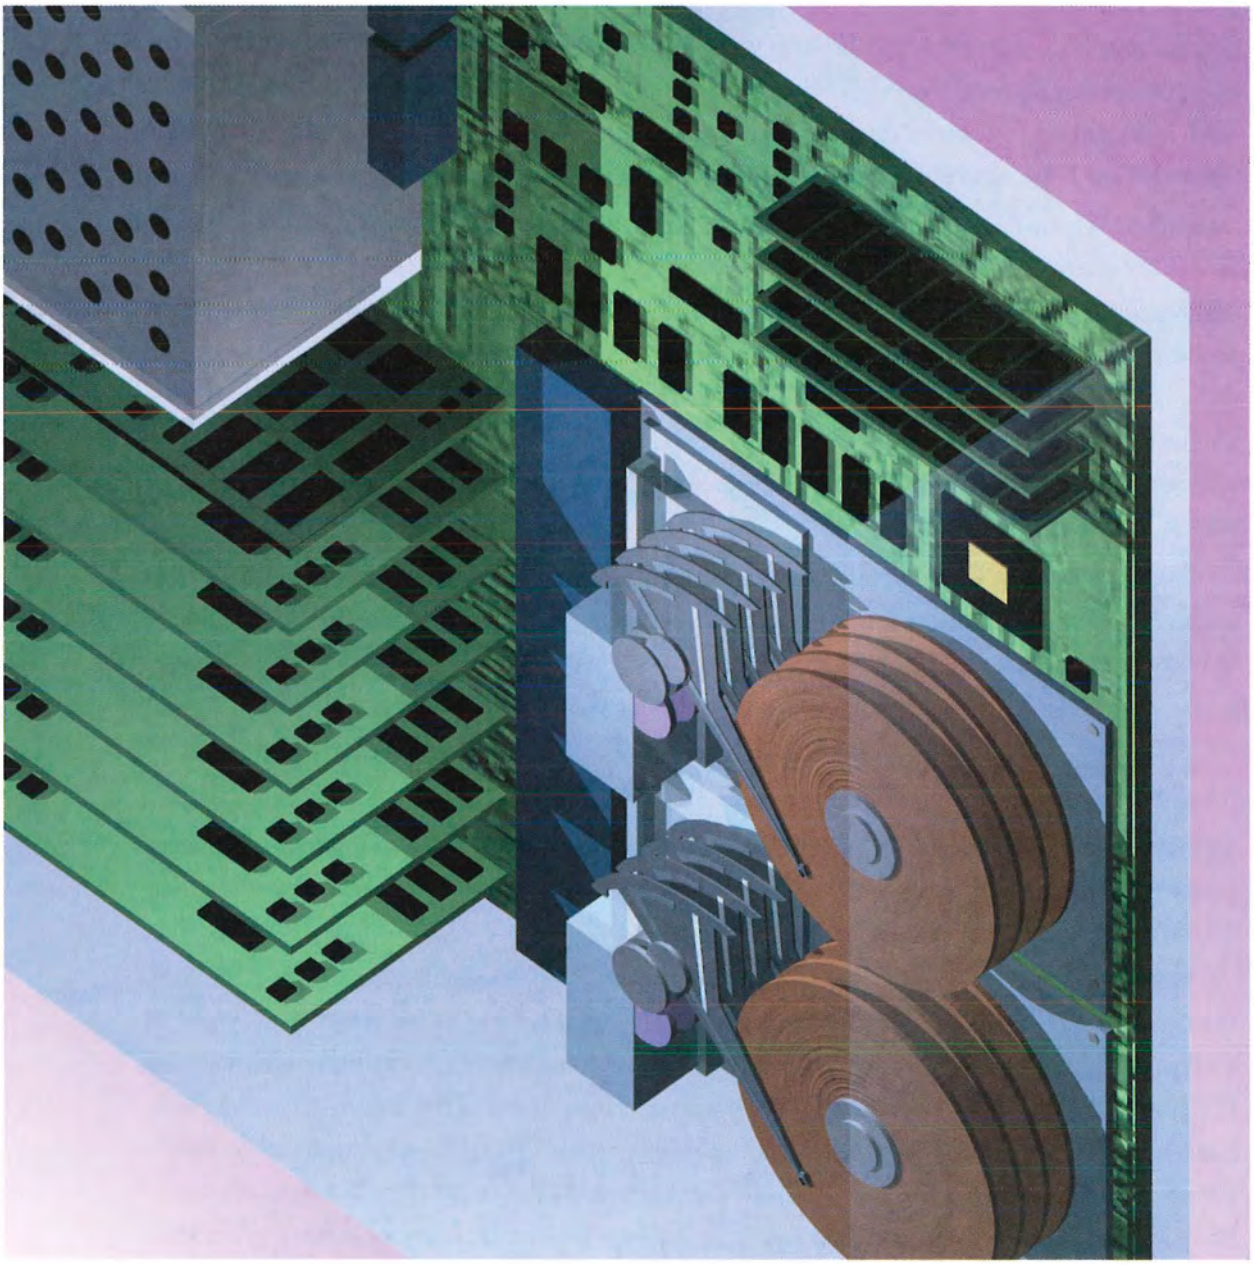
Part 5, "Using Common Internet Tools," takes a close look at common Internet tools. Here's the place to learn about basic tools and services, such as telnet, and what happens when you use FTP to download a file to your computer. You'll also learn how search engines and search tools work. And it's the section of the book that covers cutting-edge Internet technology. You'll find out about how push technology works, how the Java and ActiveX programming languages work, how JavaScript works, how agents can silently do your bidding for you, and how the Common Gateway Interface (CGI) works, which is a little-known but vital part of the Internet.

Part 6, "Enjoying Entertainment and Multimedia on the Internet," shows you how some of the most exciting parts of the Internet work—the various multimedia technologies. Whether you want to know how virtual reality or animations work, how streaming video works, how videoconferencing works, or how similar technologies work, you'll find it all here.

Part 7, "Shopping and Doing Business on the Internet," covers intranets, how the Internet works with the outside world, and how you can shop online. You'll see how companies use Internet technologies to build their own private networks, called *intranets*. And you'll take a close look at the underlying technologies that let you shop on the Web, which accounts for billions of dollars a year in sales.

Finally, Part 8, "Protecting Yourself on the Internet," covers security concerns. It explains the controversial cookie technology that lets Web servers put bits of information on your hard disk and use that information to track you. It explains how the even-more controversial FBI system called Carnivore can in essence wiretap people's use of the Internet to read their e-mail, watch their Web browsing, and more. This part also shows how firewalls work, how viruses can attack your computer, and how cryptosystems allow confidential information to be sent across the Internet. It delves into how hackers can attack Internet service providers (ISPs) using so-called "smurf attacks," and how they can attack your computer as well. And it covers the issue of pornography on the Internet and shows how parental-control software can prevent children from seeing objectionable material.

So, come along and see how the vast Internet works. Even if you're a cyberpro (and especially if you're not), you'll find out a lot you never knew.



P A R T



UNDERSTANDING THE INTERNET'S UNDERLYING ARCHITECTURE

Chapter 1: What Is the Internet?
4

Chapter 2: How Computer Networks Send Data Across the Internet
8

Chapter 3: How TCP/IP Works
12

Chapter 4: Understanding the Internet's Software Structure
16

Chapter 5: How Internet Addresses and Domains Work
20

Chapter 6: How Routers Work
28

Chapter 7: How Internet File Types Work
32

FOR the first time ever, the world is at your fingertips. From your computer, you can find information about anything you can name or even imagine. You can communicate with people on the other side of the world. You can set up a teleconference, tap into the resources of powerful computers anywhere on the globe, search through the world's best libraries, and visit the world's most amazing museums. You can watch videos, listen to music, and read special multimedia magazines. You can shop for almost anything you can name. You can do all this by tapping into the largest computer network in the world—the Internet.

The Internet isn't a single network; it is a vast, globe-spanning network of networks. No single person, group, or organization runs the Internet. Instead, it's the purest form of electronic democracy. The networks communicate with one another based on certain protocols, such as the Transmission Control Protocol (TCP) and the Internet Protocol (IP). More and more networks and computers are being hooked up to the Internet every day. Tens of thousands of these networks exist, ranging from university networks to corporate local area networks to large online services such as America Online and CompuServe. Every time you tap into the Internet, your own computer becomes an extension of that network.

The first section of this book is spent defining the Internet. It also examines the architectures, protocols, and general concepts that make it all possible.

Chapter 1, "What Is the Internet?," examines how the Internet runs. You'll look at who pays for the high-speed data backbones that carry much of the Internet's traffic and at the organizations that ensure that standards are set for networks to follow so the Internet can run smoothly. You'll also look at the various types of networks connected to the Internet.

Chapter 2, "How Computer Networks Send Data Across the Internet," explains how information travels across the Internet and describes how hardware such as routers, repeaters, and bridges sends information among networks. It also shows how smaller networks are grouped into larger regional networks—and how those large regional networks communicate among themselves.

Chapter 3, "How TCP/IP Works," covers the Internet's basic protocols for communications. You'll learn a little about basic Internet jargon, such as TCP/IP (short for Transmission Control Protocol and Internet Protocol). This chapter also explains how those protocols work and how special software such as Winsock enables personal computers to get onto a network originally designed for larger computers.

Chapter 4, "Understanding the Internet's Software Structure," looks at the Internet's client/server architecture. *Servers*—also called *hosts*—are powerful computers that perform functions such as delivering information or Web pages, hosting databases, and handling e-mail. A *client* is your own computer and the software that sits on it, such as a Web browser

or piece of e-mail software. Clients request information from servers, which do the heavy-duty processing and then send the information back to the client, which displays the information.

Chapter 5, "How Internet Addresses and Domains Work," takes the mystery out of the Internet's often confusing addressing scheme. You'll learn about Internet domains and addresses and will even be able to make sense of them. More than that, you'll learn how domain servers are capable of keeping track of all the locations on the Internet and translating addresses such as `www.zdnet.com` into Internet IP addresses such as `134.54.56.120`. You'll also learn how some computers are assigned new IP addresses by special servers every time they connect to the Internet.

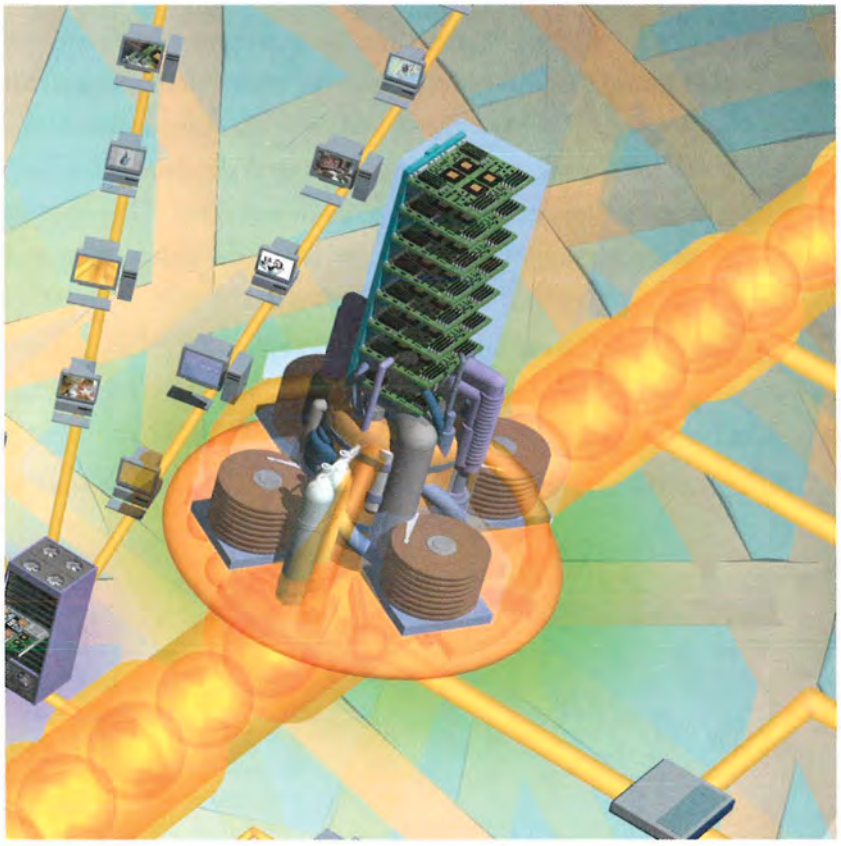
Chapter 6, "How Routers Work," details how the most basic piece of hardware on the Internet—a router—works. *Routers* are combinations of hardware and software that perform the job of ensuring that all data is sent to the proper destination. Think of routers as traffic cops of the Internet. They use the IP addresses the name servers have translated to route the data. Routers look at the addresses and then send the data to the next-closest router to the destination and so on, until the data is finally delivered. They use *routing tables* to determine how to route the traffic, and they can also adjust the routes as the traffic on the Internet changes, thus ensuring that the data is routed in the most efficient way possible.

Chapter 7, "How Internet File Types Work," gives you an understanding of the most common types of files you'll come across when you browse the Net. Compressed files, video files, graphics files—you'll learn about almost every type of file you might encounter.

Whether you're a newbie or cyberpro, this section teaches you the basics of the Internet.

CHAPTER 1

What Is the Internet?



ONE of the most frequently asked questions about the Internet is: Who runs it? The truth is that no centralized management of the Internet exists. Instead, it is a collection of thousands of individual networks and organizations, each of which is run and paid for on its own. Each network cooperates with other networks to direct Internet traffic so that information can pass among them. Together, these networks and organizations make up the wired world of the Internet. For networks and computers to cooperate in this way, however, a general agreement must take place about things such as Internet procedures and standards for protocols. These procedures and standards are laid out in requests for comment (RFCs) that are agreed on by Internet users and organizations.

A variety of groups guide the Internet's growth by helping to establish standards and by educating people on the proper way to use the Internet. Perhaps the most important is the Internet Society, a private, nonprofit group. The Internet Society supports the work of the Internet Activities Board (IAB), which handles much of the Internet's behind-the-scenes and architectural issues. The Internet Engineering Task Force (IETF) is responsible for overseeing how the Internet's TCP/IP protocols evolve. For information about the IETF, go to `www.ietf.org`.(See Chapter 3, "How TCP/IP Works," for details on TCP/IP protocols.)

The World Wide Web Consortium (W3C) develops standards for the evolution of the fastest-growing part of the Internet, the World Wide Web (find it at `www.w3.org`). The W3C is an industry consortium run by the Laboratory for Computer Science at the Massachusetts Institute of Technology (MIT).

Private companies oversee the registering of Internet domains, such as `www.zdnet.com` or `www.quepublishing.com`. These companies, called *registrars*, all must cooperate with one another to ensure that only one person or company can own a particular domain and that all the domains work properly. The registrars compete with one another as well in allowing people and businesses to register domains. Registering a domain costs money, and the registrars compete on cost and on giving extra services to those who buy domains.

Although all these types of organizations are important for holding together the Internet, at the heart of the Internet are individual local networks. These networks can be found in private companies, universities, government agencies, and online services. They are funded separately from each other and in a variety of manners, such as fees from users, corporate support, taxes, and grants. Many Internet service providers (ISPs), which provide Internet access for individuals, have networks as well. Individuals who want to access the Internet pay ISPs a monthly connection rate, so in that sense, everyone who uses the Internet helps pay for it.

The networks are connected in a variety of ways. For efficiency's sake, local networks join in consortiums known as *regional networks*. A variety of leased lines connect regional and local networks. The leased lines that connect networks can be as simple as a single telephone line or as complex as a fiber-optic cable with microwave links and satellite transmissions.

Private companies who make money by selling access to their lines build *backbones*, which are very high-capacity lines that carry enormous amounts of Internet traffic. Government agencies, such as NASA, and large private corporations pay for some of these backbones. The National Science Foundation also pays for some backbones.

How the Internet Comes Together

Organization of networks Because the Internet is a loose organization of networks, no single group runs it or pays for it all. Instead, many private organizations, universities, and government agencies pay for and run parts of it. They all work together in a democratic, loosely organized alliance. Private organizations range from small, homegrown networks to commercial online services, such as America Online and CompuServe, and private Internet service providers (ISPs) that sell access to the Internet.



Regional Network

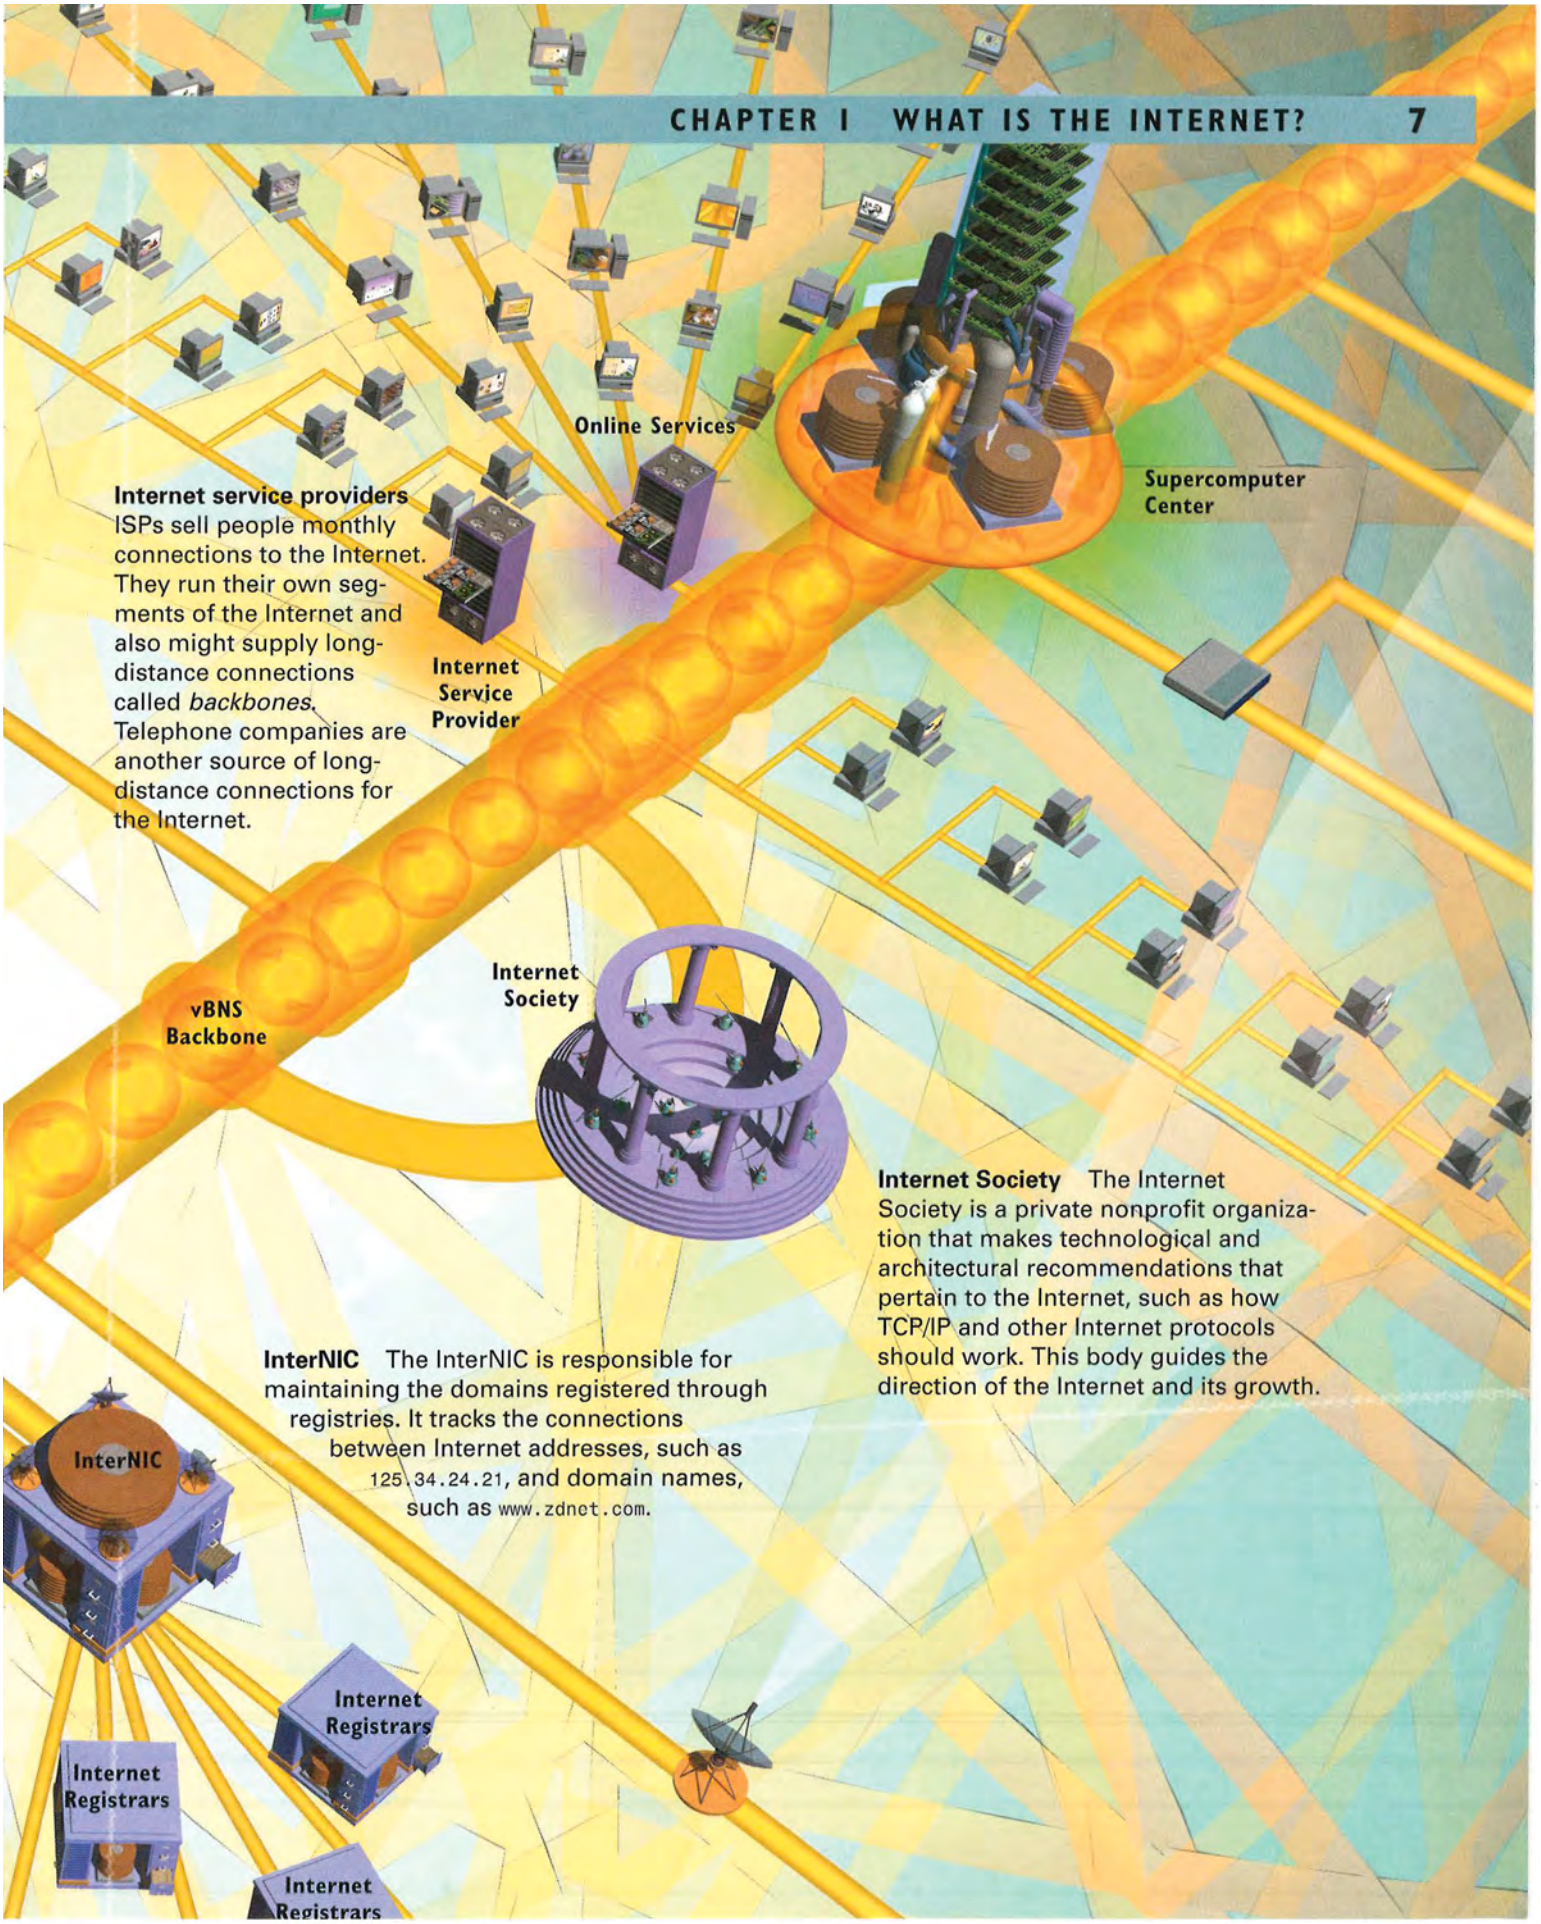
Funding the Internet Through agencies such as the National Science Foundation, the federal government pays for some high-speed backbones that carry Internet traffic across the country and the world. The high-speed vBNS (very high-speed Backbone Network Services), for example, provides a high-speed infrastructure for the research and education community by linking together supercomputer centers. Often, a large corporation or organization such as NASA provides backbones to link sites across the country or the world. The government has also funded the Internet2, a super-fast network that can transfer data at an astounding 2.4 gigabits per second, for use by universities. A number of universities already are connected to it. When finished in 2003, it will link 140 universities. Eventually, the entire Internet might run at its speed.

Regional networks Regional networks provide and maintain Internet access within a geographic area. Regional nets can consist of smaller networks and organizations within the area that have banded together to provide better service.

Registrars Private companies called Internet *registrars* are responsible for registering Internet domains, such as `www.zdnet.com`, to people and businesses. Until recently, a quasi-public company called the InterNIC had sole responsibility for doing this, but other registrars can now register domains as well.



Supercomputer Center



Internet service providers ISPs sell people monthly connections to the Internet. They run their own segments of the Internet and also might supply long-distance connections called *backbones*. Telephone companies are another source of long-distance connections for the Internet.

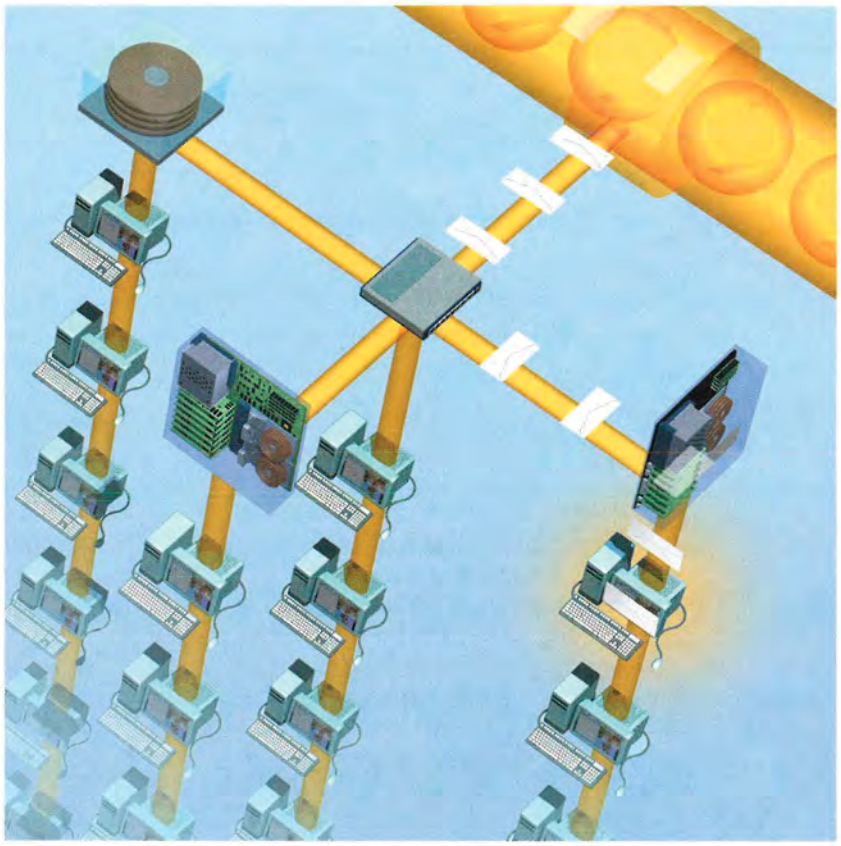
Internet Society The Internet Society is a private nonprofit organization that makes technological and architectural recommendations that pertain to the Internet, such as how TCP/IP and other Internet protocols should work. This body guides the direction of the Internet and its growth.

InterNIC The InterNIC is responsible for maintaining the domains registered through registries. It tracks the connections between Internet addresses, such as `125.34.24.21`, and domain names, such as `www.zdnet.com`.

CHAPTER

2

How Computer Networks Send Data Across the Internet



YOU might take for granted that when you send a piece of information across the Internet, it will always reach its intended destination. However, the process of sending that information is remarkably complex.

When you send information across the Internet, the *Transmission Control Protocol (TCP)*—the language computers use when communicating over the Internet—first breaks the information up into *packets*, smaller blocks of information that also contain a variety of data that helps the packets travel across the Internet. Your computer sends those packets to your local network, Internet service provider (ISP), or online service. From there, the packets travel through many levels of networks, computers, and communications lines before they reach their final destinations, which might be across town or around the world. A variety of hardware processes those packets and routes them to their proper destinations. This hardware is designed to transmit data between networks and makes up much of the glue that holds the Internet together. Five of the most important pieces of hardware are hubs, bridges, gateways, repeaters, and routers.

Hubs are important because they link groups of computers to one another and let computers communicate with each other. *Bridges* link local area networks (LANs) with one another. They enable data destined for another LAN to be sent there, while keeping local data inside its own network. *Gateways* are similar to bridges, but they also translate data from one type of network to another.

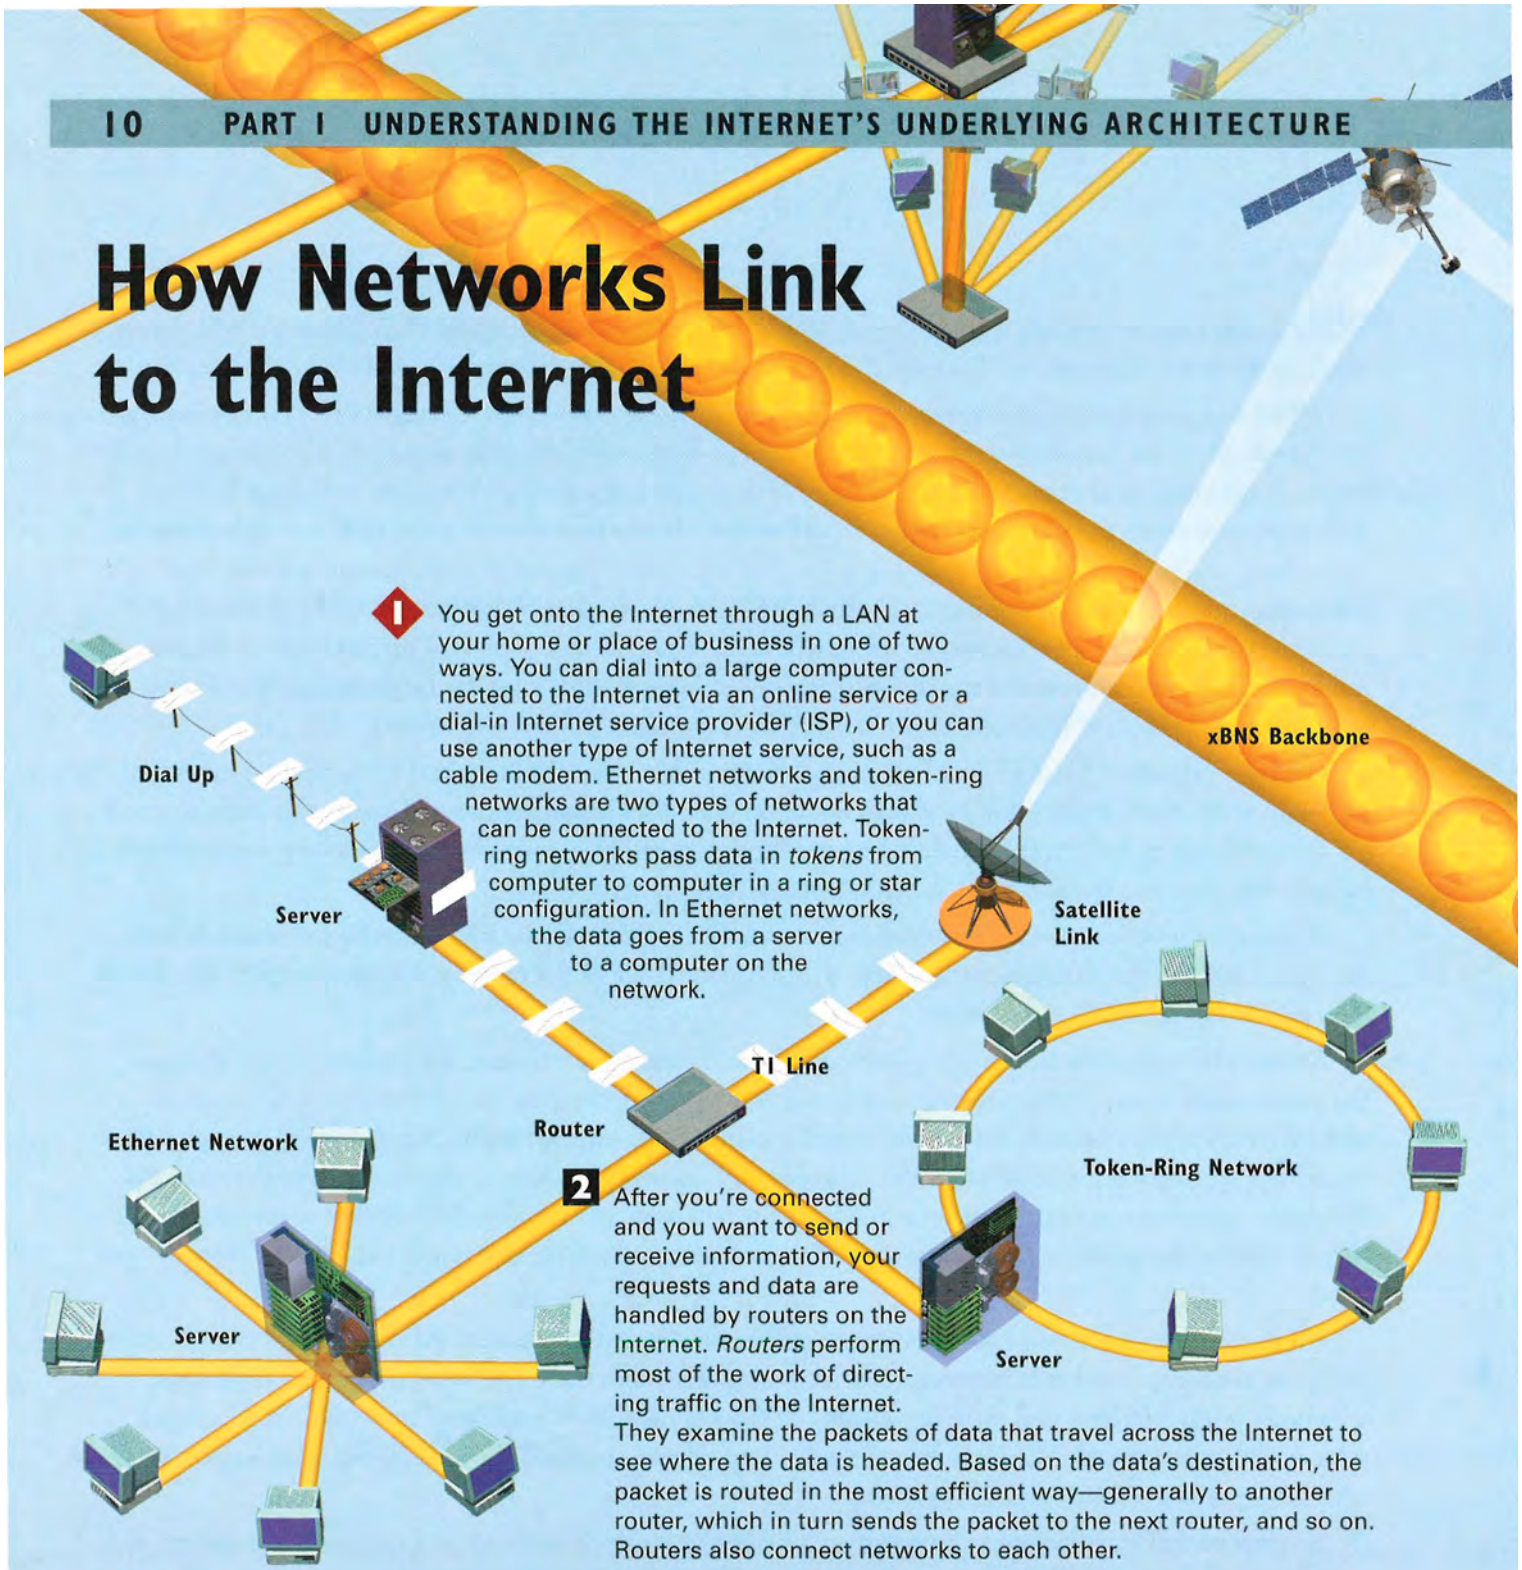
When data travels across the Internet, it often crosses great distances, which can be a problem because the signal sending the data can weaken over the distance. To solve the problem, *repeaters* amplify the data at intervals so the signal doesn't weaken.

Routers play a key role in managing Internet traffic. Their job is to ensure the packets always arrive at the proper destination. If data is being transferred among computers that are on the same LAN, routers often aren't necessary because the network itself can handle its internal traffic. Routers come into play when the data is sent between two different networks. Routers examine packets to determine their destinations. They take into account the volume of activity on the Internet, and they send the packet to another router that is closer to the packet's final destination. For more information on routers see Chapter 6, "How Routers Work."

All this hardware connects the many networks that make up the Internet. At the most local level of networks are corporate local area networks. Midlevel networks hook these LANs together using high-speed telephone lines, ethernet, and microwave links. A *regional* network is a midlevel network in a geographic area. A *wide area network (WAN)* is another type of midlevel network. A WAN consists of an organization with many networked sites linked together.

When a packet travels from a computer on a LAN in a midlevel network to a computer somewhere else on the midlevel network, a router (or a series of routers) sends the packet to its proper destination. However, if the destination lies outside the midlevel network, the packet is sent to a network access point (NAP), where it is sent across the country or the world on a backbone. High-speed backbones such as the vBNS (very high-speed Backbone Network Services) can transmit enormous amounts of data—155 megabits (millions of bits) per second (Mbps). Even faster backbones are being built that will transmit data at an astonishing 9.6 billion bits per second.

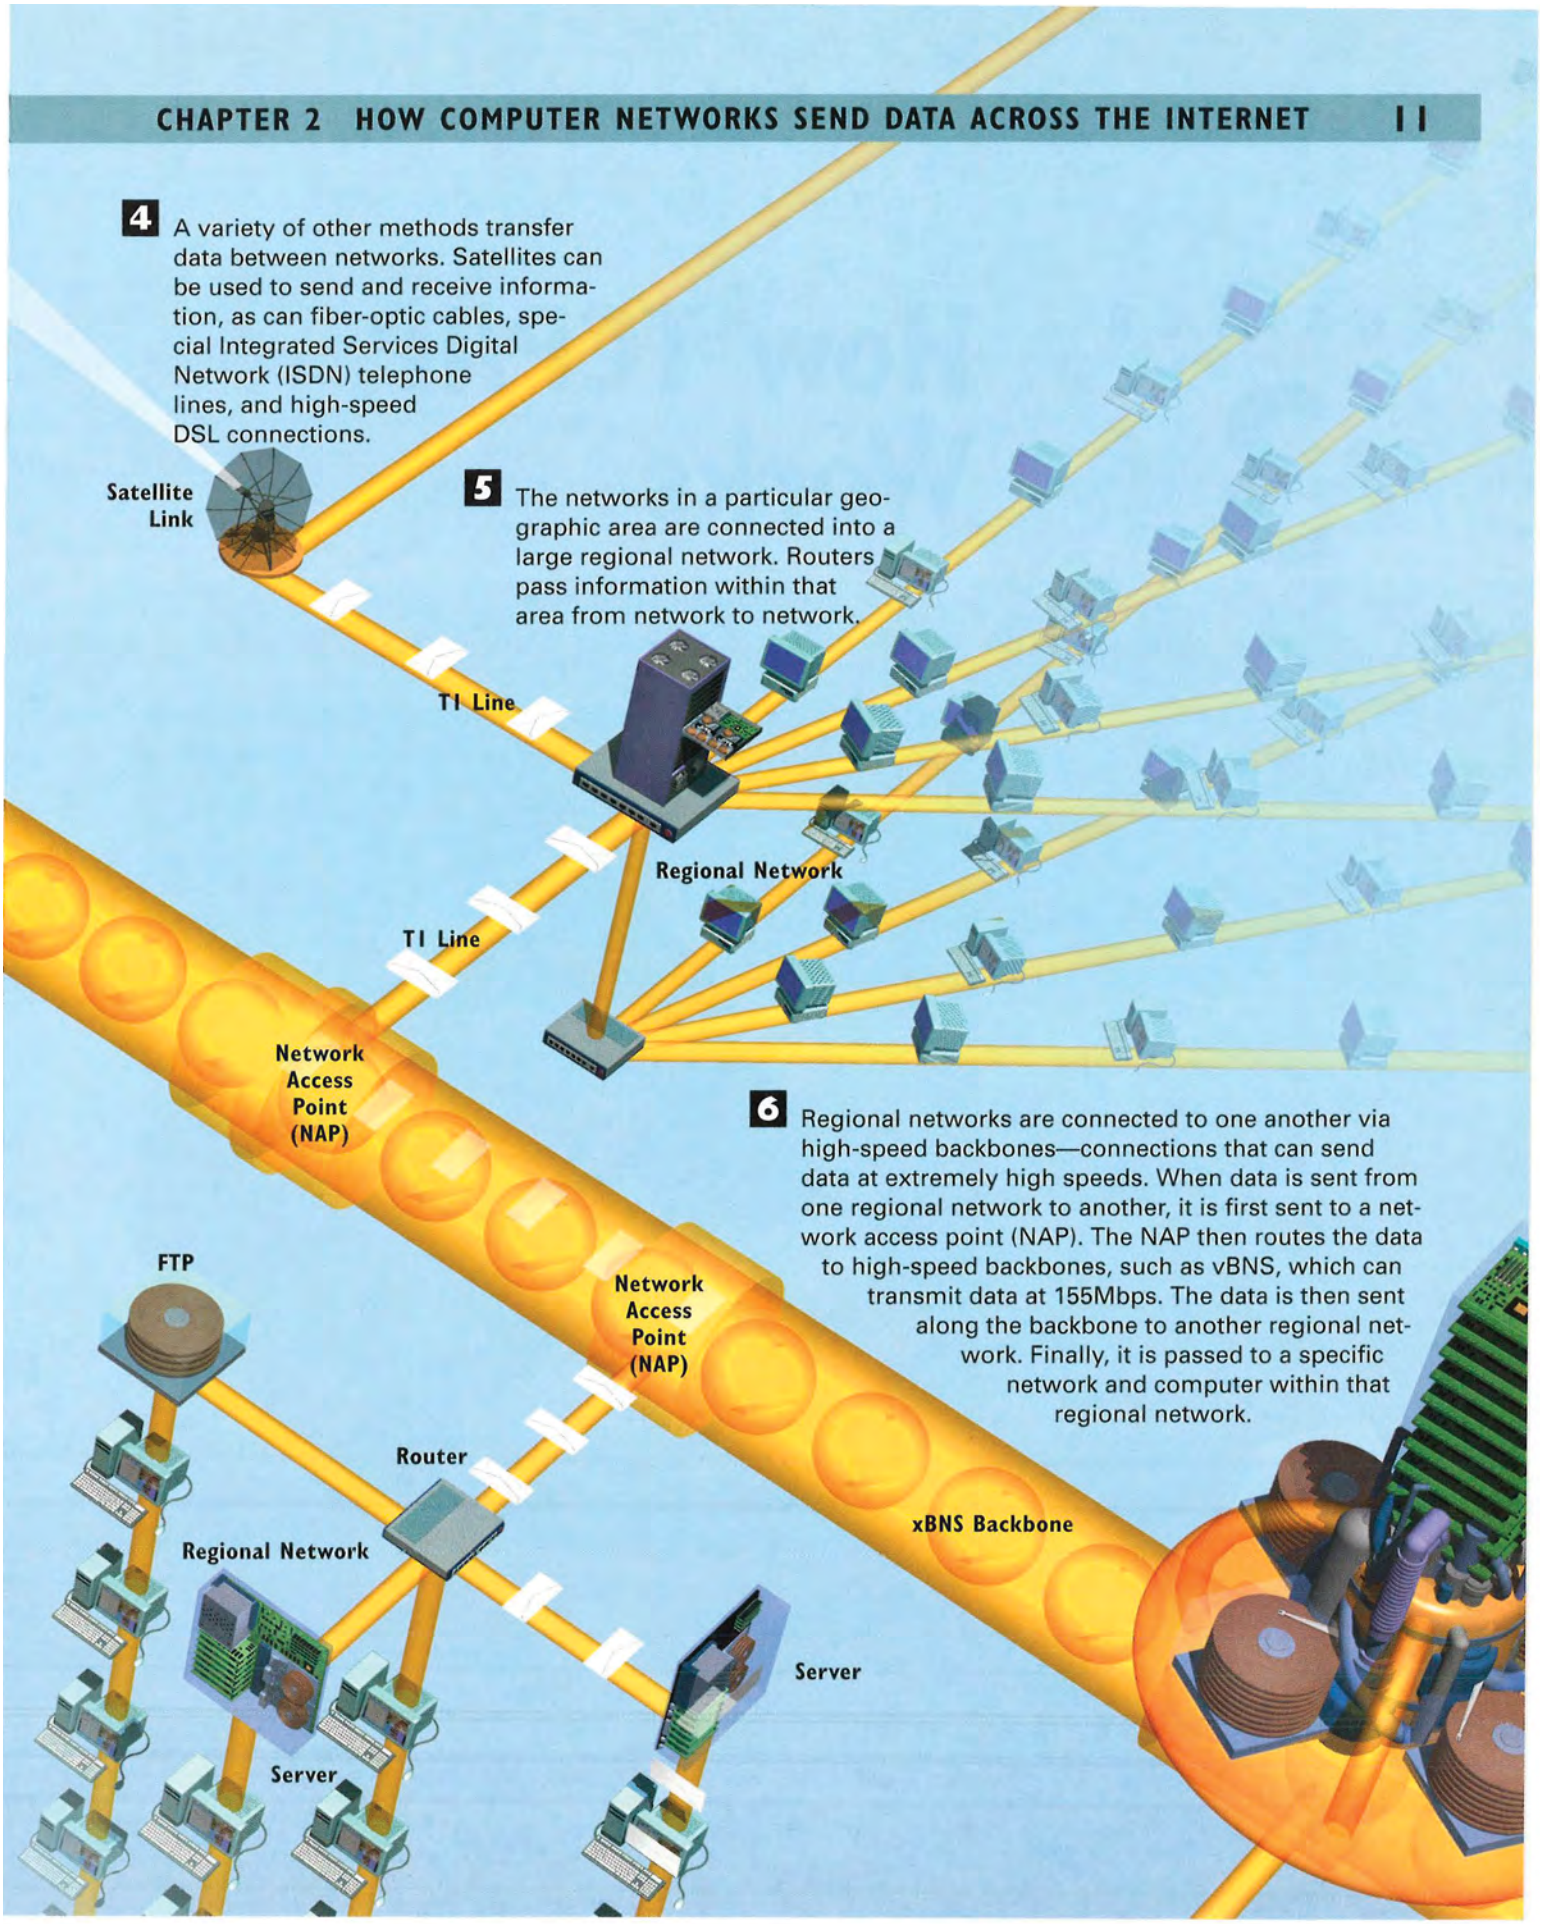
How Networks Link to the Internet



1 You get onto the Internet through a LAN at your home or place of business in one of two ways. You can dial into a large computer connected to the Internet via an online service or a dial-in Internet service provider (ISP), or you can use another type of Internet service, such as a cable modem. Ethernet networks and token-ring networks are two types of networks that can be connected to the Internet. Token-ring networks pass data in *tokens* from computer to computer in a ring or star configuration. In Ethernet networks, the data goes from a server to a computer on the network.

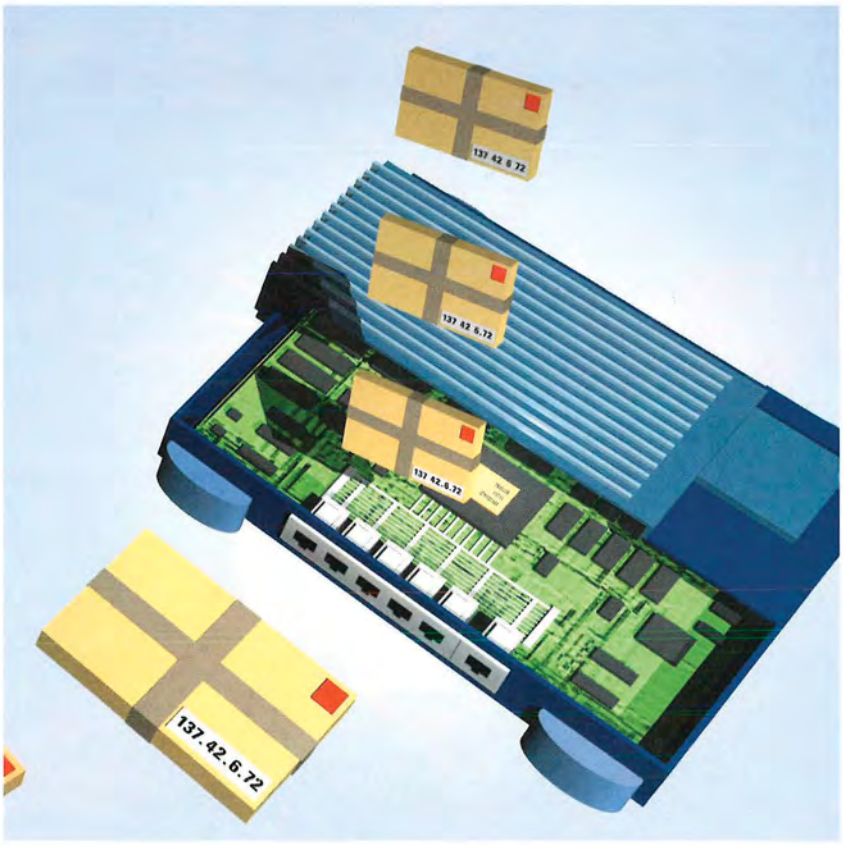
2 After you're connected and you want to send or receive information, your requests and data are handled by routers on the Internet. *Routers* perform most of the work of directing traffic on the Internet. They examine the packets of data that travel across the Internet to see where the data is headed. Based on the data's destination, the packet is routed in the most efficient way—generally to another router, which in turn sends the packet to the next router, and so on. Routers also connect networks to each other.

3 The data can be transferred between networks in a number of ways. Dedicated telephone lines can transmit data at 56Kbps (kilobits per second). An increasing number of T1 leased telephone lines carry data between networks. A T1 link can carry data at 1.544Mbps. Higher-speed T3 links, which can carry data at 44.746Mbps, are being used as well. If you dial into an ISP from home, you might connect at a lower speed than at your office, where you might have higher-speed connections such as a T1 or 56Kbps line. However, a variety of high-speed options are available for connecting to the Internet at home, such as cable modems or high-speed Digital Subscriber Line (DSL) connections.


4 A variety of other methods transfer data between networks. Satellites can be used to send and receive information, as can fiber-optic cables, special Integrated Services Digital Network (ISDN) telephone lines, and high-speed DSL connections.
Satellite Link
5 The networks in a particular geographic area are connected into a large regional network. Routers pass information within that area from network to network.
T1 Line
Regional Network
T1 Line
Network Access Point (NAP)
6 Regional networks are connected to one another via high-speed backbones—connections that can send data at extremely high speeds. When data is sent from one regional network to another, it is first sent to a network access point (NAP). The NAP then routes the data to high-speed backbones, such as vBNS, which can transmit data at 155Mbps. The data is then sent along the backbone to another regional network. Finally, it is passed to a specific network and computer within that regional network.
FTP
Network Access Point (NAP)
Router
Regional Network
xBNS Backbone
Server
Server

CHAPTER 3

How TCP/IP Works



A seemingly simple set of ideas makes it possible for computers and networks all over the world to share information and messages on the Internet: Break up every piece of information and message into pieces called *packets*, deliver those packets to the proper destinations, and then reassemble the packets into their original form after they've been delivered so the receiving computer can view and use them. That's the job of the two most important communications protocols on the Internet—the Transmission Control Protocol (TCP) and the Internet Protocol (IP). They are frequently referred to as *TCP/IP*. TCP breaks down and reassembles the packets, whereas IP is responsible for ensuring the packets are sent to the right destination.

TCP/IP is used because the Internet is what is known as a packet-switched network. In a *packet-switched network*, there is no single, unbroken connection between sender and receiver. Instead, when information is sent, it is broken into small packets, sent over many different routes at the same time, and then reassembled at the receiving end. By contrast, the telephone system is a circuit-switched network. In a *circuit-switched network*, after a connection is made (as with a telephone call, for example), that part of the network is dedicated only to that single connection.

For personal computers to take full advantage of the Internet, they need to use special software that understands and interprets the Internet's TCP/IP protocols. This software is referred to as a *socket* or a *TCP/IP stack*. For PCs, the required software is called Winsock. Many versions of Winsock are available for PCs. For Macintoshes, the software is called MacTCP. In both cases, this software serves as an intermediary between the Internet and the personal computer. Personal computers can take advantage of the simplest, most rudimentary parts of the Internet without using Winsock or MacTCP—for example, you can't view Web pages without it or listen to music over the Internet, among many other things. For full access to the Internet, TCP/IP stacks are necessary.

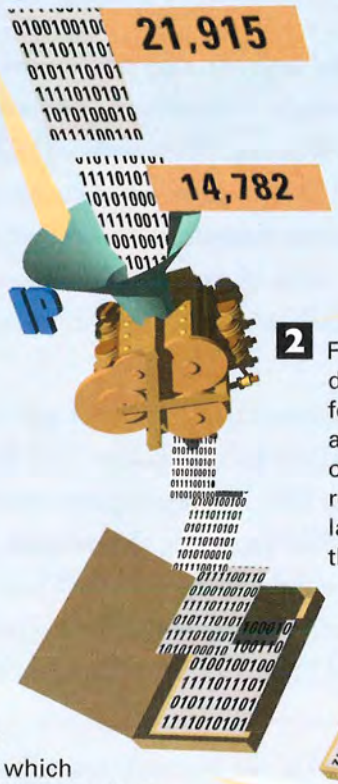
There are two primary ways that a computer can connect to the Internet and then use these protocols—through a direct connection via a local area network (LAN), cable modem, or DSL line; or by dialing in using a modem. To connect via LAN, cable modem, or DSL line, a computer needs a network card. To communicate with the network and the Internet's TCP/IP protocols, the network card requires a hardware driver—software that mediates between the network and the network card. When the computer instead dials in to the Internet by using a modem, the computer must use one of two software protocols: either Serial Line Internet Protocol (SLIP) or Point-to-Point Protocol (PPP). These protocols do the job of communicating with the Internet's TCP/IP protocols.

SLIP and PPP are designed for computers connected to the Internet over a serial connection via a modem. In general, the newer PPP provides a more error-free connection than does the older SLIP. Computers can also dial in to the Internet without using TCP/IP stacks, SLIP, or PPP. However, they won't be capable of tapping into the full power of the Internet, most specifically, the World Wide Web. To make full use of the Web, you need a computer that is part of the Internet, which requires a TCP/IP stack. The SLIP or PPP must then be used when dialing in to an ISP for Web access.

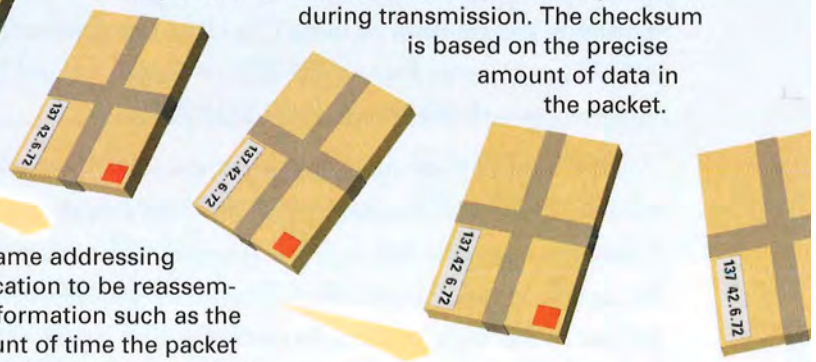
How the Internet's Basic TCP/IP Protocols Work



1 The Internet is a packet-switched network, which means that when you send information across the Internet from your computer to another computer, the data is broken into small packets. A series of switches called *routers* send each packet across the Net individually. After all the packets arrive at the receiving computer, they are recombined into their original, unified form. Two protocols do the work of breaking the data into packets, routing the packets across the Internet, and then recombining them on the other end: The Internet Protocol (IP), which routes the data, and the Transmission Control Protocol (TCP), which breaks the data into packets and recombines them on the computer that receives the information.



2 For a number of reasons, including hardware limitations, data sent across the Internet must be broken into packets of fewer than about 1,500 characters each. Each packet is given a header that contains a variety of information, such as the order in which the packets should be assembled with other related packets. As TCP creates each packet, it also calculates and adds to the header a *checksum*, which is a number that TCP uses on the receiving end to determine whether any errors have been introduced into the packet during transmission. The checksum is based on the precise amount of data in the packet.



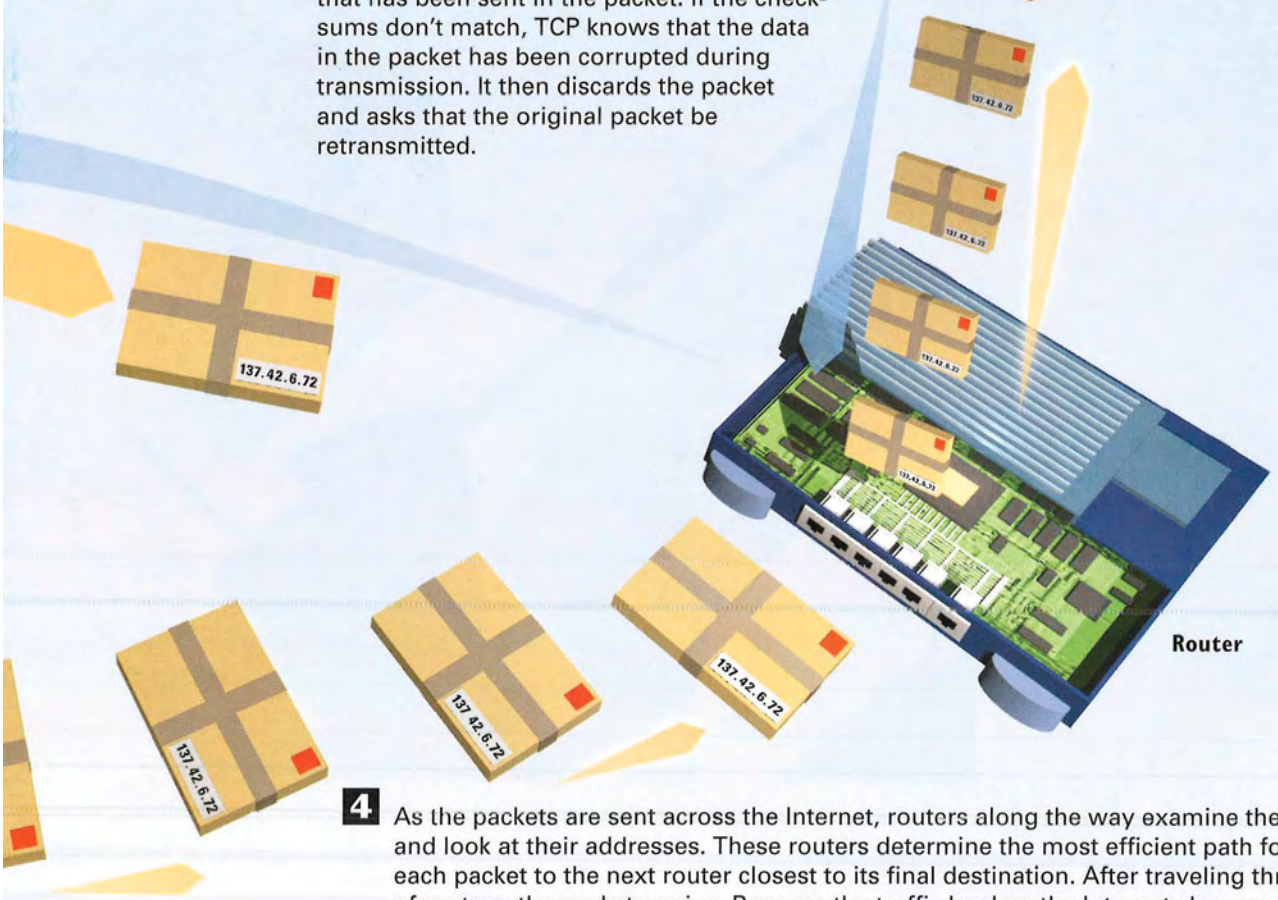
3 Each packet is put into separate IP *envelopes*, which contain addressing information that tells the Internet where to send the data. All the envelopes for a given piece of data have the same addressing information so they can all be sent to the same location to be reassembled. IP envelopes contain headers that include information such as the sender address, the destination address, the amount of time the packet should be kept before discarding it, and many other types of information.

6 When all the noncorrupt packets are received by the computer to which the information is being sent, TCP assembles them into their original, unified form.





5 As the packets arrive at their destinations, TCP calculates a checksum for each packet. It then compares this checksum with the checksum that has been sent in the packet. If the checksums don't match, TCP knows that the data in the packet has been corrupted during transmission. It then discards the packet and asks that the original packet be retransmitted.



4 As the packets are sent across the Internet, routers along the way examine the IP envelopes and look at their addresses. These routers determine the most efficient path for sending each packet to the next router closest to its final destination. After traveling through a series of routers, the packets arrive. Because the traffic load on the Internet changes constantly, the packets might be sent along different routes and the packets might arrive out of order.

CHAPTER

4

Understanding the Internet's Software Structure



THE Internet works on the client/server model of information delivery. In a client/server model, the client computer connects to a server computer on which information resides; the client depends on the server to deliver information. In effect, the client requests the services of the larger computer. These services can involve searching for information and sending it back to the client, such as when a database on the Web is queried. Other examples of these services are delivering Web pages and handling incoming and outgoing e-mail. Whenever you use the Internet, you're connected to a server computer and requesting the use of that server's resources.

Typically, the *client* is a local personal computer or the software that runs on it, and the *server* (also known as the *host*) is usually a more powerful computer that houses the data. Hosts and clients can be of many makes and manufacturers and run a wide variety of operating systems.

The connection to the server is made in many ways—via a LAN (local area network), via a phone line, via a cable modem or DSL modem, wirelessly, or any one of other countless ways. A primary reason to set up a client/server network is to allow many clients to access the same applications and files that are stored on a server. Also, as previously mentioned, these clients can be any one of many types of computers accessing the Internet in many ways.

In the case of the Internet's World Wide Web, the client is actually the browser on your PC and the server is a host computer located somewhere on the Internet. Typically, the browser sends the server a request for a specific Web page. The server processes that request and sends an answer back to the browser (again, most often in the form of a Web page).

The connection between the client and server is maintained only during the actual exchange of information. Thus, after a Web page is transferred from the host (or server) computer, the HTTP connection between that computer and the client is broken. (HTTP stands for Hypertext Transfer Protocol; it's the protocol used by the World Wide Web.) Even though the HTTP connection is closed, the ISP maintains the TCP/IP connection to the Internet.

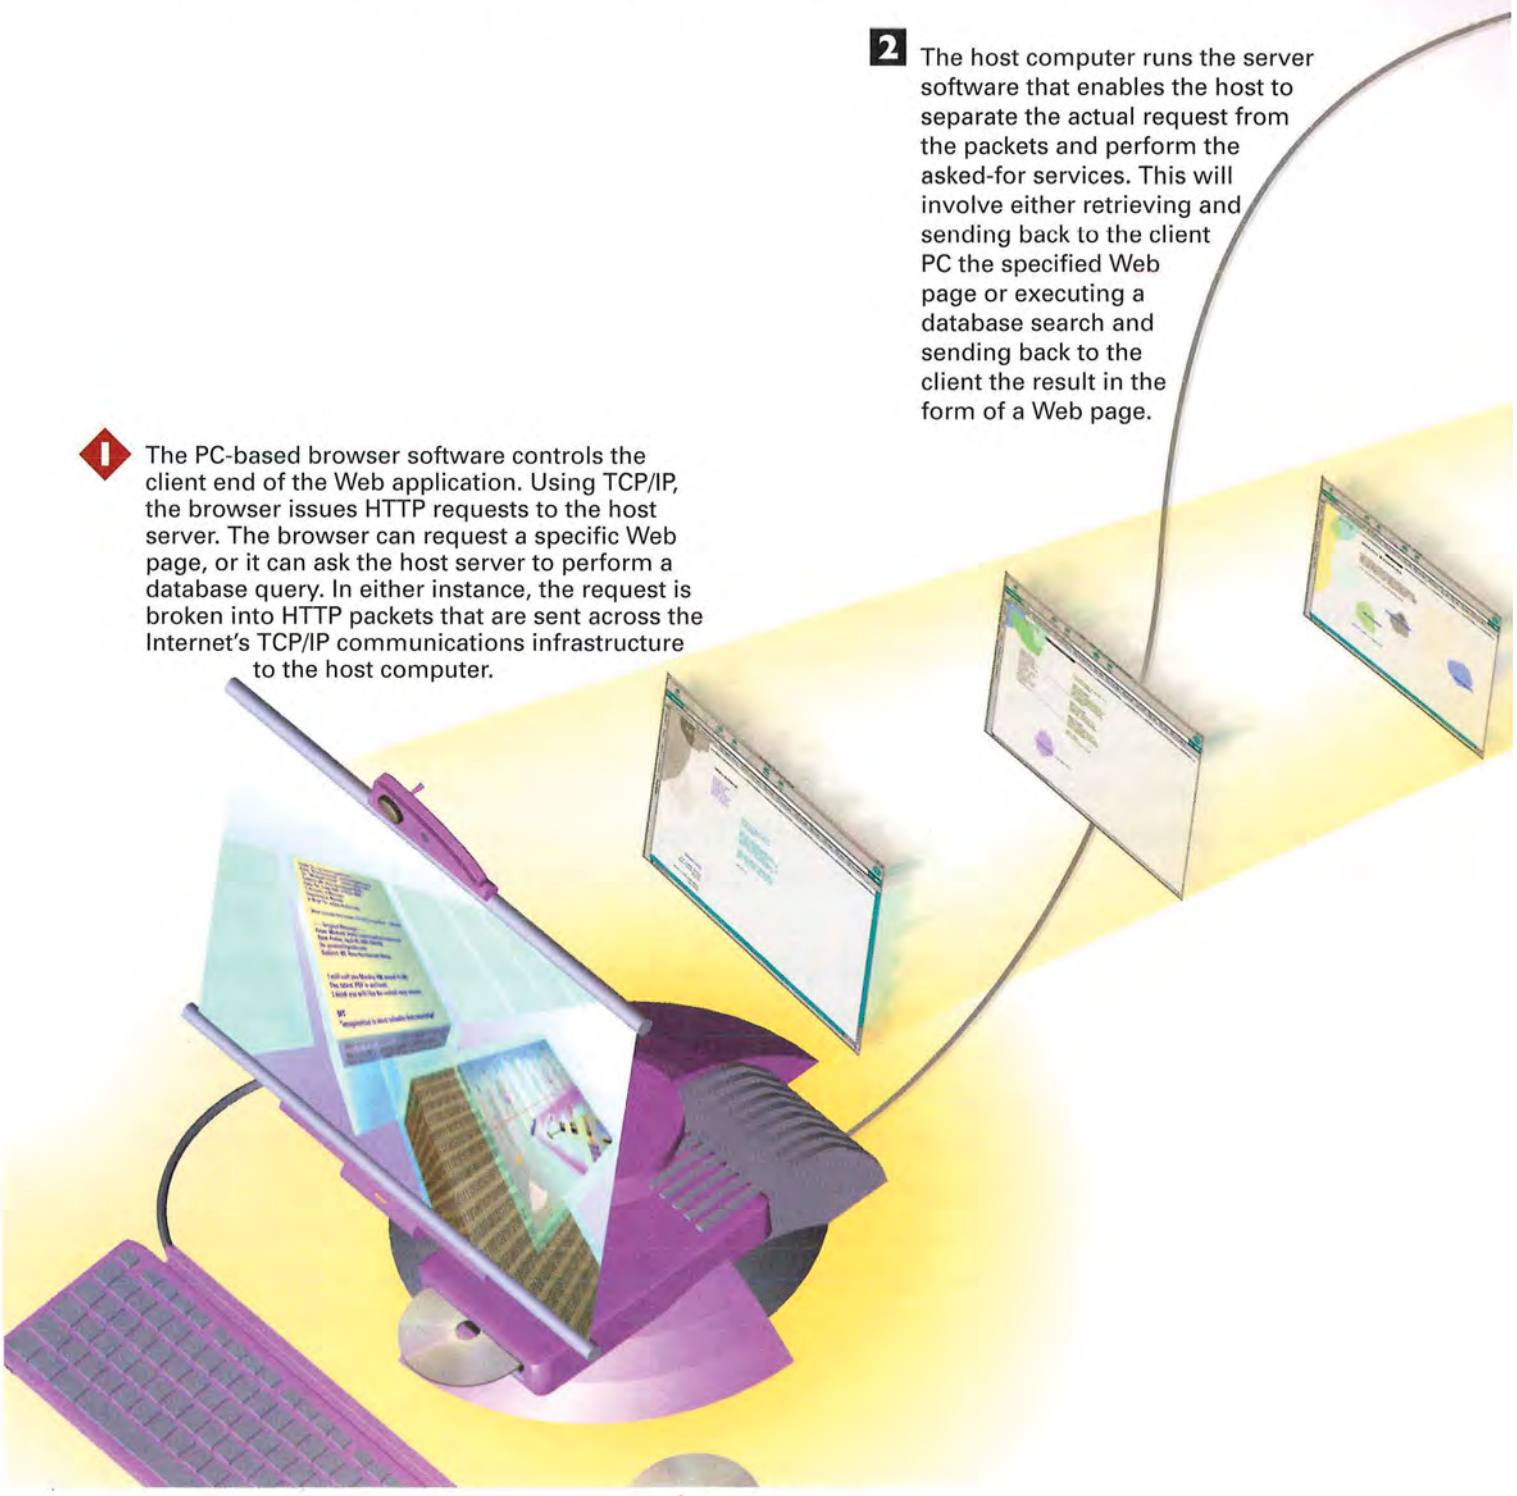
The client/server model enables the desktop PC to run the browser software to search the Web, yet still access host servers around the Internet to execute search and retrieval functions. Many of these functions are accomplished through technology called the Common Gateway Interface (CGI). For more information about how CGI works, see Chapter 35, "How CGI Scripting Works." In essence, this architecture enables the Web to be conceived of as a limitless file storage medium and database, distributed among thousands of host computers, all accessible by any individual PC.

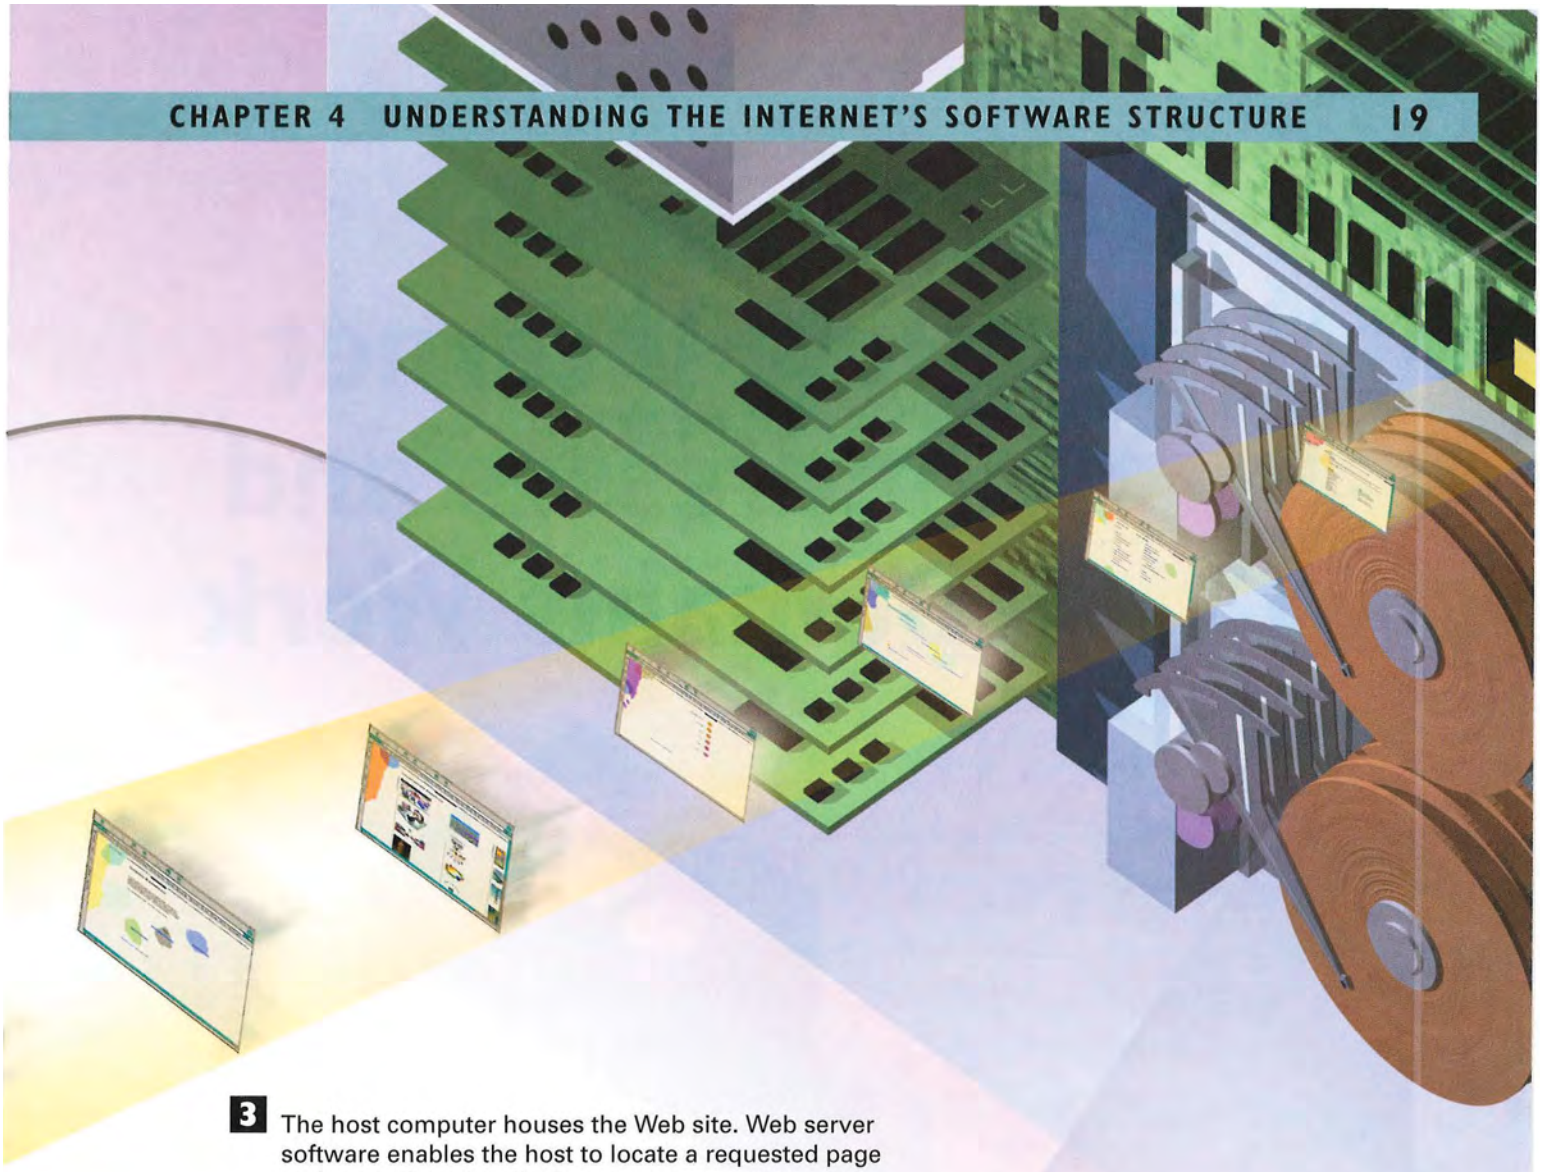
The following illustration shows how the Web runs on a client/server architecture. Keep in mind that all other resources on the Internet run on the client/server model as well. For example, in e-mail transactions the client would be the e-mail software on your computer, whereas the server would be the e-mail server to which you connect.

How Client/Server Architecture Works

2 The host computer runs the server software that enables the host to separate the actual request from the packets and perform the asked-for services. This will involve either retrieving and sending back to the client PC the specified Web page or executing a database search and sending back to the client the result in the form of a Web page.

1 The PC-based browser software controls the client end of the Web application. Using TCP/IP, the browser issues HTTP requests to the host server. The browser can request a specific Web page, or it can ask the host server to perform a database query. In either instance, the request is broken into HTTP packets that are sent across the Internet's TCP/IP communications infrastructure to the host computer.





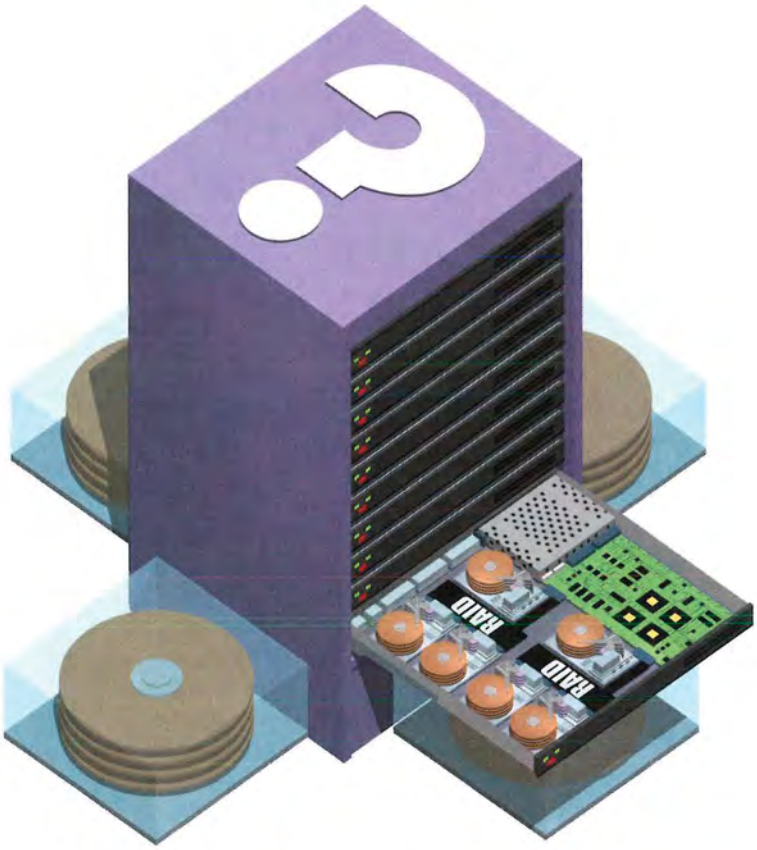
3 The host computer houses the Web site. Web server software enables the host to locate a requested page and return it to the client.

4 Database and other similar applications on the host computer are accessed and controlled through a common gateway interface. When the server receives the request for the database query, it forwards the request to a separate, dedicated database server or application for processing via the CGI.



CHAPTER

5 How Internet Addresses and Domains Work



THE heart of how the Internet works is the Domain Name System (DNS), the way in which computers can contact each other and do things such as exchange e-mail or display Web pages. When someone on the Internet wants to contact a location—for example, to visit a Web site—he types in an address, such as `www.metahouse.com`. (A specific location on the Internet is also referred to as a *uniform resource locator*, or *URL*.) The DNS translates the plain English address, `www.metahouse.com`, into series of numbers called an IP (for Internet Protocol) address. An IP address, such as `123.23.43.121`, marks the location of a computer on the Internet similar to the way a house number and street mark the location of where you live.

In this example, metahouse.com is called a *domain*. To work most efficiently, the Internet has been organized into a number of major domains. Major domains refer to the letters at the end of a plain English address, such as `.com`. A number of common domains are `.com` (commercial), `.edu` (education), `.gov` (government), `.mil` (military), `.net` (Internet service providers and networks—companies and groups concerned with the organization of the Internet), and `.org` (organization).

Domains are organized in a hierarchical manner, so that beneath major domains are many minor domains. As an example of how the DNS and domains work, look at NASA's SPACElink Internet address: `spacelink.nasa.gov`. The top domain is `.gov`, which stands for government. The domain just below that is `.nasa`, which is the NASA domain. Then, `spacelink` identifies the NASA computer that runs the SPACElink program. SPACElink's numeric IP address has changed through the years, but its Internet address has stayed the same.

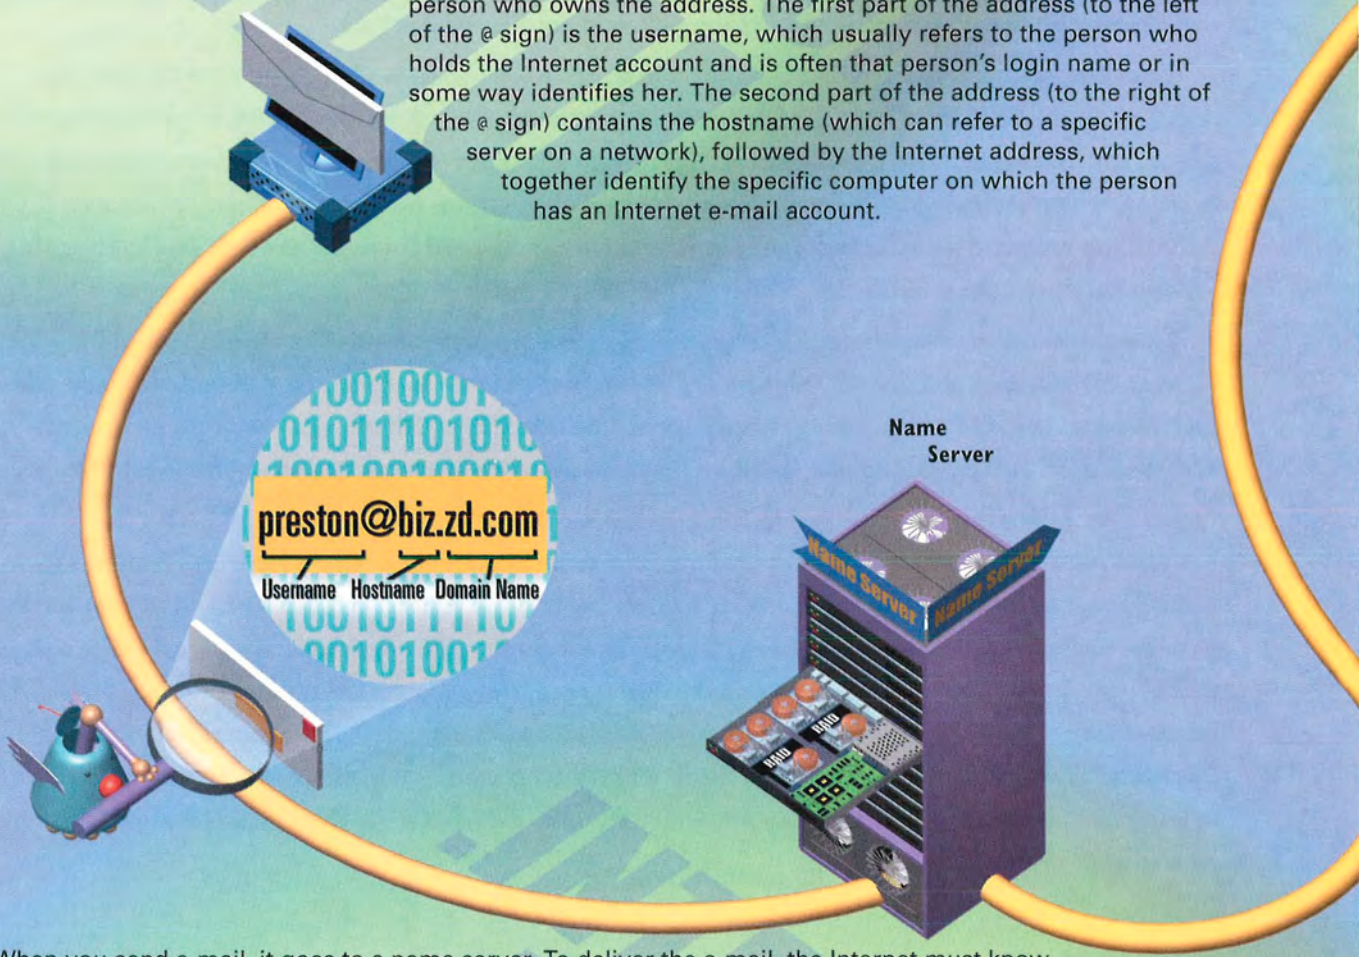
Computers called *name servers* are responsible for keeping track of such changes and translating them between IP addresses and domain addresses. The Internet can't understand alphanumeric Internet addresses, such as `pgralla@ziffdavis.com`, so name servers translate those addresses into their proper numeric IP addresses, such as `163.52.128.72`. Name servers contain tables that match alphanumeric Internet addresses to numeric IP addresses.

When you connect your computer to the Internet, your computer needs to have an IP address to do common things, such as browsing the Web. Depending on how your computer is set up and your service provider operates, you might have a static address or a dynamic address. A *static* IP address never changes, so if you have one, you will have the same IP address every time you connect to the Internet. However, because the Internet has only a limited number of IP addresses, many ISPs use dynamic addresses. With a *dynamic* IP address, you are given a different IP address every time you connect. In this way, ISPs don't have to have an individual address for every subscriber. Instead, they can share their pool of addresses among all subscribers.

Understanding Internet Addresses and Domains

1 The Internet Protocol (IP) delivers mail based on the specific address of the recipient's e-mail server. This address is expressed as four numbers, separated by periods (called dots), such as 163.52.128.72. However, because it would be difficult to remember such complex addresses, you can instead use Internet addresses made up of words, letters, and numbers. Computers called domain name servers translate the alphabetical address into a numerical address, so e-mail can be sent to the proper location.

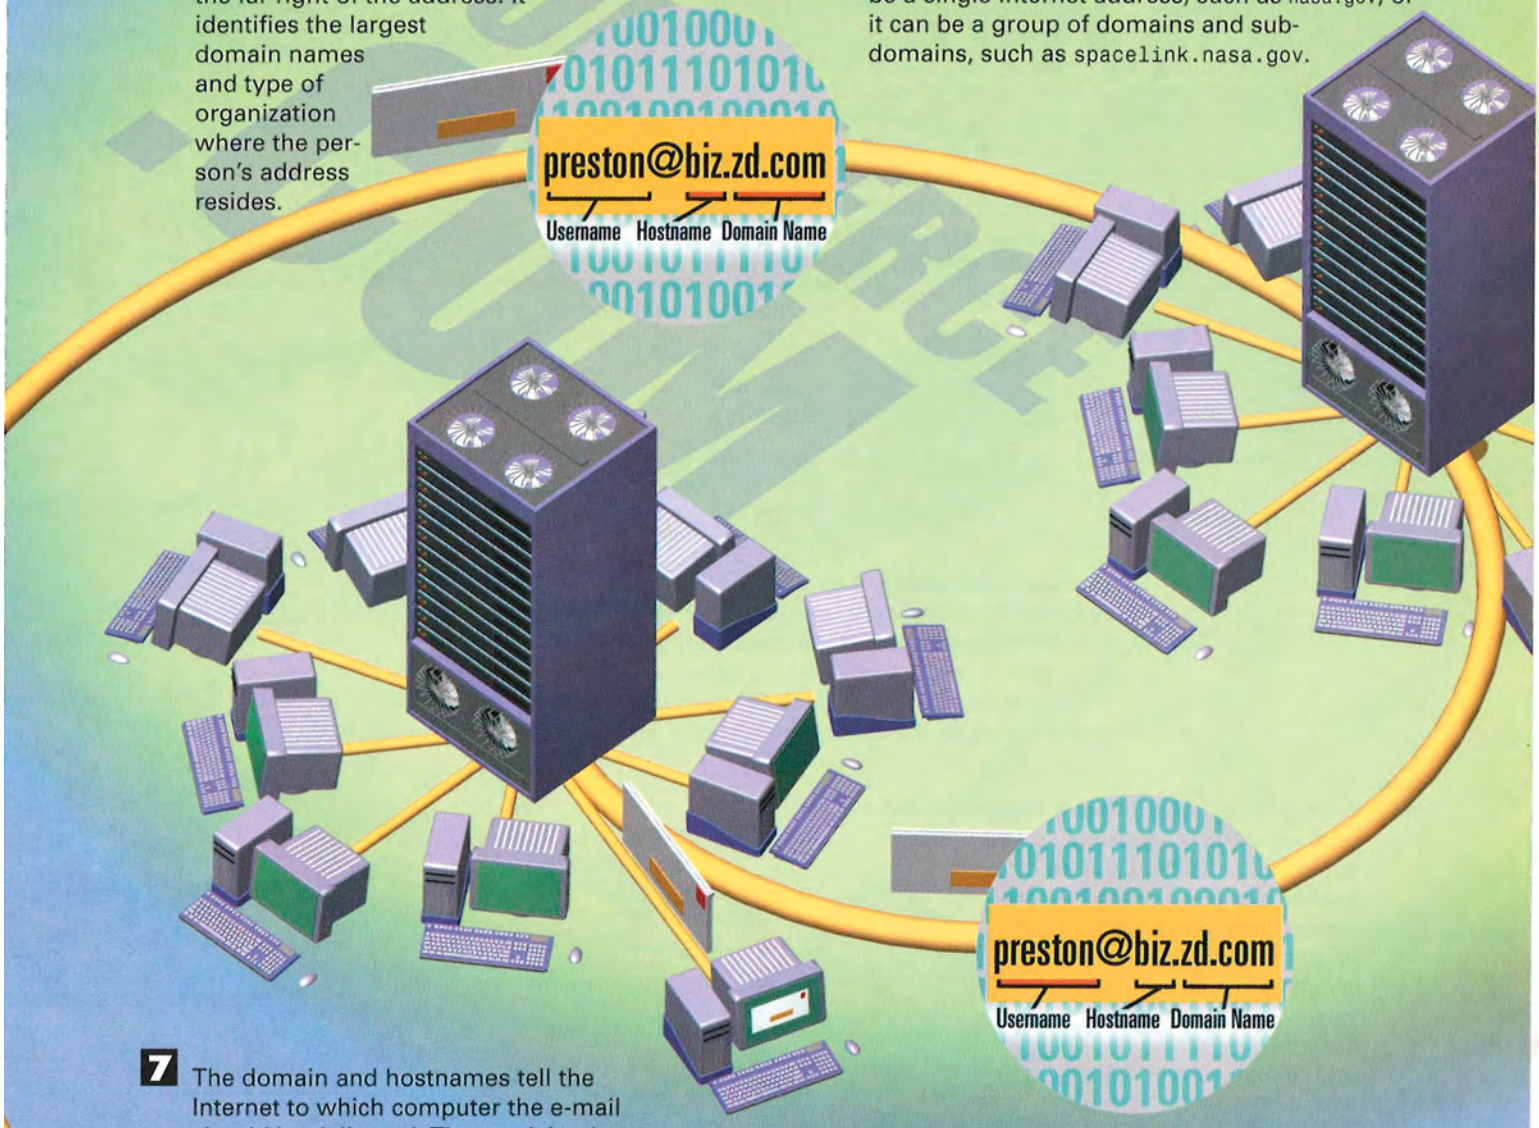
2 An Internet address is made up of two major parts separated by an @ (at) sign. The address can tell you a good deal of information about the person who owns the address. The first part of the address (to the left of the @ sign) is the username, which usually refers to the person who holds the Internet account and is often that person's login name or in some way identifies her. The second part of the address (to the right of the @ sign) contains the hostname (which can refer to a specific server on a network), followed by the Internet address, which together identify the specific computer on which the person has an Internet e-mail account.



3 When you send e-mail, it goes to a name server. To deliver the e-mail, the Internet must know the numeric IP address, such as 163.52.128.72. Name servers look up the alphabetical address and substitute the numeric IP address for it so the e-mail can be delivered properly.

4 The Domain Name System (DNS) divides the Internet into understandable groups, or domains. Note the portion of the domain section at the far right of the address. It identifies the largest domain names and type of organization where the person's address resides.

5 To the left of the largest domain is specific information about the organization, which tells routers to which network the e-mail should be sent. It can be a single Internet address, such as `nasa.gov`, or it can be a group of domains and subdomains, such as `spacelink.nasa.gov`.

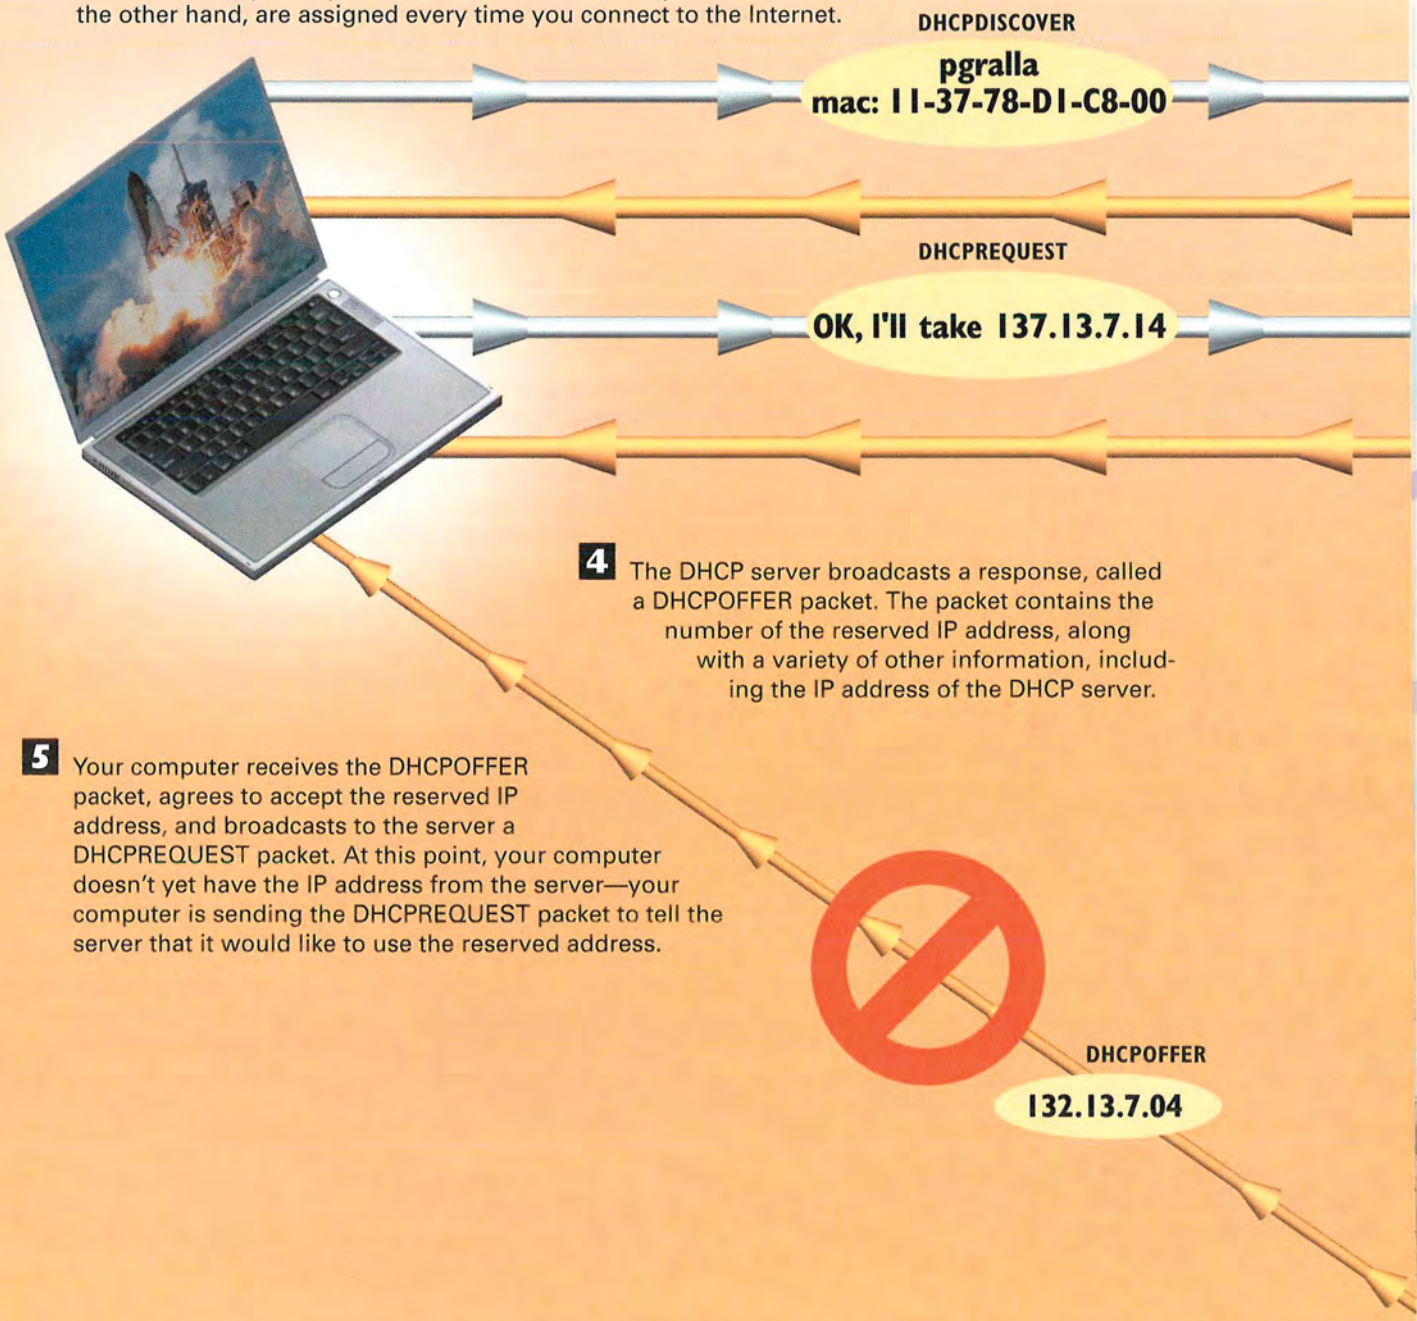


6 To the left of the Internet address is the hostname, which tells routers to which specific computer within the domain the e-mail should be delivered.

7 The domain and hostnames tell the Internet to which computer the e-mail should be delivered. The receiving host computer looks at the username and delivers the mail to the proper e-mail box.

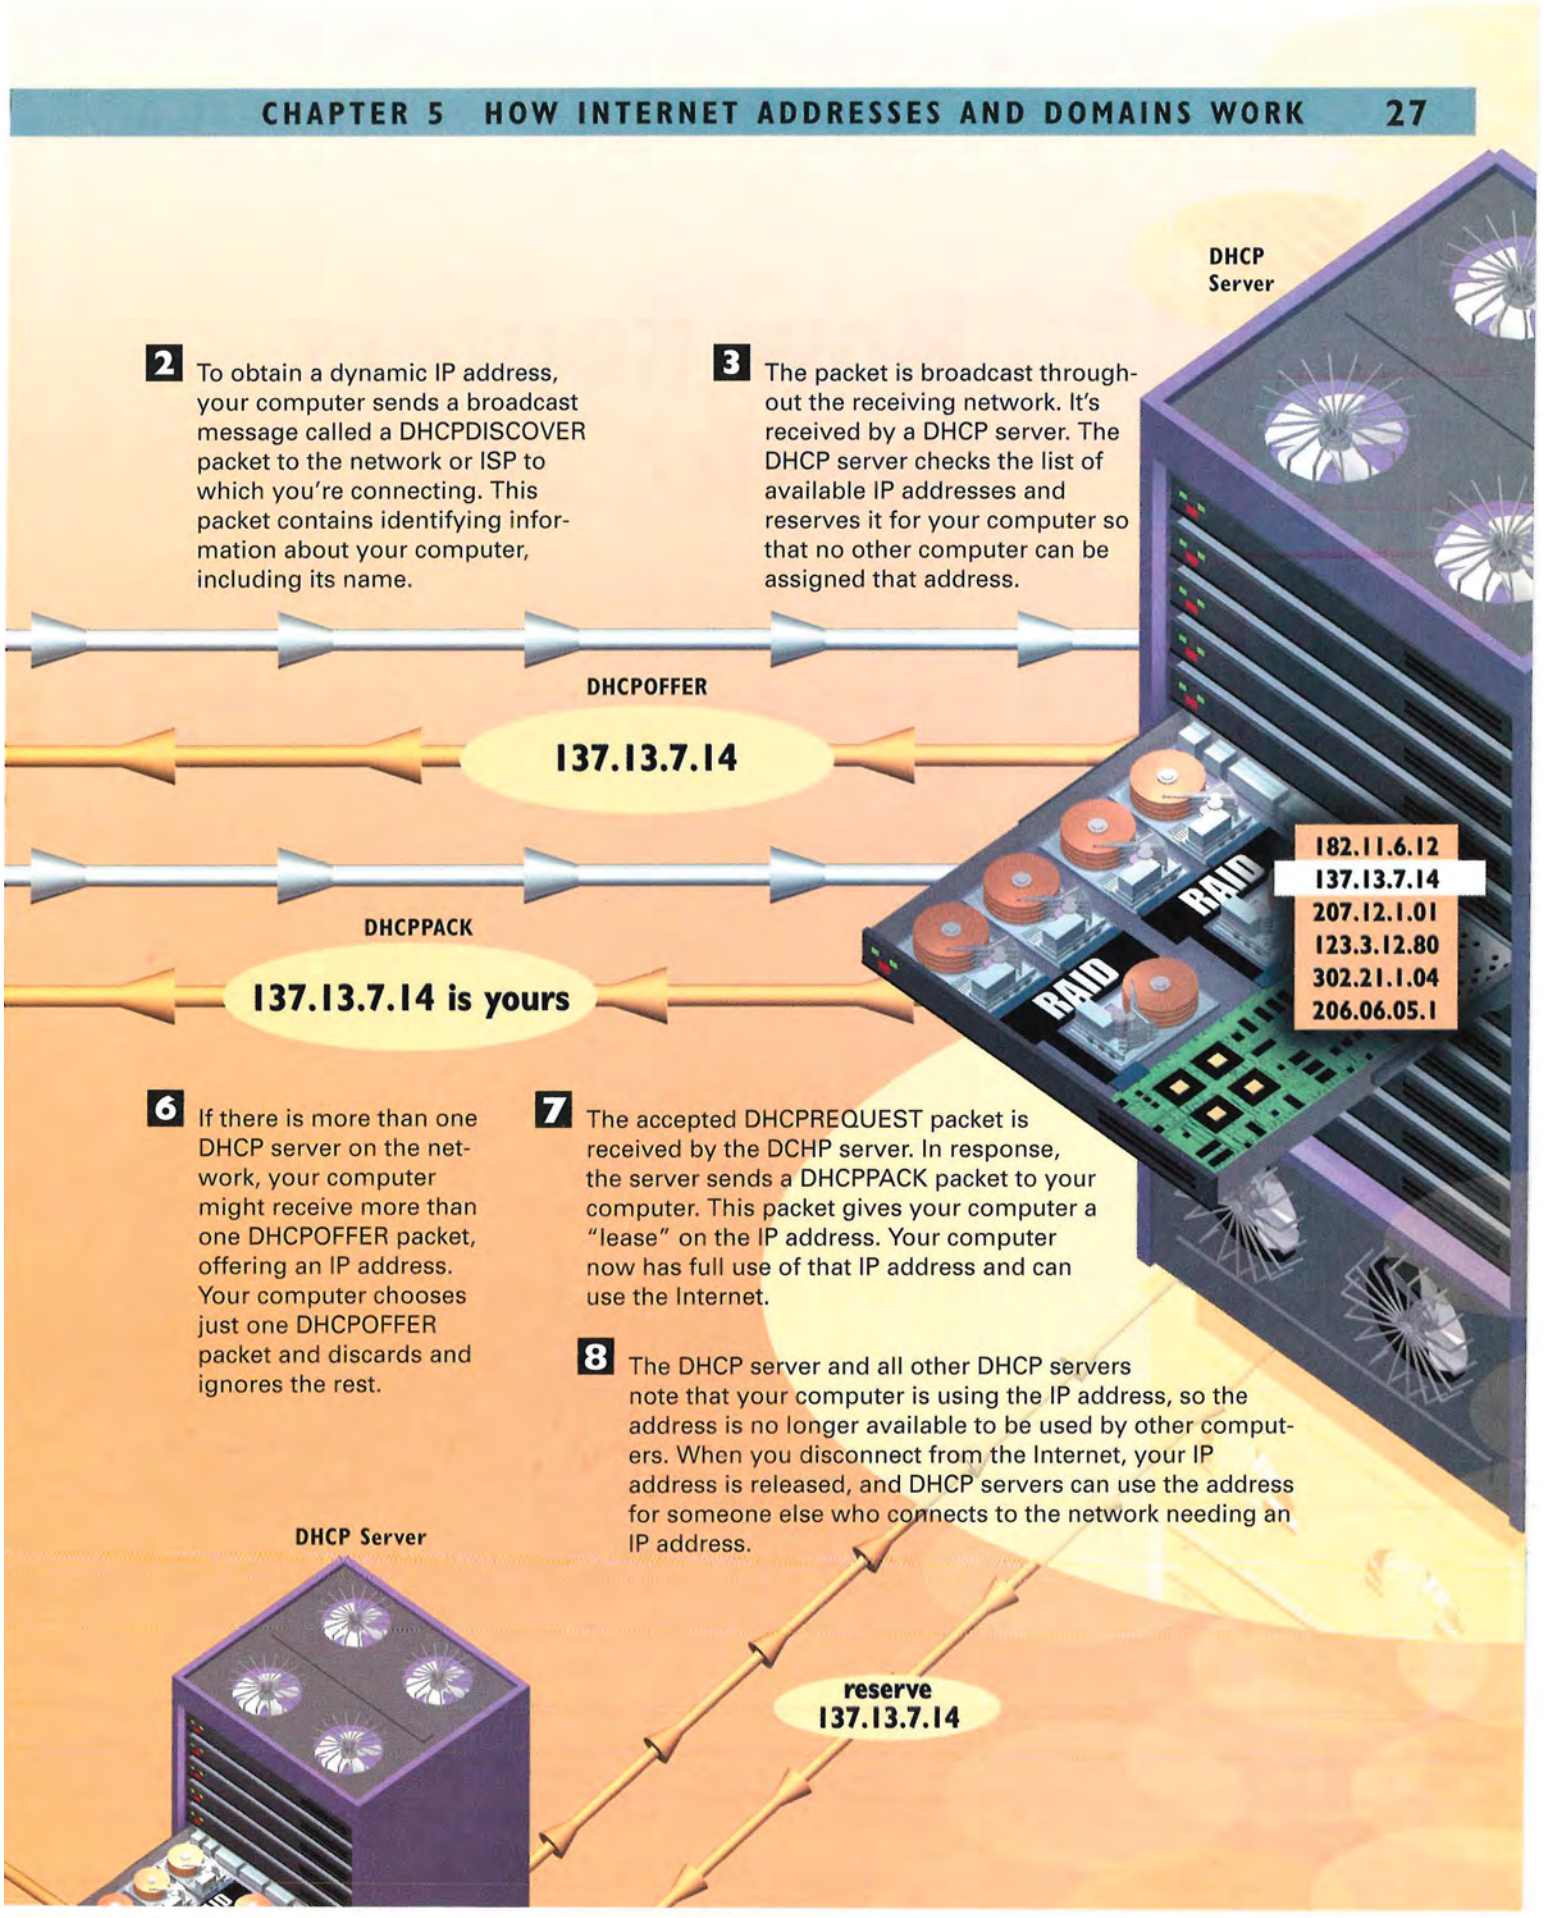
How Static and Dynamic IP Addresses Work

1 When you connect to the Internet, your computer must be identified by a number called an IP address. IP addresses come in two varieties: static and dynamic. When you have a static IP address, your computer keeps the same IP address every time you connect to the Internet. Dynamic addresses, on the other hand, are assigned every time you connect to the Internet.



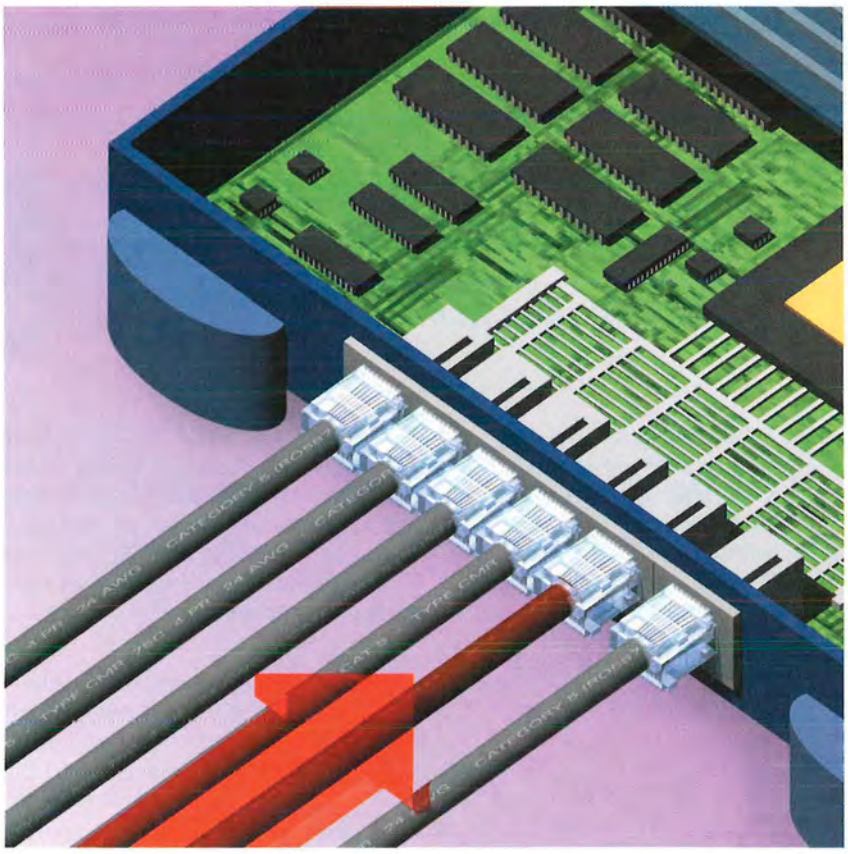
4 The DHCP server broadcasts a response, called a DHCPOFFER packet. The packet contains the number of the reserved IP address, along with a variety of other information, including the IP address of the DHCP server.

5 Your computer receives the DHCPOFFER packet, agrees to accept the reserved IP address, and broadcasts to the server a DHCPREQUEST packet. At this point, your computer doesn't yet have the IP address from the server—your computer is sending the DHCPREQUEST packet to tell the server that it would like to use the reserved address.


DHCP Server
2 To obtain a dynamic IP address, your computer sends a broadcast message called a DHCPDISCOVER packet to the network or ISP to which you're connecting. This packet contains identifying information about your computer, including its name.
3 The packet is broadcast throughout the receiving network. It's received by a DHCP server. The DHCP server checks the list of available IP addresses and reserves it for your computer so that no other computer can be assigned that address.
DHCPOFFER
137.13.7.14
DHCPPACK
137.13.7.14 is yours
RAID
RAID
182.11.6.12
137.13.7.14
207.12.1.01
123.3.12.80
302.21.1.04
206.06.05.1
6 If there is more than one DHCP server on the network, your computer might receive more than one DHCPOFFER packet, offering an IP address. Your computer chooses just one DHCPOFFER packet and discards and ignores the rest.
7 The accepted DHCPREQUEST packet is received by the DCHP server. In response, the server sends a DHCPPACK packet to your computer. This packet gives your computer a "lease" on the IP address. Your computer now has full use of that IP address and can use the Internet.
8 The DHCP server and all other DHCP servers note that your computer is using the IP address, so the address is no longer available to be used by other computers. When you disconnect from the Internet, your IP address is released, and DHCP servers can use the address for someone else who connects to the network needing an IP address.
DHCP Server
reserve
137.13.7.14

CHAPTER 6

How Routers Work



ROUTERS are the traffic cops of the Internet. They ensure all data gets sent to where it's supposed to go via the most efficient route. When you sit down at your computer and connect to the Internet and send or receive data, generally that information first must go through at least one router—and often more than one router—before it reaches its final destination.

Routers open the IP packets of data to read the destination address, calculate the best route, and then send the packet toward its final destination. If the destination is on the same network as the sending computer, such as within a corporation, the router sends the packet directly to the destination computer. If the packet is going to a destination outside the local network, the router instead sends the packet to another router closer to the destination. That router in turn sends the packet to a yet-closer router, and so on, until the packet reaches its final destination.

When routers determine the next router to receive packets, factors such as traffic congestion and the number of *hops* (routers or gateways on any given path) come into play. The IP packet carries a segment that holds the maximum hop counts it can travel, and a router will not use a path that would exceed that predefined number of hops.

Routers have two or more physical ports: receiving (input) ports and sending (output) ports. When an input port receives a packet, a software routine called a *routing process* is run. This process looks inside the header information in the IP packet and finds the address to which the data is being sent. It then compares this address against an internal database called a *routing table*. The routing table has detailed information about the ports to which packets with various IP addresses should be sent. Based on what it finds in the routing table, the router sends the packet to a specific output port. This output port then sends the data to the next router or to the destination itself.

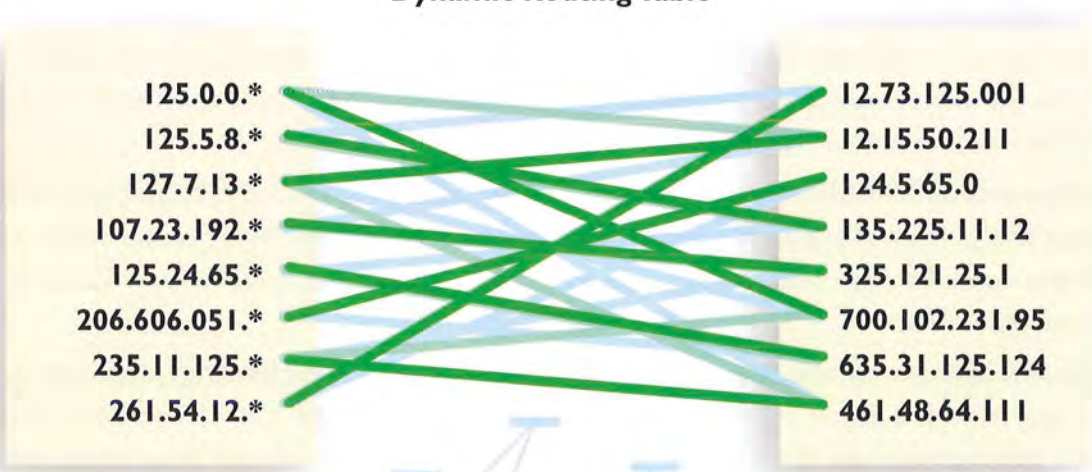
There are two kinds of routing tables: a *static* routing table and a *dynamic* routing table. The static table is simpler and specifies specific paths for packets to use to get to their final destinations. The dynamic table allows packets to have multiple routes to get to their destinations.

At times, packets are sent to a router's input port faster than the port can process them. When this happens, the packets are sent to a special holding area called an *input queue*, which is an area of RAM on the router. That specific input queue is associated with a specific input port. A router can have more than one input queue. Each input port processes packets from the queue in the order in which they were received, so the first packets sent in are the first to get processed and sent out.

If the number of packets received exceeds the capacity of the queue (called the *length* of the queue), packets might be lost. When this occurs, the TCP protocol on the sending and receiving computers have the packets re-sent.

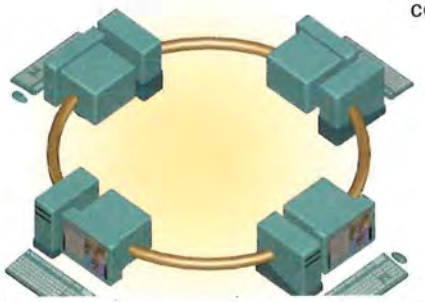
How Routers Send Data to Their Destination

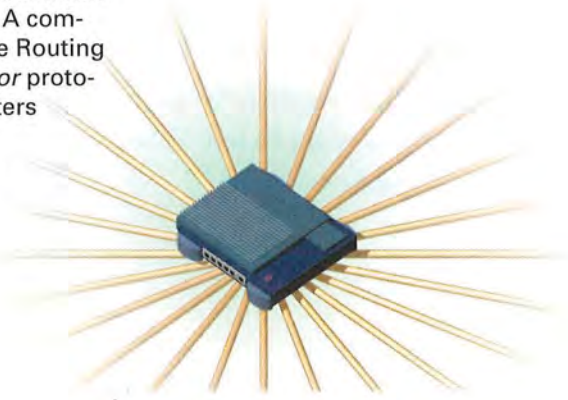
1 A router has input ports for receiving IP packets and output ports for sending those packets toward their destinations. When a packet comes to an input port, the router examines the packet header and checks the destination in it against a routing table—a database that tells the router how to send packets to various destinations.

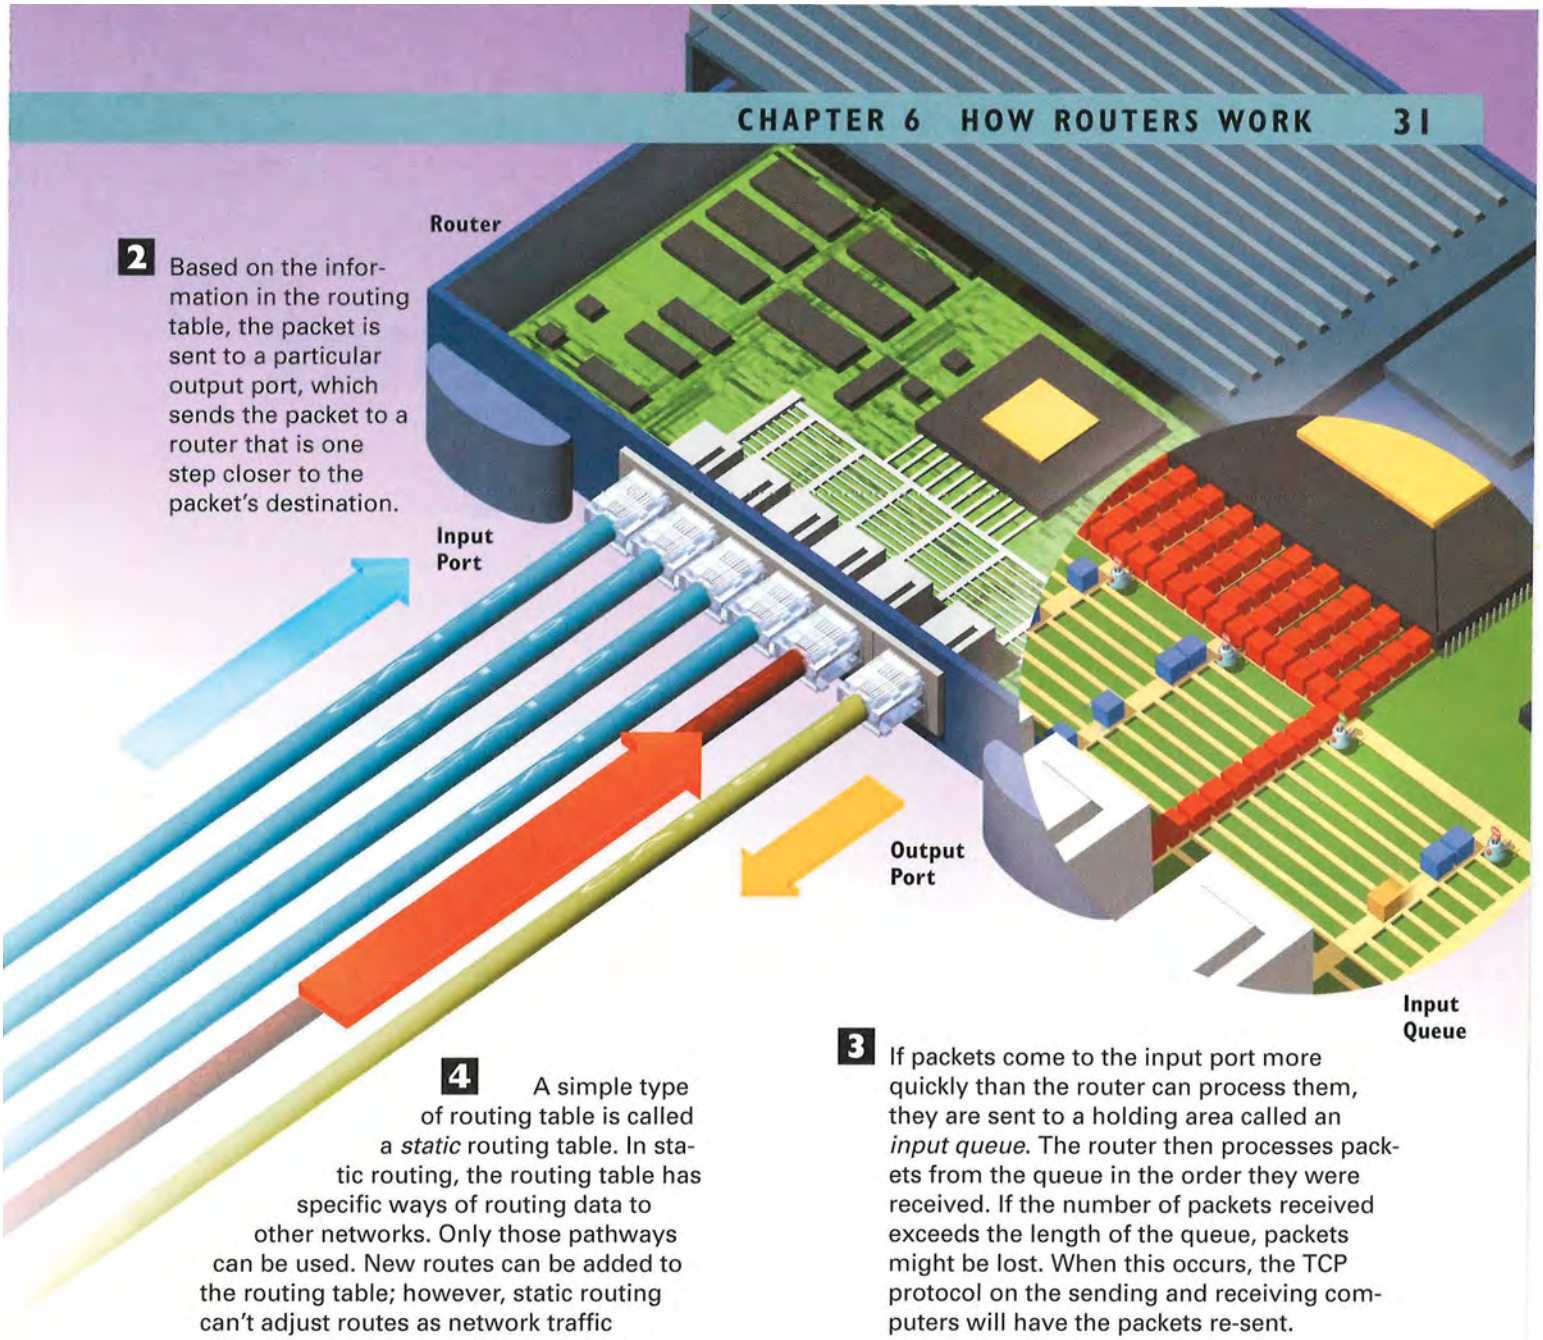


5 Dynamic routing is more useful than static routing. It allows a packet to have multiple routes to reach its final destination. Dynamic routing also enables routers to change the way they route information based on the amount of network traffic on some paths and routers. In dynamic routing, the routing table is called a *dynamic* routing table and changes as network conditions change. The tables are built dynamically by routing protocols, constantly changing according to network traffic and conditions.

6 Two broad types of routing protocols exist: interior and exterior. *Interior* routing protocols are typically used only on routers in a company's intranet, or internal network. These interior routing protocols route traffic bound only for inside the intranet. A common interior routing protocol is the Routing Information Protocol (RIP). *Exterior* protocols typically are used for routers located on the Internet. A common exterior routing protocol is the Exterior Gateway Protocol (EGP).







2 Based on the information in the routing table, the packet is sent to a particular output port, which sends the packet to a router that is one step closer to the packet's destination.

3 If packets come to the input port more quickly than the router can process them, they are sent to a holding area called an *input queue*. The router then processes packets from the queue in the order they were received. If the number of packets received exceeds the length of the queue, packets might be lost. When this occurs, the TCP protocol on the sending and receiving computers will have the packets re-sent.

4 A simple type of routing table is called a *static* routing table. In static routing, the routing table has specific ways of routing data to other networks. Only those pathways can be used. New routes can be added to the routing table; however, static routing can't adjust routes as network traffic changes so it isn't an optimal alternative for many routers.

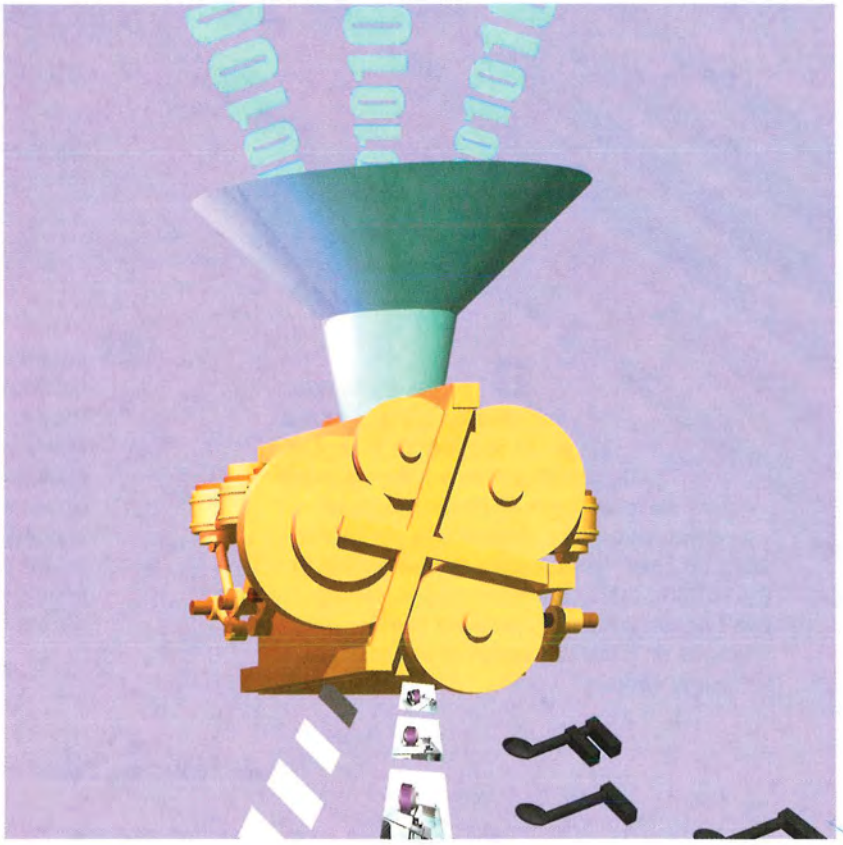
Static Routing Table

125.0.0.*	12.73.125.001
125.5.8.*	12.15.50.211
127.7.13.*	124.5.65.0
107.23.192.*	135.225.11.12
125.24.65.*	325.121.25.1
206.606.051.*	700.102.231.95
235.11.125.*	635.31.125.124
261.54.12.*	461.48.64.111

CHAPTER

7

How Internet File Types Work



THERE are millions of files on the Internet that show you pictures, let you hear music and sounds, let you to watch videos, and let you read articles and run software on your computer. Some of these files can be used on any computer, regardless of whether it is a Macintosh or an IBM-compatible PC. Graphics files, sound files, and video files, for example, can be played on many kinds of computers, as long as those computers have the special software (often called players or plug-ins) required to view, run, or listen to them. Some files require special hardware to work. Other files, such as software programs you download, can run on only a PC or a Mac, or they might require extra memory.

In general, on the Internet you'll find two types of files: ASCII (American Standard Code for Information Interchange) files and binary files. ASCII files are often referred to as *plain text files*, *text files*, or *ASCII text files*. They contain ASCII codes—computer codes that represent the characters on your screen, such as the uppercase letter A or a dollar sign ($). ASCII text files contain nothing but this simple character data. Binary files, by way of contrast, contain special coded data and can be run or read only by specific computers and software.

Interestingly enough, though, Web pages are in fact made up of plain text files. Although they are plain text files, they use special codes in the Hypertext Markup Language (HTML) that instructs your browser how to display the Web page.

In contrast to plain ASCII text files, some files contain sophisticated formatting and graphical information. Files that adhere to the PostScript page description language are of this type. These files are in fact ASCII files, but they contain information about how to format and print the file. HTML pages—the pages that make up the World Wide Web—are also plain ASCII text files that contain coding information that tells your browser how to display a Web page.

Other types of files contain formatting and graphical information, such as those in the Adobe Acrobat PDF format, but many of them are binary files instead of ASCII files. To view or print these types of files, you need special software readers. In the case of PostScript files, you generally need a special PostScript printer that can print files in that language. Files of this type can be viewed in two ways—either when you're online and connected to the Internet or when you're offline and not connected. To view the files when online, you use special plug-in modules or helper applications for your Web browser. To view the files when offline, you need software readers, although in some instances you also can use the same online viewers as well.

Sound, visual image, animation, and video files are also common on the Internet. These files are all binary files, which are made up of bits represented by 0s and 1s. Files of this type are often large, requiring special software readers, players, and sometimes hardware to play and read them. Some of these files can be played or viewed only on a specific type of computer, whereas others can be played or viewed on many types of computers if you have the right software. Some of the files, such as streaming audio files and streaming video files, can be viewed while you're online, connected to the Internet. Other types of files must be viewed with special readers or players when you're offline.

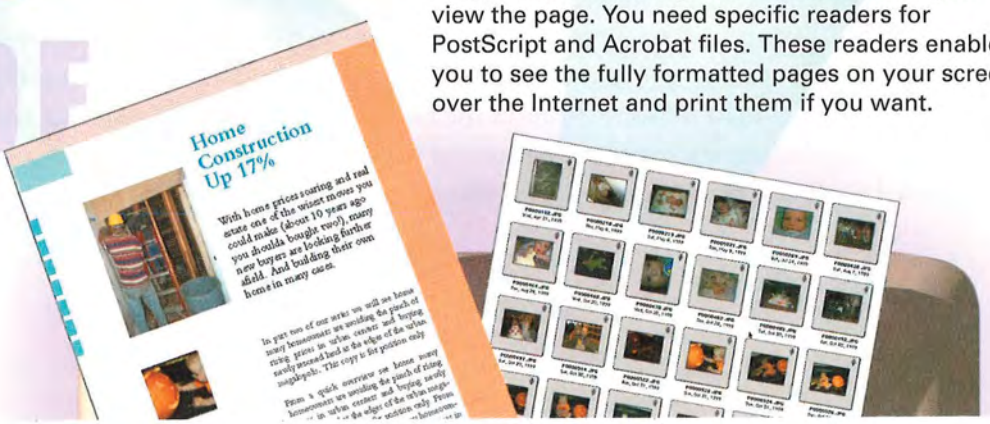
Understanding Internet File Types

File types Many types of files exist on the Internet. Each file has information contained in its header. The header is simply the first line of a file or a specific number of bytes at the beginning of a file. Software that can read or view the files looks into the header of a file to distinguish which kind of file it is and then processes the file accordingly.



Graphics formats You'll find graphics files in many formats, including GIF, JPEG, PNG, PCX, and TIF. GIF and JPEG are two of the most common graphics formats. They enable the exchange of image files between different types of computers and can be downloaded to your computer quickly. A GIF reader or JPEG reader can read binary GIF and JPEG files and display them on your computer. Graphics programs typically read both types of files. A GIF file contains the binary data that displays an image when viewed with the proper reader. Your computer's video card takes information from the reader and displays it on your computer's monitor. Both GIF and JPEG are commonly used graphics formats in Web pages. Your Web browser has the capability to read and display these files on Web pages without needing any added software.



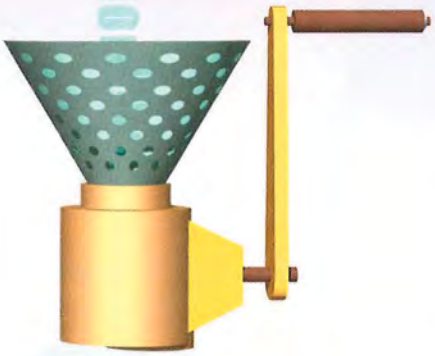


Document types Several file types, such as PostScript (which ends in a .PS extension) and Adobe Acrobat (.PDF) files, contain complex information about documents. This includes placement of pictures, size and type of fonts, and the complex shapes and formatting information necessary to view the page. You need specific readers for PostScript and Acrobat files. These readers enable you to see the fully formatted pages on your screen over the Internet and print them if you want.

Multimedia files You'll find many kinds of binary multimedia files on the Internet that let you listen to or play sounds, music, and videos. Windows sound files (which end in a .WAV extension) and Macintosh sound files are common types of files that you can download to your computer and play with a sound player, as are music files in the MP3 format (these end in an .mp3 extension). Other sound files, such as those created by Real Audio software, enable you to listen to the sounds while the file is downloading. This process is called *streaming audio*. Windows .AVI files are common animation files. Macintosh QuickTime and MPEG files (which end in an .MPG extension) are common video formats. Other animation files enable you to watch the video while the file is downloading. This process is called *streaming video*. You need special software to use all these types of files.



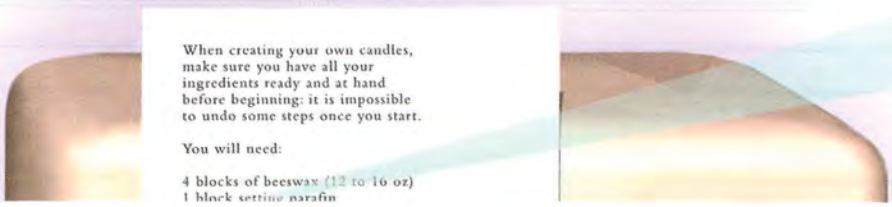
Downloadable software Many executable programs can be found on the Internet that you can download to your computer and use just like any other type of software. Because these software files can be large and take a long time to download, they often are compressed to make them download more quickly. After the file is on your computer, you uncompress it with special decompression software and run it like any other program. Frequently, PC software has been compressed with PKZIP and has a .ZIP extension, whereas Macintosh software can be compressed with a variety of compression software. In general, PC files can be used only with PCs, whereas Macintosh files can be used only with Macintoshes.

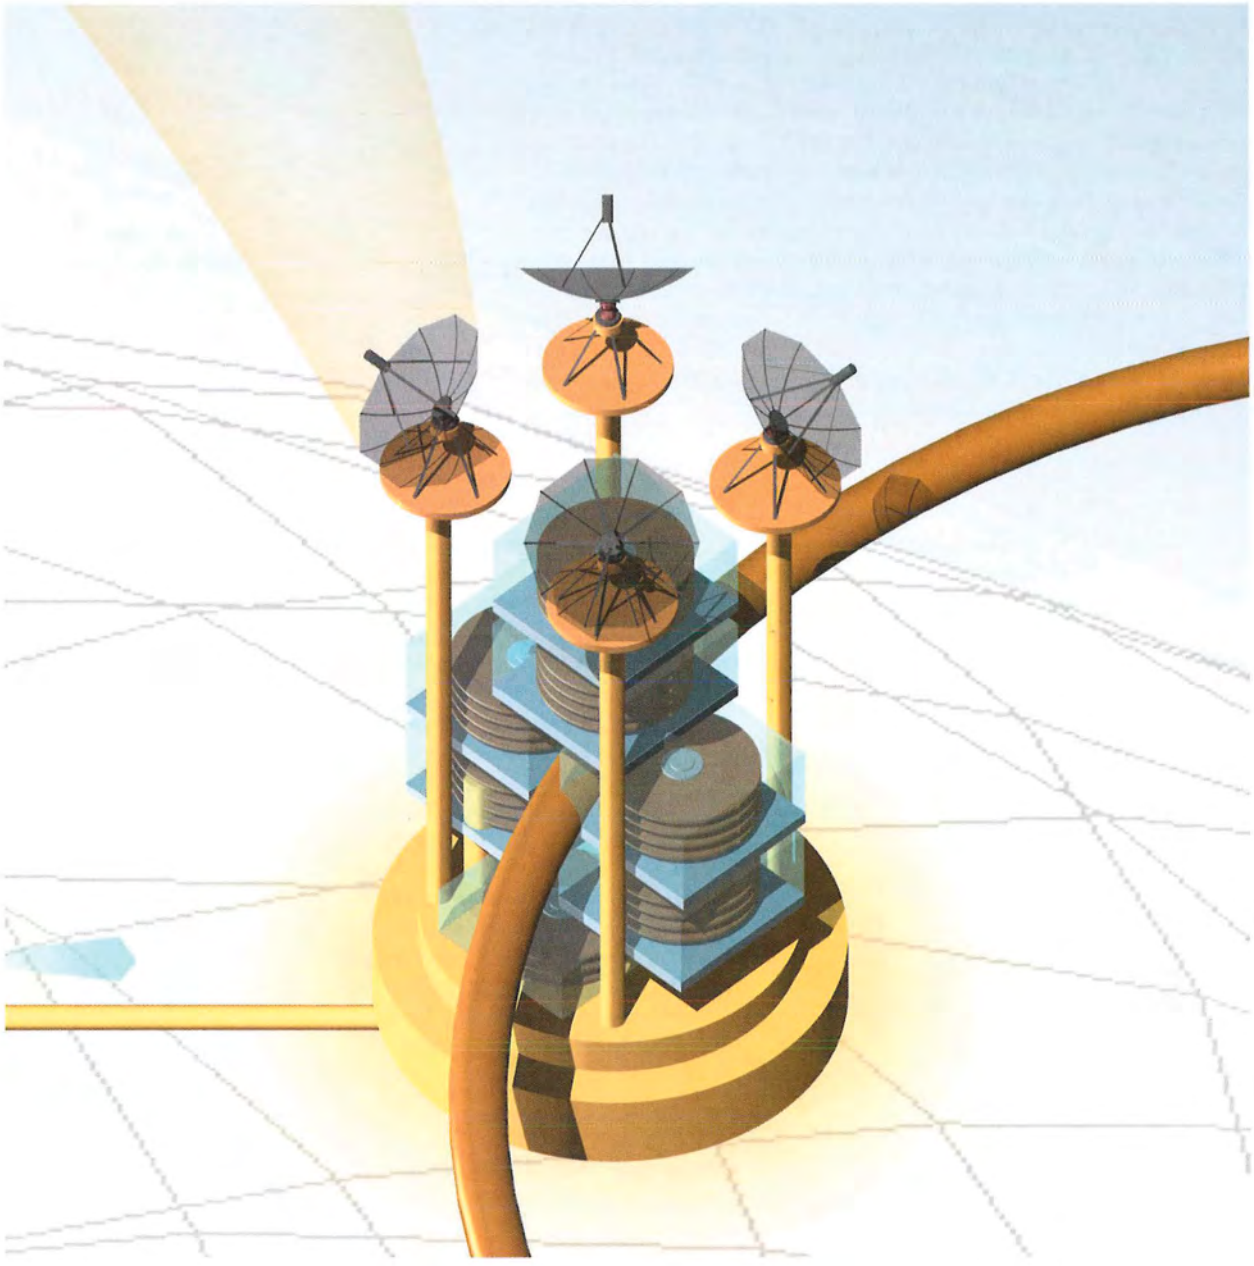






Plain text You'll find many ASCII text files on the Internet that you can read with a text editor or word processor. These can be articles, FAQs (frequently asked questions and their answers), or any other type of informational files.





PART 2

CONNECTING TO THE INTERNET

Chapter 8: How Computers Connect to the Internet
40

Chapter 9: How a Modem Works
44

Chapter 10: How Internet/Television Connections Work
48

Chapter 11: How the Digital Subscriber Line (DSL) Works
56

Chapter 12: How ISDN Works
60

Chapter 13: How You Connect to the Internet from Online Services
64

Chapter 14: How Wireless and Palmtop Connections Work
68

Chapter 15: How Home Networks Work
78

YOU can connect to the Internet in many ways—and many more ways pop up practically every day. They range from simple telephone dial-in connections to high-speed cable and digital subscriber lines (DSL) to satellite connections, TV connections, connections at work and home via local area networks (LANs), and even connections via cellular telephones. This part of the book looks at the myriad ways people and computers gain access to the Internet.

One general rule is true about Internet connections: the faster, the better. People want the fastest connection possible because many pictures, sounds, and videos are available on the Internet. Today, the two most common ways you can connect to the Internet are through a corporate or university LAN or over telephone lines. Direct connections over LANs are generally faster than telephone-line connections—although a connection via new telephone technology called DSL is changing that.

This section looks at all the ways computers can connect to the Internet. Chapter 8, "How Computers Connect to the Internet," provides an overview of the types of Internet connections that are possible. Not only will you examine various kinds of network connections and phone-line connections, but you'll also look at newer, high-speed connections that are becoming available; cable modems and DSL are two such examples.

Chapter 9, "How a Modem Works," examines how the most common way of accessing the Internet works—via a modem. You'll get an inside look at what happens when you dial the phone and make your connection.

Chapter 10, "How Internet/Television Connections Work," looks at what might become one of the primary ways many of us connect to the Internet—through some type of television connection. This chapter looks at using cable television as a way to get onto the Internet and how to use an ordinary television to browse the Internet. It also peers inside a newer technology—Internet-enhanced TV—that adds interactivity to your television set and makes your Internet experience more compelling.

Chapter 11, "How the Digital Subscriber Line (DSL) Works," looks at one of the newer high-speed ways to connect to the Internet—DSLs. These connections enable you to use existing telephone lines to hook up to the Internet at extremely high speeds. To use it, you need a special DSL modem, and your phone company needs special modems as well.

A related technology to DSL is Integrated Services Digital Network (ISDN). ISDN enables you to connect at high speeds to the Internet via normal, existing telephone lines. Not every area in the country has ISDN yet. In some areas it can be an expensive option, but in others it is reasonably priced. It requires a special modem, as well. Chapter 12, "How ISDN Works," explains that technology.

Online services, such as CompuServe, America Online, and The Microsoft Network, at one time competed directly with the Internet for users. Now, instead of competing, those online services have joined the Internet club and provide Internet access to users. An online service can be one of the simplest ways to get an Internet connection. When you connect to an online service, you also get a full connection to the Internet. Chapter 13, "How You Connect to the Internet from Online Services," examines how online services make that Internet connection.

Chapter 14, "How Wireless and Palmtop Connections Work," looks at two of the newest ways to access the Internet. With a satellite connection, you can connect to the Internet using a satellite dish, much like the one now used for satellite TV. In fact, the dish can do double-duty—giving you both satellite TV and a high-speed Internet connection. The connection isn't yet as high speed as a cable modem or DSL connection, but it is much faster than a traditional modem. This chapter also looks at how palmtop devices, such as the Palm from 3COM can access the Web using a traditional modem. More exciting is the ability to connect to the Internet without any wires at all—using cellular technology. So this chapter discusses how cellular palmtops and cellular telephones can get onto the Web. Those connections aren't as high a speed as other ways of getting onto the Internet yet, but one day, they might be faster than normal phone-line connections.

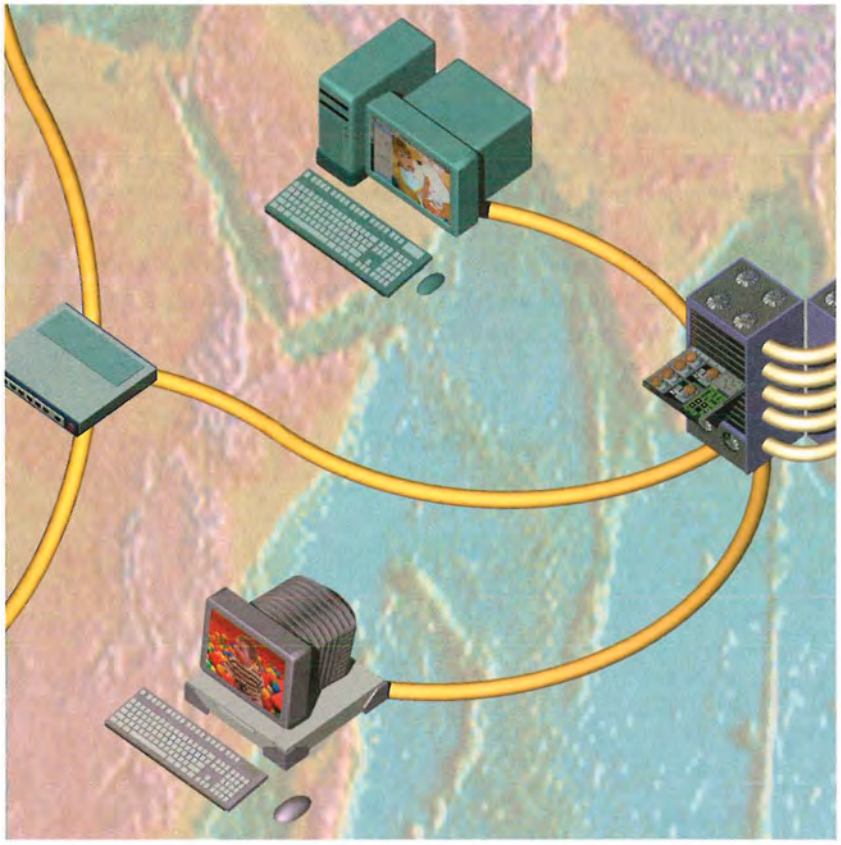
Chapter 15, "How Home Networks Work," examines a fast-growing phenomena—using networks at home to connect to the Internet. Many homes have more than one computer, and users want to be able to share a high-speed connection, such as a cable modem or DSL modem. Home networks make that possible. And increasingly, other devices at home will be connected to the Internet, not just PCs—and this chapter looks at them as well.

As you'll see in this part of the book, connecting to the Internet will become increasingly easy—and at increasingly higher speeds. Not only will you be able to do things more quickly on the Internet, but you will also be able to access entirely new services that will contain video, animation, and other high-bandwidth content. And as satellite connections and the capability to connect wirelessly via small handheld devices and cellular telephones come online, the Internet will eventually be able to be accessed anywhere—not just when you happen to be sitting in front of a PC.

CHAPTER

8

How Computers Connect to the Internet



THERE are many different ways your computer can connect to the Internet, ranging from dial-in connections to LANs (local area networks) to wireless connections to connections over cable TV wires. If you are connected to a LAN or campus network at your business or school, you already might be connected to the Internet. If the LAN or campus network you're on is connected to the Internet via a router or bridge, your computer is also connected to the Internet. Often, this offers higher-speed access than when you dial in to the Internet. If you're not connected to the Internet via a network, a variety of options are available for accessing the Internet, and many more appear every day.

The most common option is to find an Internet service provider (ISP) or online service, such as America Online, that lets you dial in to it with your home computer. When you connect this way, you establish a Serial Line Internet Protocol (SLIP) or Point-to-Point Protocol (PPP) connection. (SLIP is an older type of connection, and its use is fading away.) When you dial in to the Internet with these protocols, your computer becomes a part of the Internet and you can access all of its vast resources, such as browsing the Web. PPP is newer and more stable than SLIP and can retransmit packets if they get garbled, which is a fairly common situation you'll encounter when sending information over telephone lines.

Many ways are available to connect to the Internet in addition to dialing in using a traditional modem. Increasingly popular are special cable modems that enable you to use your coaxial television cable to access the Internet. Your local cable company—or another private company—might be able to provide you with this service. Speeds can be up to 100 times faster than regular modem speeds.

Also increasingly popular are Digital Subscriber Line (DSL) connections to the Internet. Similar to cable modems, these offer very high-speed connections to the Internet. The price of DSL will vary according to the connection speed.

Integrated Services Digital Network (ISDN) lines provide relatively high-speed connections to the Internet. Typically, these speeds range from 64Kbps to 128Kbps. You'll need a special ISDN modem and an ISP that offers ISDN access.

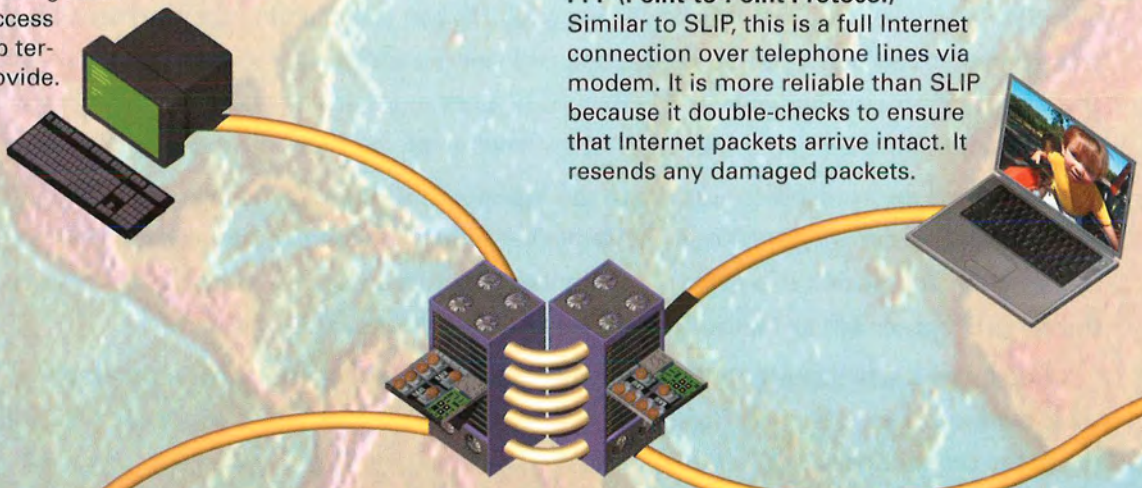
The newest type of connections are wireless connections. You can now send and receive e-mail, browse the Web, and access other Internet resources via a cellular phone or cellular palmtop connection. These connections are slower than normal dial-in connections but are becoming popular because you can access the Internet with them wherever you are.

If you're not yet connected and are interested in finding out more about an Internet connection, start at your local college or library. Contact an ISP and find out about their fees. Many national and local service providers exist, so shop around. Lastly, the major online services now include Internet access, so you easily can use them to connect to the Net.

Connecting Your Computer to the Internet

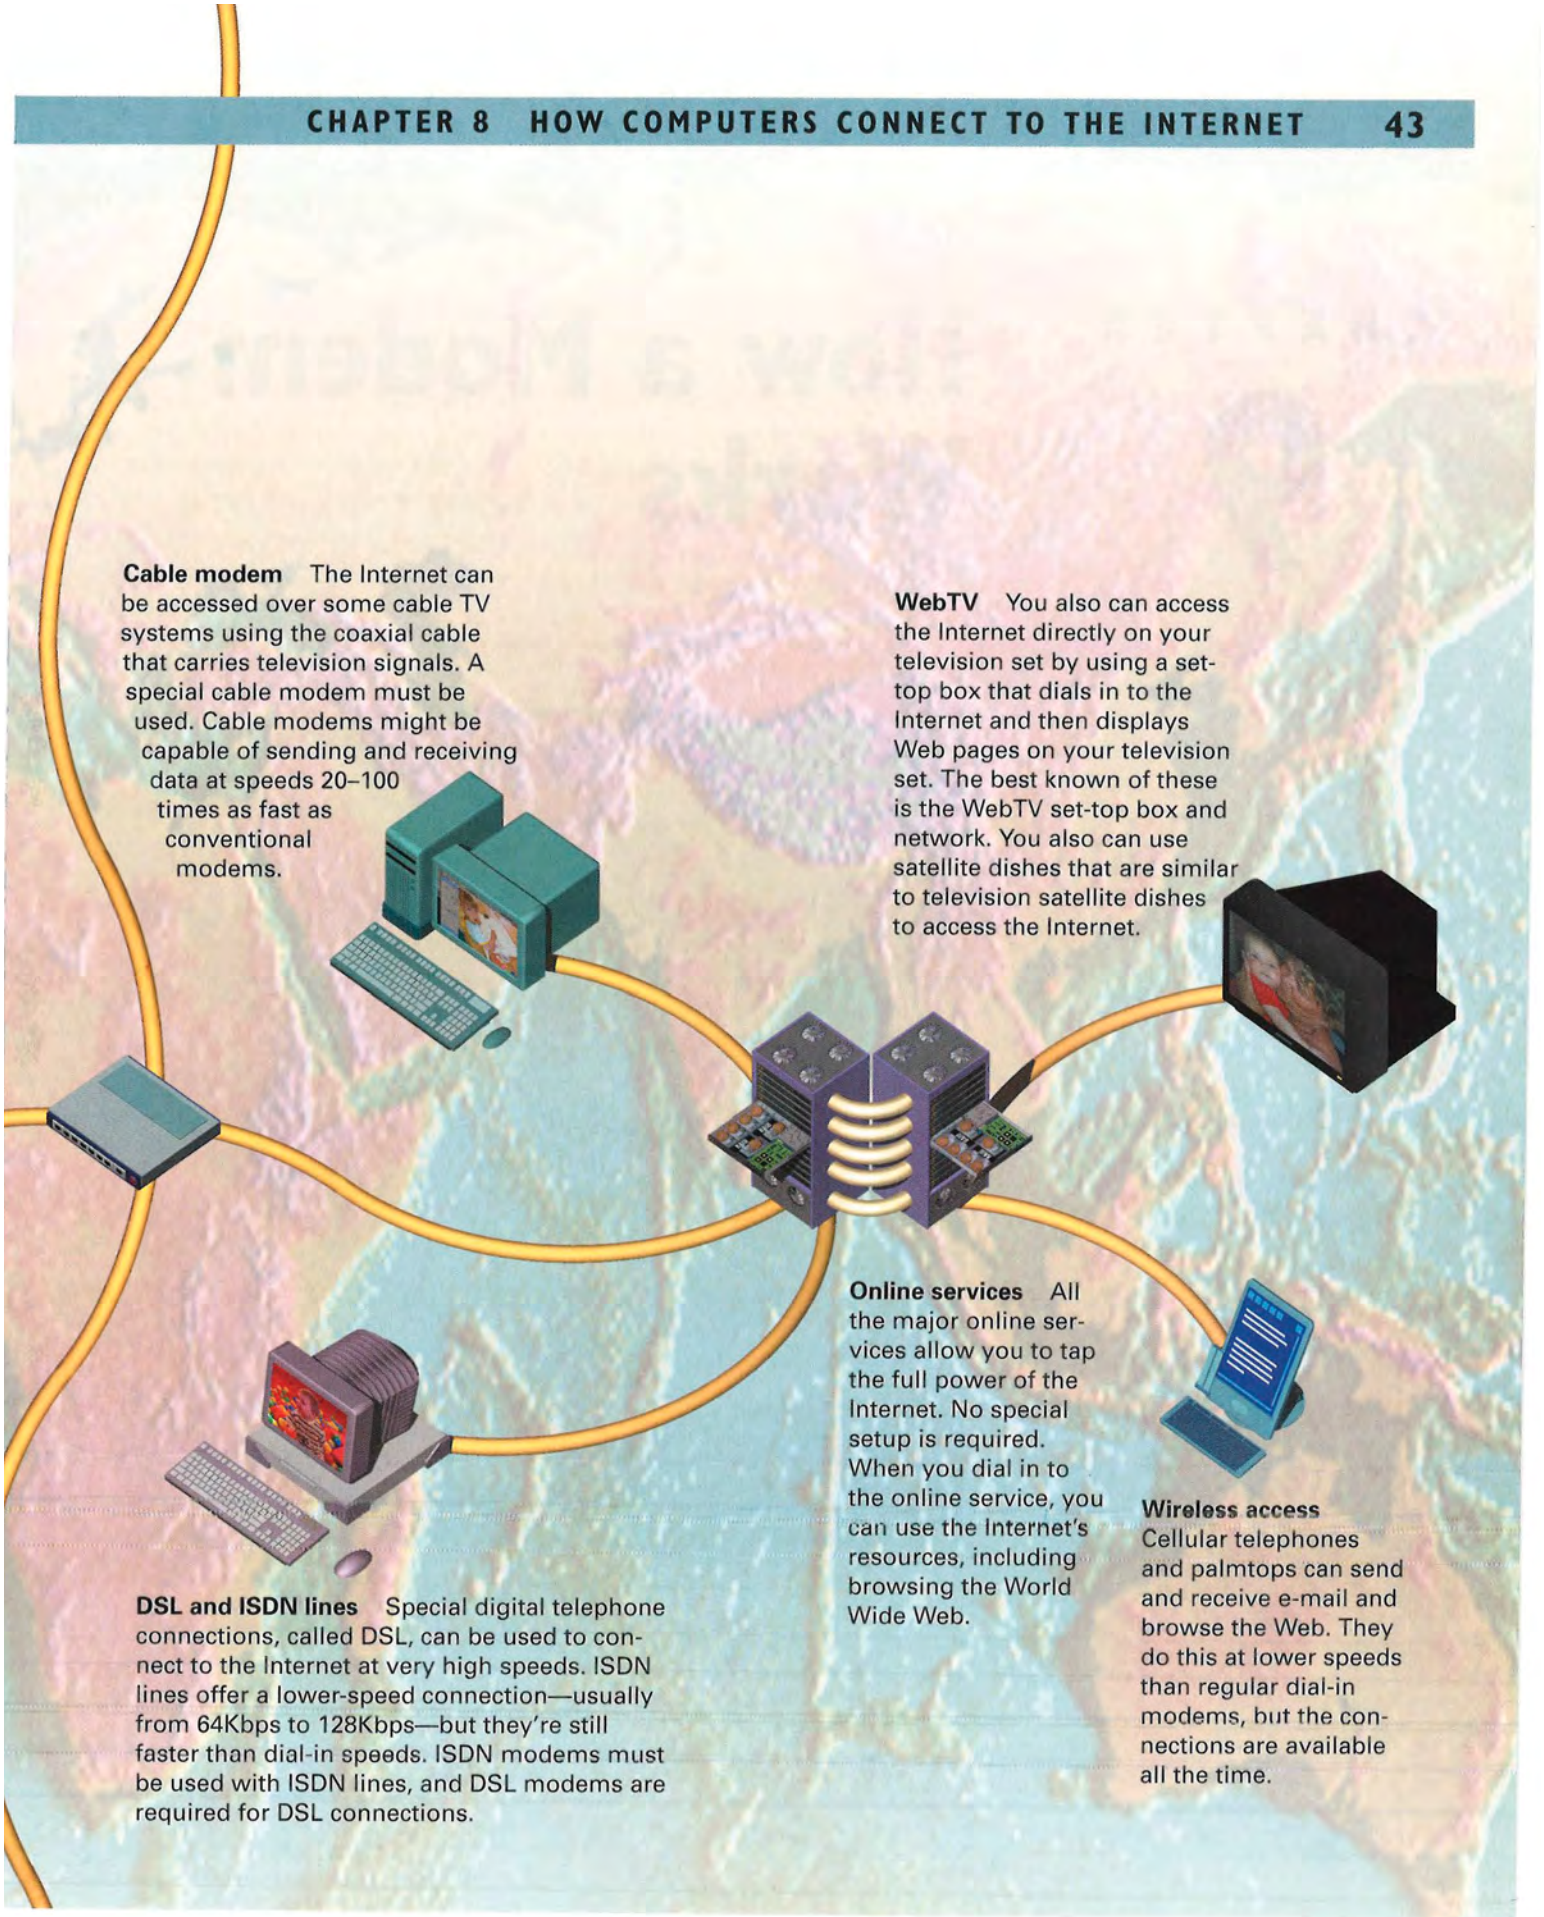
Dumb terminal A terminal that is connected to a mainframe, minicomputer, or other kind of large computer. This type of connection often can be found in libraries or universities, although those institutions are moving toward offering fuller Internet access than dumb terminals provide.

PPP (Point-to-Point Protocol) Similar to SLIP, this is a full Internet connection over telephone lines via modem. It is more reliable than SLIP because it double-checks to ensure that Internet packets arrive intact. It resends any damaged packets.



SLIP (Serial Line Internet Protocol) A full Internet connection over telephone lines that sends Internet packets via 9,600bps modems or better.

Direct connection LANs or large computers such as mainframes can be directly connected to the Internet. When a LAN is connected to the Internet, all the computers on the network can have full access to the Net. This type of access is common inside corporations.



Cable modem The Internet can be accessed over some cable TV systems using the coaxial cable that carries television signals. A special cable modem must be used. Cable modems might be capable of sending and receiving data at speeds 20–100 times as fast as conventional modems.

WebTV You also can access the Internet directly on your television set by using a set-top box that dials in to the Internet and then displays Web pages on your television set. The best known of these is the WebTV set-top box and network. You also can use satellite dishes that are similar to television satellite dishes to access the Internet.

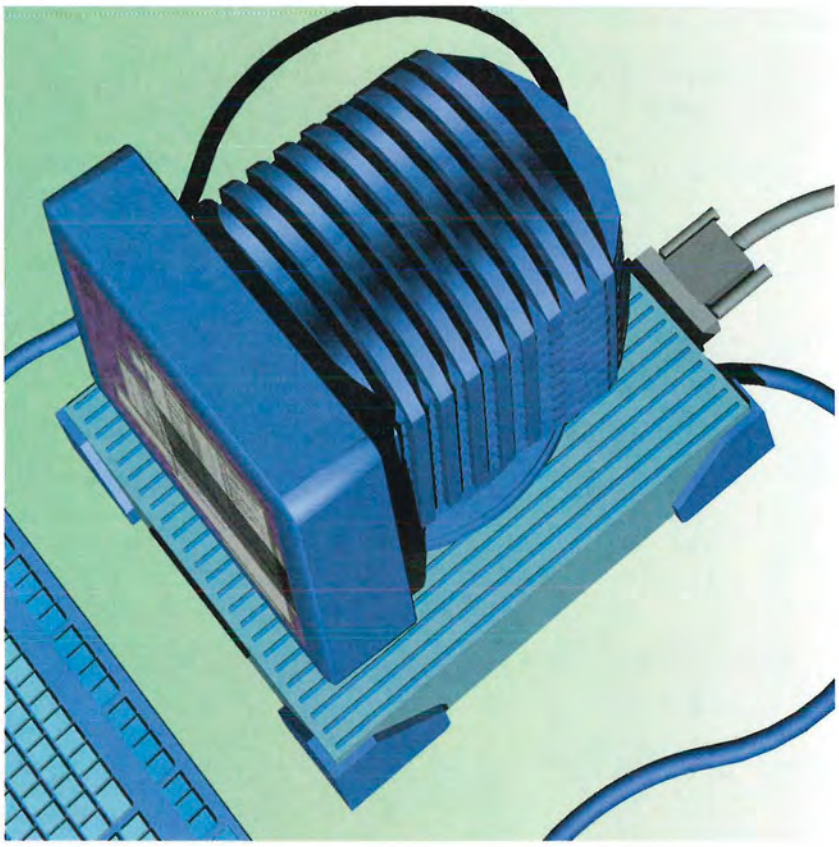
Online services All the major online services allow you to tap the full power of the Internet. No special setup is required. When you dial in to the online service, you can use the Internet's resources, including browsing the World Wide Web.

Wireless access Cellular telephones and palmtops can send and receive e-mail and browse the Web. They do this at lower speeds than regular dial-in modems, but the connections are available all the time.

DSL and ISDN lines Special digital telephone connections, called DSL, can be used to connect to the Internet at very high speeds. ISDN lines offer a lower-speed connection—usually from 64Kbps to 128Kbps—but they're still faster than dial-in speeds. ISDN modems must be used with ISDN lines, and DSL modems are required for DSL connections.

CHAPTER 9

How a Modem Works



MOST people connect to the Internet in the most old-fashioned of ways—using a modem. Modems enable your computer to send information to, and receive information from, other computers. In doing so, modems enable you to do things such as hook up to the Internet, browse the World Wide Web, and send and receive e-mail.

Typically, when you use your modem to connect to the Internet, you dial in to an Internet service provider (ISP), for example America Online. When you dial in to and connect to your ISP, you're in fact connecting to a modem attached to a larger computer called a *server*. ISPs typically have banks of hundreds or thousands of modems that accept dial-ins from subscribers trying to connect. If you've ever tried to dial in and been unable to connect (and to be frank about it, who hasn't?), the modems might all be busy, or the receiving computer might be handling too many simultaneous calls.

Modems are controlled by your computer and communications software by a set of command called the *AT command set* (also known as the Hayes command set, named after one of the original modem manufacturers, Hayes). It's a language that instructs the modem on what to do at various points during a communications session, such as opening up a line and sending out tones that the telephone system can understand.

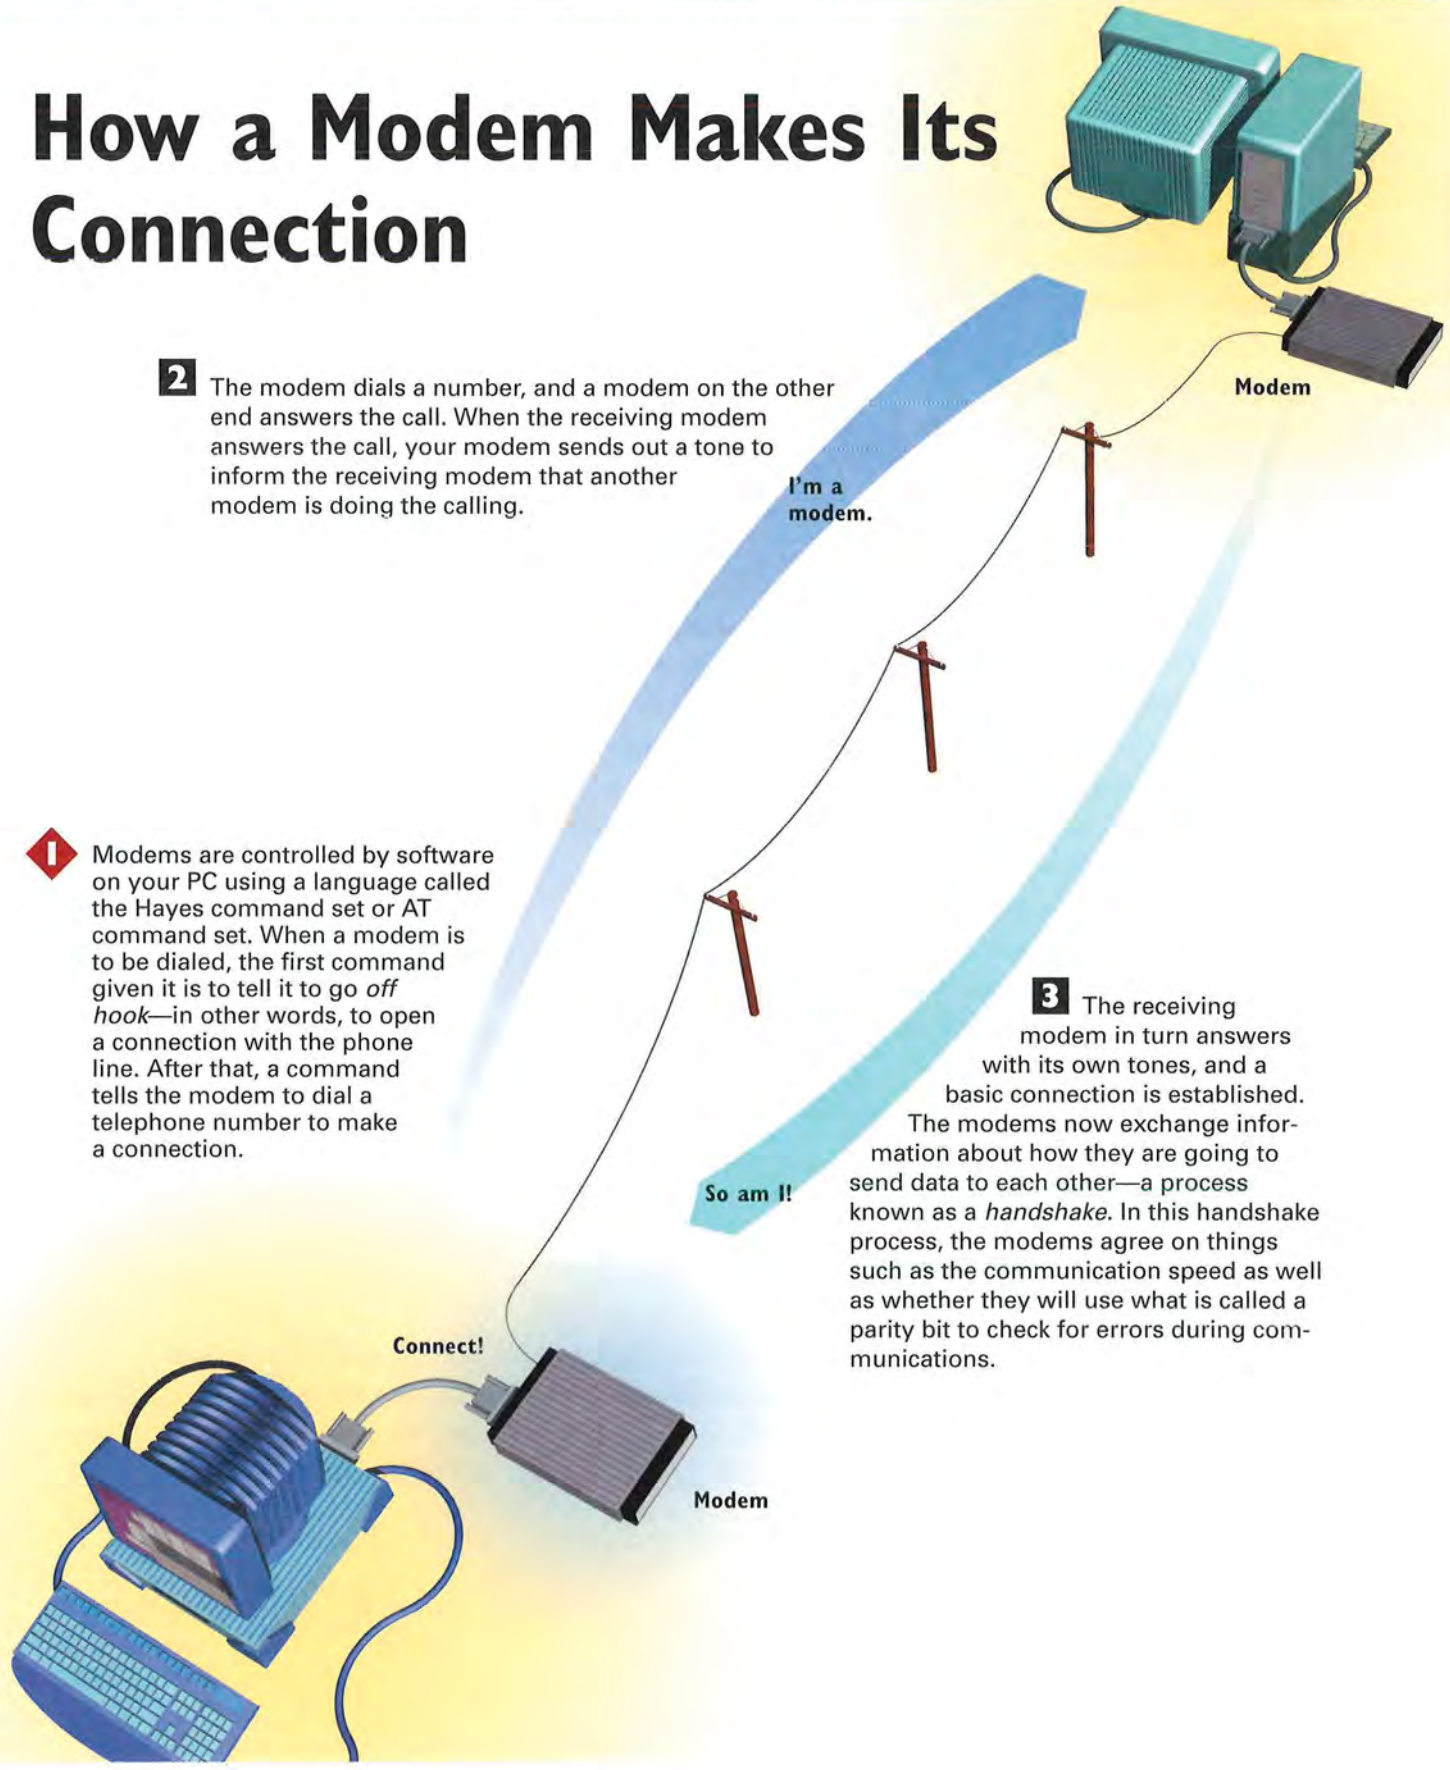
Two types of modems exist—external modems and internal modems. *External* modems are connected to your computer via a serial port, whereas internal modems are attached directly to the computer's motherboard. Typically, external modems cost more, not because they're better, but because they're more expensive to manufacture. External modems need to be put inside a case of some sort and often have lights on them. But however they're attached, they work very much the same.

The lights on an external modem are not there just to look pretty—they have specific purposes. From looking at whether lights are on or off, or are steady or flickering, for example, you'll be able to see things such as when your modem takes control of the phone line, when it has made a connection, when it is sending and receiving data, and whether it's connected at the highest speed possible.

In the modem's earliest days, they could transfer data at only a very slow rate—300 bits per second. Today's modems typically are rated at a much higher speed, 56 kilobits per second. Generally, though, when you dial in to the Internet using your computer, you won't connect at 56 kilobits per second and will usually come in at a lower speed. Things such as line noise force modems to connect at lower speeds, so that data isn't lost during the call. And even in perfect conditions, you can't connect at higher than about 53 kilobits per second.

Increasingly, people use much higher-speed devices, such as cable modems and DSL modems, to connect to the Internet. (Even though these devices are called modems, they aren't true modems and use a different technology to connect to the Internet.) But still, for the foreseeable future the majority of people will continue to connect to the Internet the old-fashioned way—over telephone lines, using a modem.

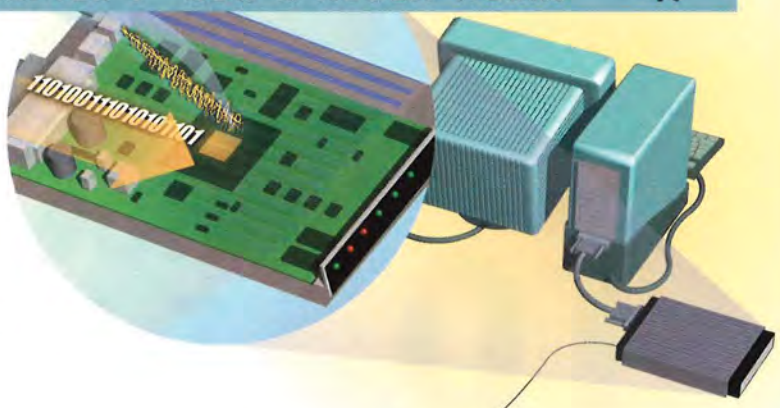
How a Modem Makes Its Connection



2 The modem dials a number, and a modem on the other end answers the call. When the receiving modem answers the call, your modem sends out a tone to inform the receiving modem that another modem is doing the calling.

1 Modems are controlled by software on your PC using a language called the Hayes command set or AT command set. When a modem is to be dialed, the first command given it is to tell it to go *off hook*—in other words, to open a connection with the phone line. After that, a command tells the modem to dial a telephone number to make a connection.

3 The receiving modem in turn answers with its own tones, and a basic connection is established. The modems now exchange information about how they are going to send data to each other—a process known as a *handshake*. In this handshake process, the modems agree on things such as the communication speed as well as whether they will use what is called a parity bit to check for errors during communications.



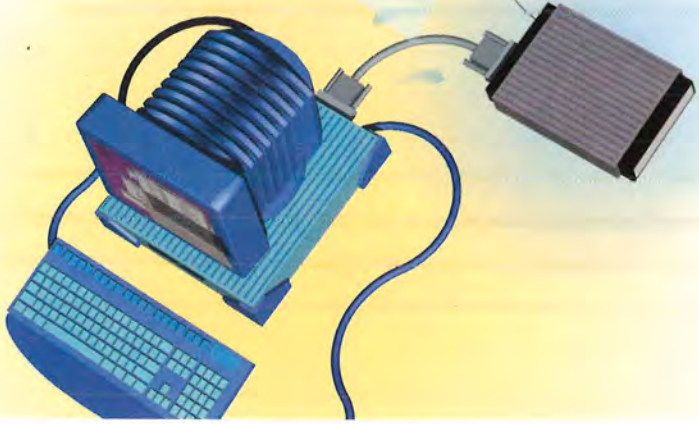
6 The analog signals reach the receiving modem where they are changed by the receiving modem—demodulated—from analog data back to digital data and sent into the computer via the serial port. This modulating and demodulating of data is what gives the modem its name: MOdulating/DEModulating.

4 Now that the connection is established, the modems can start exchanging data. Your computer works with digital data—binary bits of information that are either on or off. Your PC sends this binary data to be communicated to the modem.





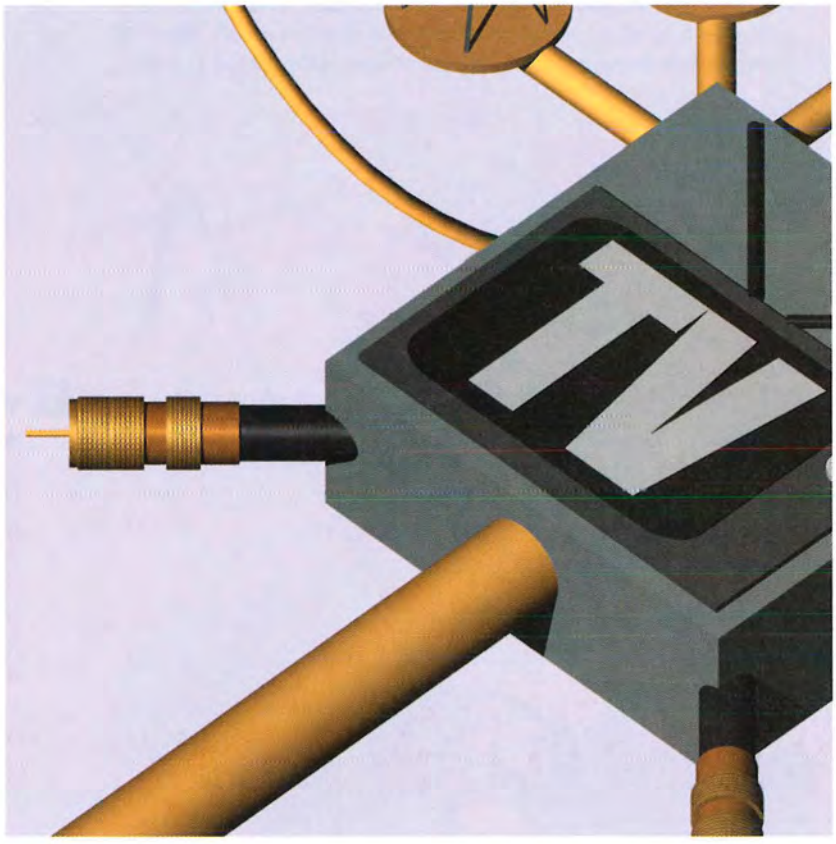
5 The telephone system, unlike a computer, works with analog information—streams of continuous electric current that vary in frequency and strength. The digital data in your computer must be changed to analog information so it can be sent via the telephone system. The modem takes the digital information and changes it—modulates it—into analog signals. The information travels along the telephone system as analog data.



CHAPTER

10

How Internet/ Television Connections Work



WHEN people first began talking about the "information superhighway" several years ago, it wasn't the Internet they were talking about. Instead, it was television—and specifically cable television—that they believed would change the way we lived and worked. There were visions of 500 TV channels, "interactive television," shopping from home, and customized news available whenever you wanted it. This information superhighway was going to hook us all together electronically so we could more easily communicate and get information, services, goods, and entertainment. It didn't pan out that way. Instead, the Internet has become an information superhighway that can do almost everything people imagined could be done by using cable TV.

How much longer the Internet will be the sole driving force behind the information superhighway remains to be seen. Every day, TV and the Internet are drawing closer to each other. The Internet is gaining more TV-like qualities, such as the capability to play videos and music, and even the capability to broadcast live video feeds. Television technology is also developing to be capable of using the Internet to add interactivity to the television experience.

In fact, television and the Internet are already merging. Soon you'll be able to watch a sporting event and simultaneously chat with others while you're watching—all on the same screen. Furthermore, when a batter comes up to bat during a baseball game, you'll be able to use the Net to get detailed statistics about the batter, and even past videos and highlights of his career.

Television and the Internet have already begun to merge in a very real way, primarily through cable modems, TVs that connect to the Net (enabling you to browse the Net through your TV set, such as with WebTV), and interactive TV (which uses the Internet to deliver information to people on their TV screens).

Cable modems offer extremely high-speed access to the Internet. They enable you to access the Net using the existing coaxial TV cable that comes into your house. Cable modems can receive data at up to 1.5 million bits per second and can send data at up to 300,000bps—far faster than normal modems and ISDN lines. They offer T1-like speeds, but at a fraction of the cost of T1 lines. They can deliver these high speeds because they are sent through high-capacity cable lines. Because Internet data and the normal cable signal coexist on the same lines, you can access the Net and watch TV simultaneously. The illustration in this chapter describes how MediaOne provides cable modem access to its customers.

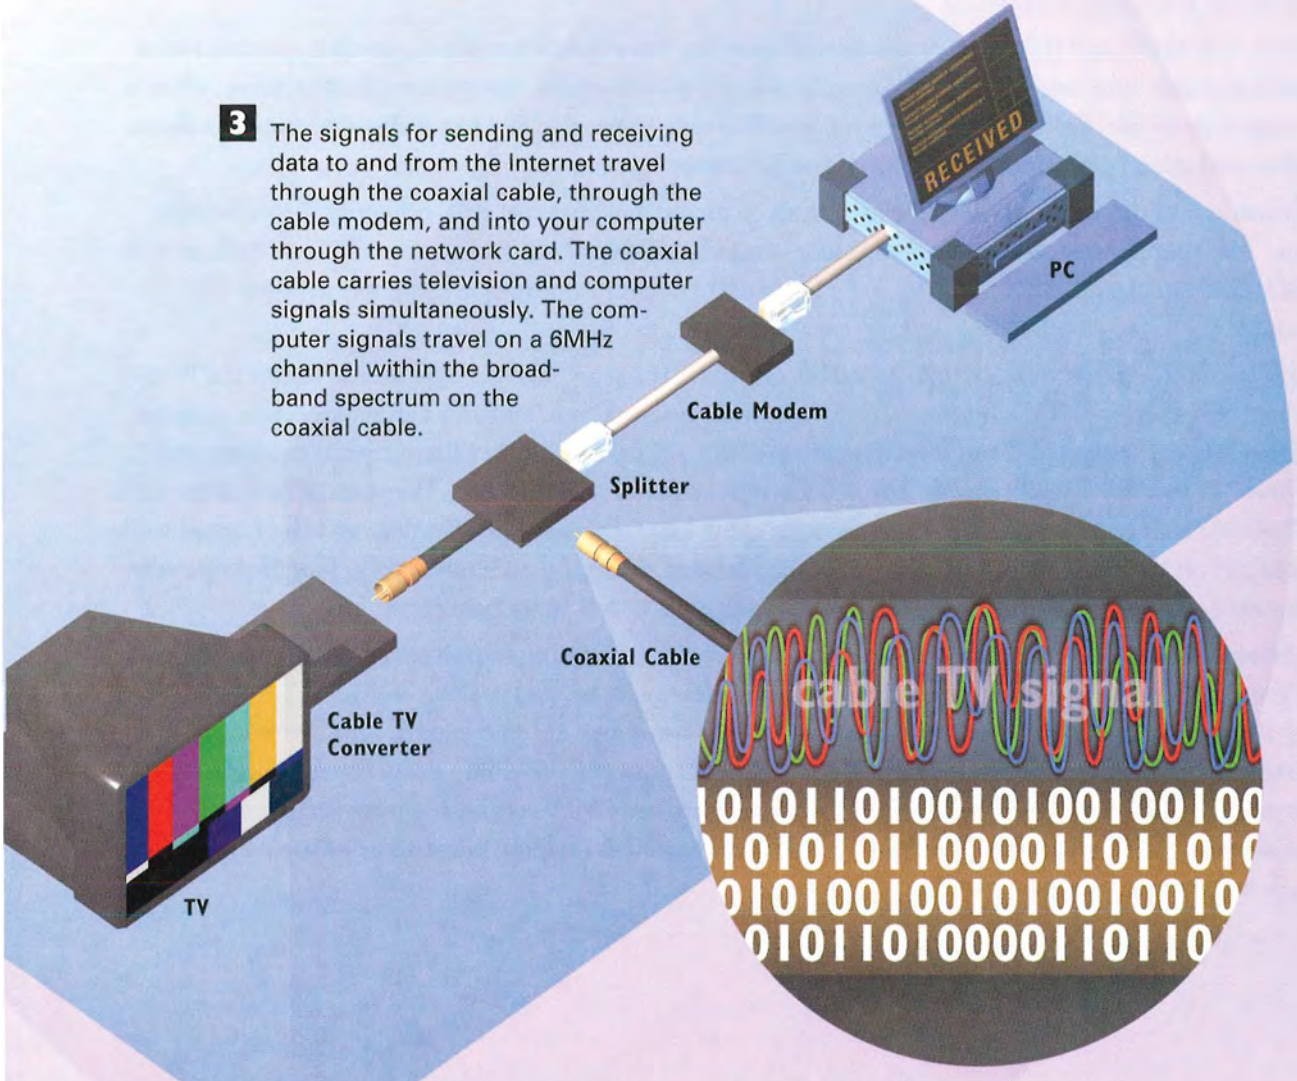
A different technology enables you to browse the Web using your television set. A *set-top box* connects your TV to the Net via a modem, takes the signal from the modems, and sends it to the TV. A remote control–like device enables you to browse the Web while watching TV at the same time. Interactive TV goes a step further and uses a set-top box to enable interactivity with your PC, using Internet technologies. At this point, no single standard exists for how Internet-enhanced TV works, nor even an agreement on which features such a service should have. The illustration in this chapter shows some of the common ways that such a service will work.

How Cable Modems Work

1 The coaxial cable—sometimes called a *broadband wire*—that comes into a home from a pole is divided into two connections by a splitter inside the home. One part of the wire goes to the normal set-top box that provides cable TV access. The other part of the wire goes to a cable modem, sometimes called a *broadband modem*.

2 The cable modem attaches to an ethernet network card inside the computer. The network card is configured the same as any other network card on a computer that is attached to the Internet and has a network address.

3 The signals for sending and receiving data to and from the Internet travel through the coaxial cable, through the cable modem, and into your computer through the network card. The coaxial cable carries television and computer signals simultaneously. The computer signals travel on a 6MHz channel within the broadband spectrum on the coaxial cable.

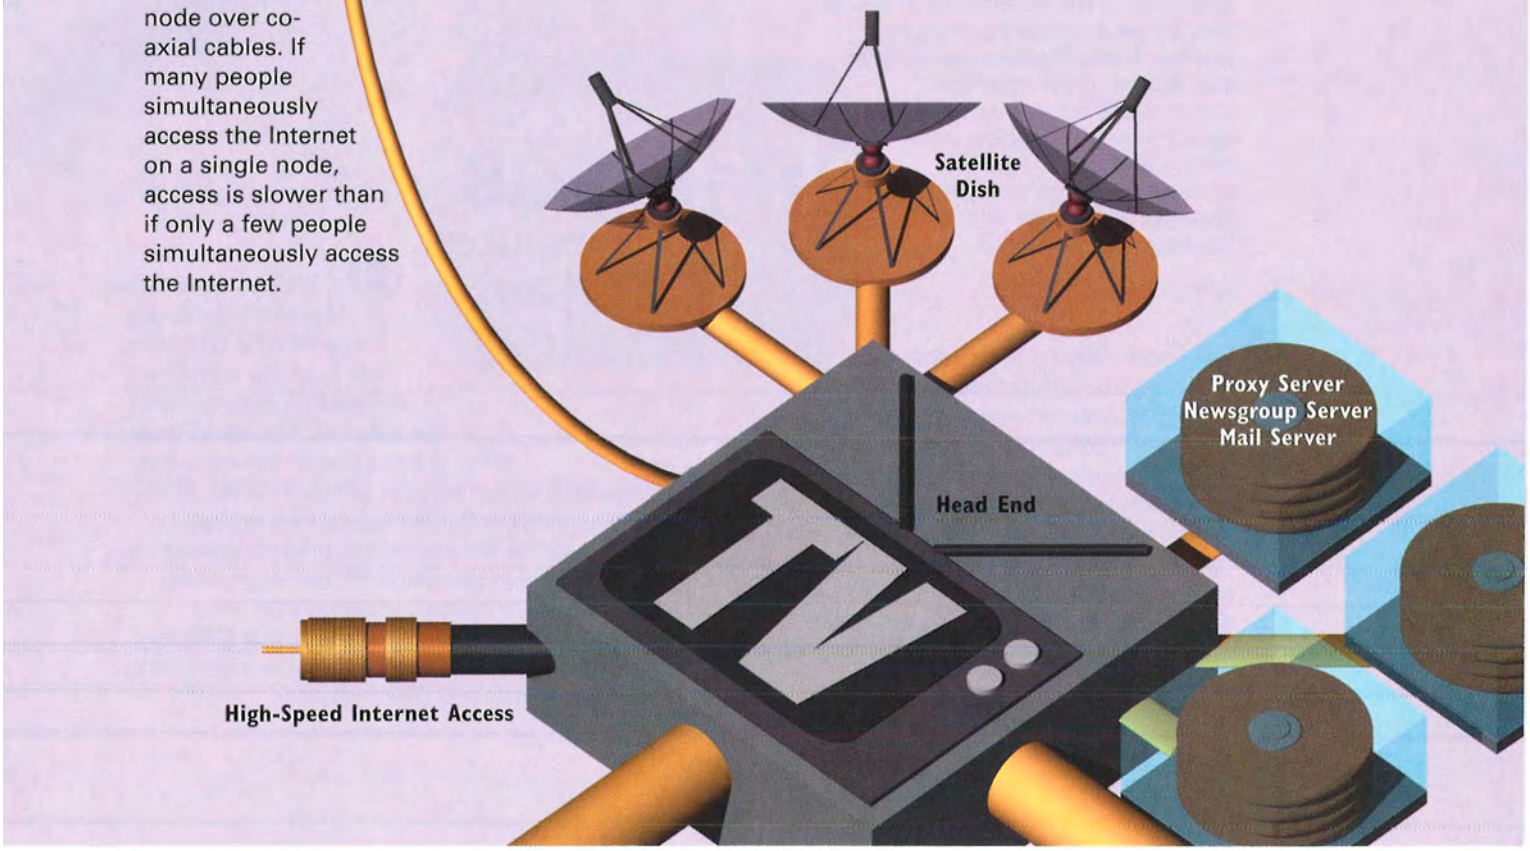


4 The cable company divides each town or city into neighborhoods of about 500 homes, all of which are located on a single local area network, or *node*. All of those 500 homes share the node. Both television and Internet data travel to and from those 500 homes to the node over coaxial cables. If many people simultaneously access the Internet on a single node, access is slower than if only a few people simultaneously access the Internet.

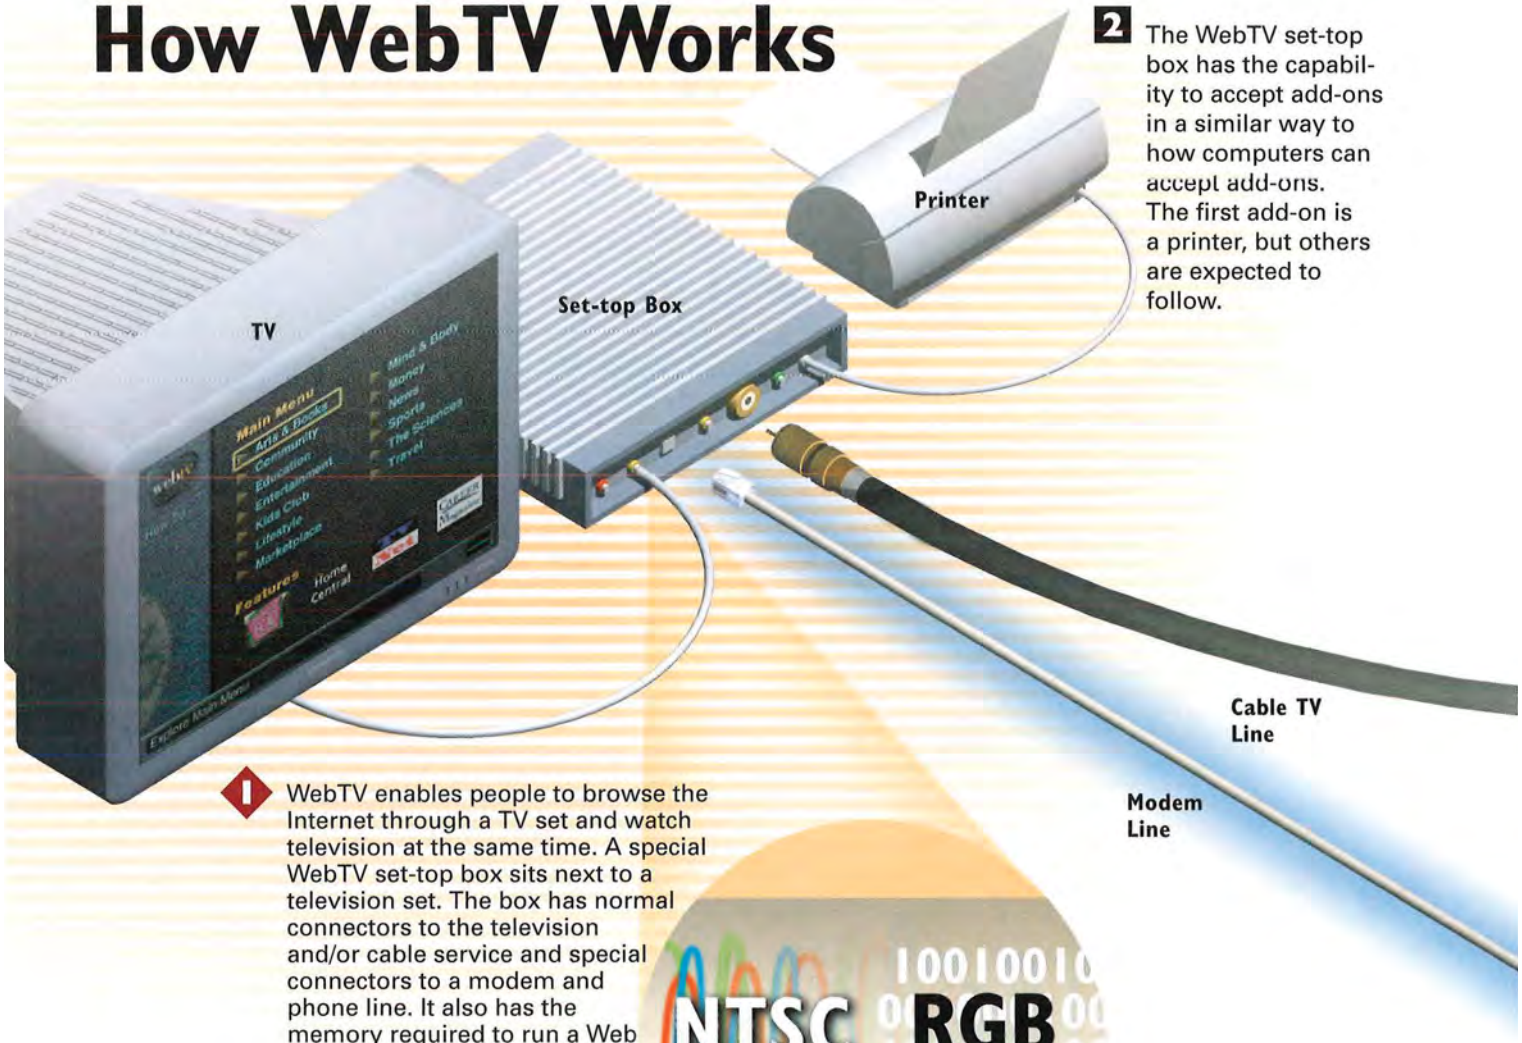
5 The nodes are connected via high-speed fiber-optic lines to a head-end cable facility. A single head end typically handles all the nodes for 4–10 towns. The head end is responsible for delivering television programming and Internet access to the cable customers.

6 The head end receives television transmissions from satellites and has Internet access via high-speed links to the Internet. These feeds provide the cable programming and Internet access to cable customers.

7 The head end also has high-speed Internet servers. A news server offers access to Internet Usenet newsgroups at a high speed because people access the server via the high-speed fiber-optic wires and coaxial cables instead of having to go out over the slower Internet. Additionally, *proxy servers* cache in their memory the most current versions of the more frequently accessed sites on the Internet. In that way, cable modem customers can get high-speed access to the sites because they get them via high-speed fiber-optic wires and coaxial cables instead of having to go out over the slower Internet.

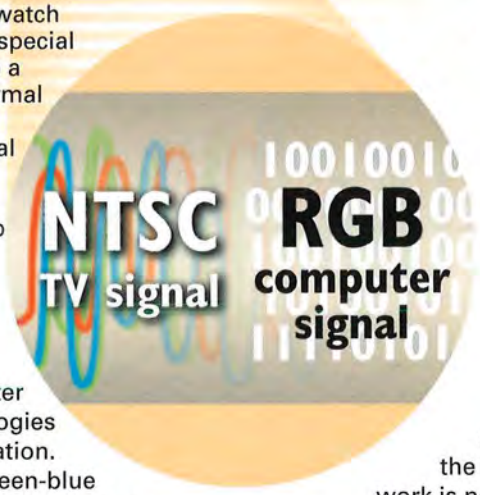


How WebTV Works



1 WebTV enables people to browse the Internet through a TV set and watch television at the same time. A special WebTV set-top box sits next to a television set. The box has normal connectors to the television and/or cable service and special connectors to a modem and phone line. It also has the memory required to run a Web browser along with other hardware and software.

2 The WebTV set-top box has the capability to accept add-ons in a similar way to how computers can accept add-ons. The first add-on is a printer, but others are expected to follow.



3 Television screens and computer monitors use different technologies to display pictures and information. Computer monitors use red-green-blue (RGB) technology, in which the three colors are combined to form all the colors you see on your monitor. Televisions use National Television Standards Committee (NTSC) technology. The WebTV set-top box converts the signal from the RGB of the Web to the NTSC used by televisions so the signal is as clear as possible.

4 When you want to access the Internet, the WebTV box connects to the WebTV network to get access to the Internet. The WebTV network is not a single network, but instead a virtual private network (VPN). WebTV makes deals with Internet service providers (ISPs) across the country to provide access via local points of presence (POPs); encryption technology knits it all together into what appears to be a single network. This VPN has many Web servers on it to provide capabilities such as e-mail and other services.

Remote





7 The ability to browse the Web is controlled by a simple remote control unit that enables people to scroll, move around the screen, and click objects. The control also allows people to browse the Web, watch television, or view the Web and television simultaneously.

6 WebTV uses its own proprietary Web browser, which adheres to HTML standards and can run common plug-ins. The browser is designed for optimal display on a television set. Some sites, such as *TV Guide*'s, have developed Web sites that make the best use of the browser.

5 To speed up access to Internet sites, the WebTV network runs a series of proxy servers. These proxy servers cache commonly accessed Web pages. When someone who uses WebTV goes to one of those pages, the page is delivered from the proxy server, which delivers the page at a higher speed than if it were accessed from the Internet.

How Internet-Enhanced TV Works

1 To receive Internet-Enhanced TV, a television set requires a special set-top box that can receive and deliver the Internet-enhanced services. The set-top box has a powerful processor and a hard disk inside it.





2 Internet-Enhanced TV can work with many types of TV connections, including cable connections, satellite connections, and possibly DSL connections as well. The connection, such as the cable connection shown here, connects directly to the set-top box, not to the television itself.

3 When a TV station sends out a broadcast, it contains the normal TV broadcast, as well as an accompanying enhanced TV signal with Internet-related information. Televisions without the special set-top box display the normal TV broadcast and ignore the enhanced signal.

4 The enhanced signal can contain a variety of extra information, and it can be sent using the Hypertext Markup Language (HTML), the language of the Web. For example, the information might contain statistics about a batter during a baseball broadcast, or it could allow viewers to answer questions during a quiz show. The HTML pages also could contain TV-specific HTML commands that can be displayed only via set-top boxes and not with Web browsers.

5 The set-top box interprets the HTML information and formats it in a way that can be displayed on a TV set. This information can be overlaid on the regular TV show so that you can see the information at the same time that you're watching the show.

6 Using a special remote control, the viewer can interact with the television—for example, by clicking a link to get baseball statistics or taking a poll during a news show. When a viewer clicks n a link, the link behaves as does any other on a Web page and delivers the information or page requested. The request goes back through the set-top box, is sent to the Web page, and is then delivered via the enhanced TV signal back to the set-top box.

CHAPTER 11

How the Digital Subscriber Line (DSL) Works



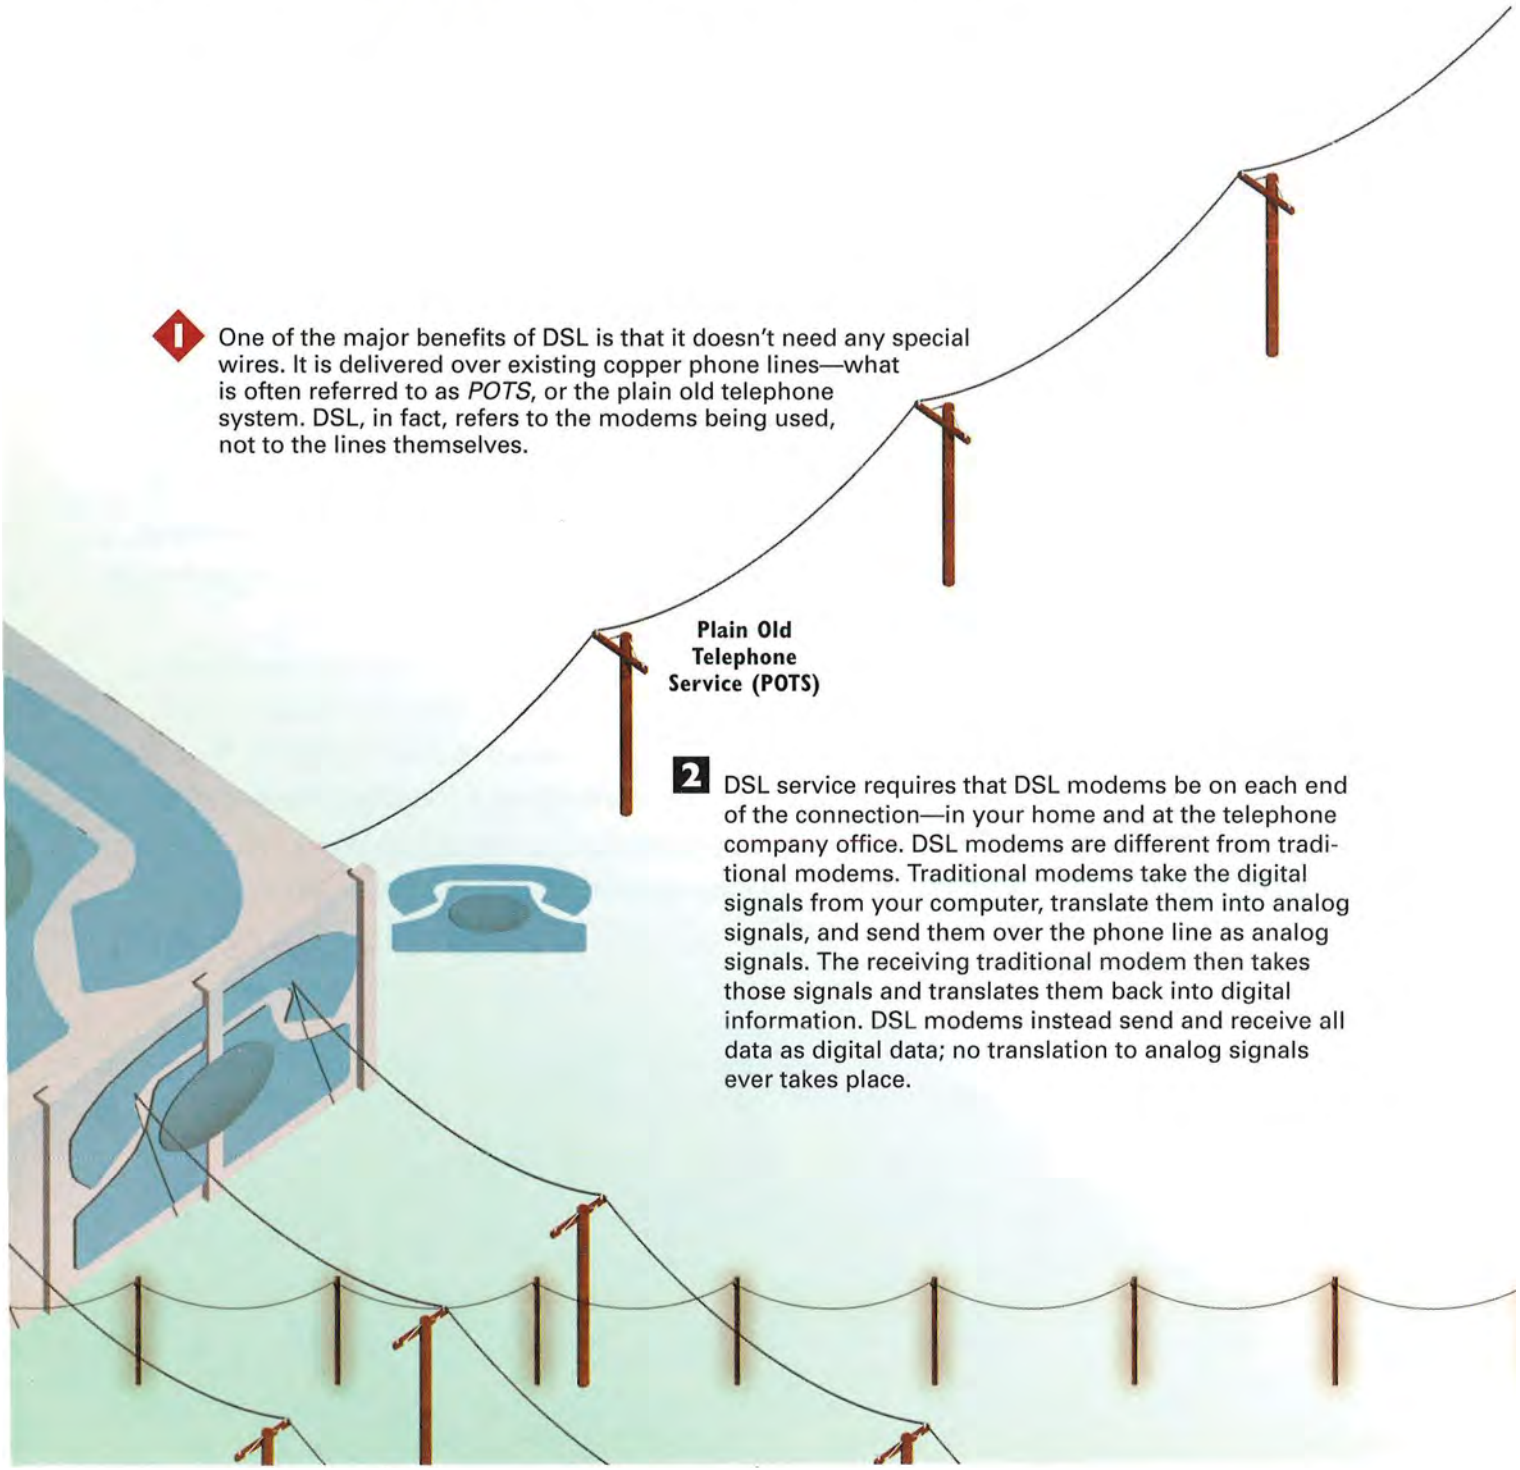
IN recent years, a plethora of options could be found for gaining high-speed access to the Internet. Among the newest of those options is the Digital Subscriber Line (DSL), which might eventually prove to be one of the most popular. Several types of DSL technologies are available, but they all work on the same principles. They enable you to use your existing telephone lines to access the Internet at very high speeds—potentially as high as 55 million bits per second (Mbps)—using VDSL (Very high data rate DSL). However, that kind of extremely high-speed connection is used primarily by large corporations because of its cost. Home users will probably get speeds of 1.5Mbps, which is still very fast—about 50 times the speed of a 28,800-bits-per-second (bps) modem. Additionally, DSL technologies enable you to use the Internet at high speeds and simultaneously talk on your telephone, using the same, single phone line.

DSL technologies require that DSL modems be used on each end of the phone line. In fact, the term *DSL* doesn't really refer to a phone line because an ordinary, existing copper phone line can be used for DSL. Instead, it refers to the DSL modems themselves. More confusing still, DSL modems aren't really modems at all. Traditional modems modulate the digital signals from your computer into analog signals so that they can be sent over telephone lines, and then they demodulate them back into digital signals so your computer can understand them—hence the name modem, from the terms MOdulate/DEModulate. DSL modems don't convert signals from digital to analog. Instead, they transmit and receive all data as digital signals. Because the signals don't need to be converted and can be sent as digital data instead of analog data, data can be sent at much higher speeds than data sent via traditional modems.

DSL technology has one drawback; it requires that your house (and DSL modem) be located within a certain distance from the telephone company office and its DSL modem. In cities, this should not be a major problem, but it could be a problem in rural areas. The exact distance required depends on the type and speed of DSL service you use. Higher speeds require that you be closer to the phone company office.

Although many types of DSL service are available, probably the most popular for people at home is standard DSL service. With it, you can receive data at a faster rate than you can send data—for example, you can receive data at 1.5Mbps and send data at 640Kbps.

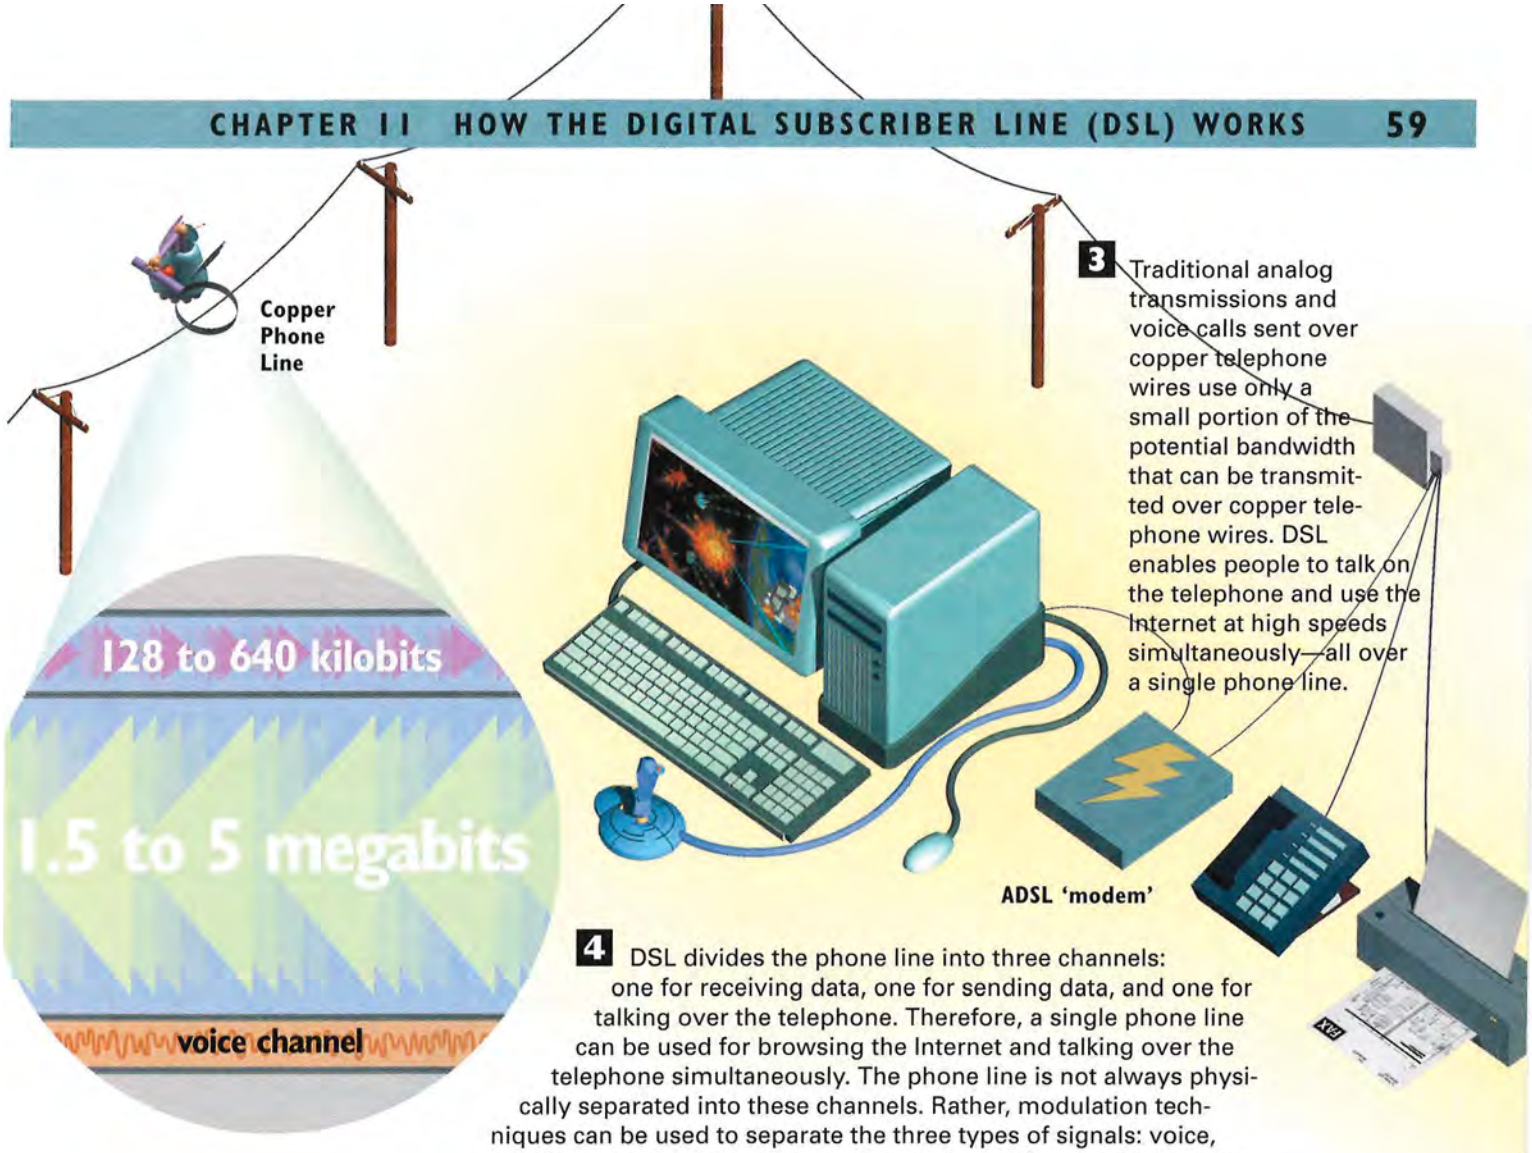
How DSL Makes Its High-Speed Connection



1 One of the major benefits of DSL is that it doesn't need any special wires. It is delivered over existing copper phone lines—what is often referred to as *POTS*, or the plain old telephone system. DSL, in fact, refers to the modems being used, not to the lines themselves.

Plain Old Telephone Service (POTS)

2 DSL service requires that DSL modems be on each end of the connection—in your home and at the telephone company office. DSL modems are different from traditional modems. Traditional modems take the digital signals from your computer, translate them into analog signals, and send them over the phone line as analog signals. The receiving traditional modem then takes those signals and translates them back into digital information. DSL modems instead send and receive all data as digital data; no translation to analog signals ever takes place.



3 Traditional analog transmissions and voice calls sent over copper telephone wires use only a small portion of the potential bandwidth that can be transmitted over copper telephone wires. DSL enables people to talk on the telephone and use the Internet at high speeds simultaneously—all over a single phone line.

4 DSL divides the phone line into three channels: one for receiving data, one for sending data, and one for talking over the telephone. Therefore, a single phone line can be used for browsing the Internet and talking over the telephone simultaneously. The phone line is not always physically separated into these channels. Rather, modulation techniques can be used to separate the three types of signals: voice, sending, and receiving. The sending and receiving channels can be divided into a variety of speeds. One form of ADSL might be able to receive data at 1.5Mbps, for example, and send it at 640Kbps.

5 To work properly, your DSL modem must be located within a certain distance from the phone company's answering DSL modem. The exact distance varies according to the precise DSL service and speed being offered, and even according to the gauge of the copper telephone wire. For an 8.448Mbps service, for example, the phone company's DSL modem needs to be within 9,000 feet of your modem. At 2.048Mbps, on the other hand, the modem can be located 16,000 feet away.



CHAPTER 12

How ISDN Works



ONE common complaint about the Internet is that ordinary telephone connections are too slow. Even "high-speed" modems that can connect at 28,800Kbps can be too slow to take advantage of the rich graphics now available. The Internet is full of graphics, videos, sounds, and other multimedia files that are very large and take a very long time to be sent across the Internet to your computer.

One affordable way to get a faster Internet connection is to use ISDN (Integrated Services Digital Network). It's a technology that has been around for years, but the advent of the Internet and the need for higher-speed connections has made it more popular.

ISDN enables high-speed connections to the Internet using existing copper cable telephone wires. It isn't available to everyone in the country yet because it requires telephone companies to install special ISDN digital-switching equipment. However, every day more areas of the country offer ISDN service. If ISDN is available in your area, you'll pay extra for the service compared to a normal telephone line. To use ISDN and the Internet, the number you're dialing must be equipped for ISDN access. Many private Internet dial-in providers and online services allow for ISDN access.

ISDN also requires that you have an ISDN modem. Although it looks like a modem, an ISDN modem isn't really a modem at all. Instead, it's a *terminal adapter*, a piece of hardware that lets you send and receive digital signals over ISDN phone lines. Normal modems transform digital signals from your computer into analog signals that can be sent over normal analog telephone lines. Because ISDN is a digital technology, only digital data is sent by the ISDN adapter. Some ISDN devices include the capability to function as a regular modem. This capability is necessary because not every place you can dial into with your computer lets you use ISDN; some services presently allow only regular modem access.

ISDN enables you to use your line for voice calls as well as for sending and receiving data, and in fact, depending on where you get the service, it might support caller ID so you can see the phone number of people who are calling. In some setups, you can even have more than one telephone number sharing a single line. Then, depending on the call, it can be routed to different devices sharing the same ISDN line. For example, you can have a fax number, and all calls to that number automatically will go to your fax machine.

Although ISDN offers higher connection speeds than traditional modems, it is slower than DSL and cable modems. Because of this, and because there have been problems rolling the service out across the country, it hasn't been as popular as its higher-speed DSL and cable modem competitors. Still, when DSL and cable modems aren't available, or in areas where ISDN service is less expensive, it's a good way to get onto the Internet at higher speeds than traditional modems.

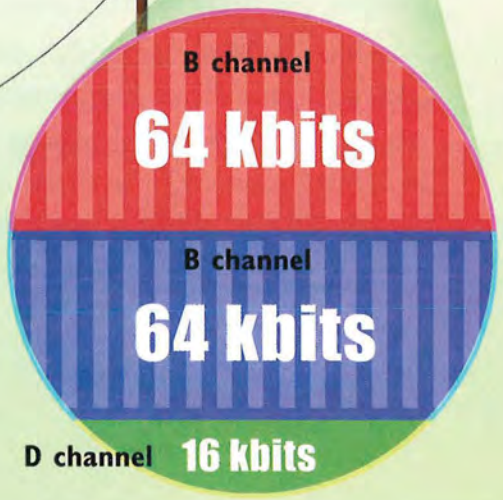
How ISDN Makes Its High-Speed Connection

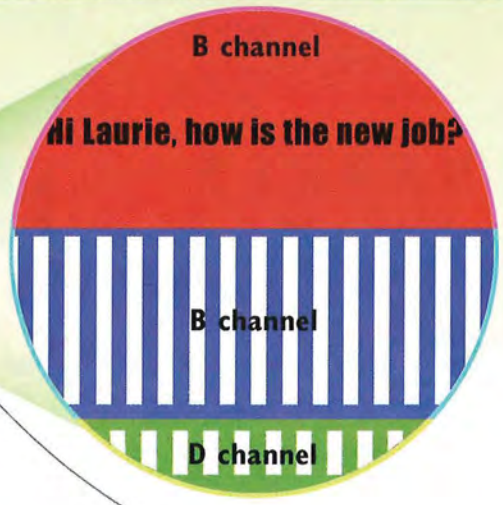
1 To use ISDN service, a computer requires an ISDN modem. In fact, this "modem" isn't really a modem at all and should be called a terminal adapter instead. The ISDN "modem" sends digital information from your computer over ISDN telephone lines and receives digital information from telephone lines.



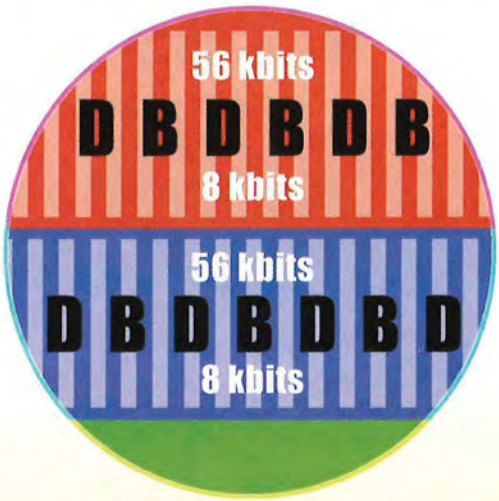
2 Unlike normal telephone lines, ISDN lines do not receive power from the telephone company. Instead, they must use power from an outside source, so they plug into a power outlet. Because of this, in the event of a blackout or power failure, ISDN lines will not work. Normal telephones usually work during blackouts because they receive their power directly over the phone line.

3 A number of variations of ISDN service exist, but by far the most common is known as Basic Rate Interface (BRI). With BRI, your telephone line is divided into three logical channels. These logical channels are the means by which data is sent and received over your telephone line. BRI has two 64Kbps B (bearer) channels and one 16Kbps D (data) channel. This is commonly referred to as 2B+D.





4 The B channels are used to send your data and can be used simultaneously. So, for example, you can be talking on the telephone through one B channel while cruising the Internet on your computer over the other B channel. The D channel is used to send signaling information for routing the data that is being sent over the B channels.



Out-of-band Signaling



In-band Signaling

5 Not all telephone companies have the capability to use the D channel to send signaling information. In cases where the D channel can't be used for signaling information, that information is instead delivered over the B channels. In this instance, each B channel can send and receive data at only 56Kbps because the other 8Kb in each channel is used for signaling information. When the B channels carry signaling information, it is referred to as *out-of-band signaling*. When the D channel carries signaling information, it is referred to as *in-band signaling*.

6 When accessing the Internet, ISDN generally uses the Point-to-Point Protocol (PPP). However, with the Multilink PPP (MP) protocol, the two B channels can be combined into a single channel for a speed of 128Kbps when in-band signaling is used and 112Kbps when out-of-band signaling is used.

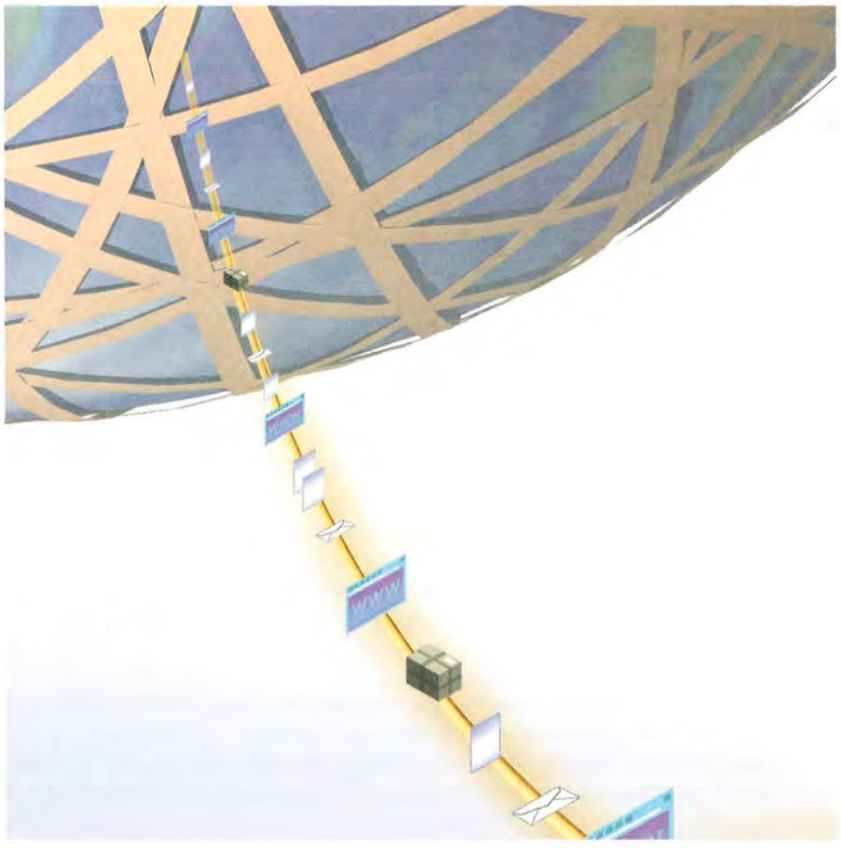


Multilink PPP

CHAPTER

13

How You Connect to the Internet from Online Services



THERE are many ways to get access to the Internet; one is by using an online service such as CompuServe or America Online. These online services have long provided their own unique content, special areas, and services available only to their subscribers. They use proprietary software and interfaces to give subscribers access to their resources. Unlike most of the Internet, the content, areas, and services the online companies provide are not always free. To get them, you might have to pay a monthly subscription fee to the online service. These online services are different from Internet service providers (ISPs) that offer only access to the Internet and don't have their own private areas and services.

The services let you access the Internet in a number of ways. To make using the Internet's resources easier, the services often use their own proprietary software or user interface. The user interface requires you to issue a command that is then sent out over a gateway to the Internet. (A *gateway* is a device that sends data from network to network and can translate between the two if they use different kinds of data and protocols.) The information you've asked for is retrieved, sent back over the gateway to your online service, and is displayed for you using the online service's proprietary software.

Most online services also allow you to use your own client software, such as FTP software or newsgroup readers, to use the Internet. To do this, you essentially bypass the online service's proprietary interface. You use the online service as you would a dial-in ISP. First, you dial into the online service to establish a TCP/IP connection. Then, you launch the client software on your computer. It is the TCP/IP connection that gives you access to the Internet's resources.

Online services also let you browse the Web. Some let you use only their own proprietary Web browsers. Others have their own proprietary Web browsers but also allow you to browse the Web using any Web browser of your choice. To use your own Web browser, you first dial into the service to establish a TCP/IP connection. Then, you launch your own Web browser, and you can browse the Web just as you would with any other dial-in ISP.

In addition to letting you browse the Web, online services also let you use other Internet resources, such as Telnet (which lets you gain access to a distant computer on the Internet) and Usenet newsgroups (public Internet discussion areas). In some instances, you gain access to those resources using the online service's proprietary software, whereas in other instances you use normal Internet client software.

Increasingly, the online services are moving toward eliminating their proprietary architecture altogether. The trend is to build the service on the same framework as the Internet, such as the TCP/IP protocol and related technologies. The Microsoft Network, for example, is now based entirely on TCP/IP. America Online and CompuServe are both incorporating elements of Internet technology inside their services and one day might go entirely to an Internet-based model.

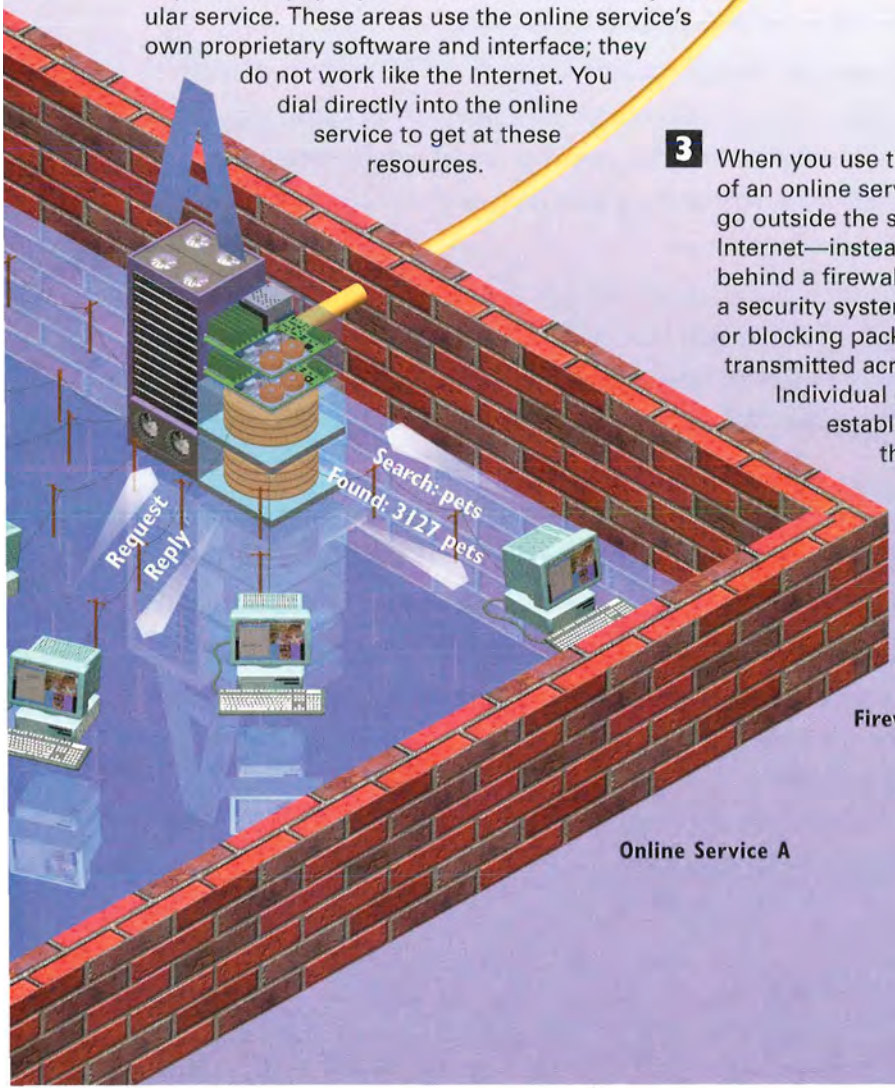
How Online Services Connect to the Internet

1 Online services such as CompuServe and America Online provide a very convenient way to use Internet resources. These services have their own software that makes accessing Internet resources such as Telnet and FTP easy. Some also allow you to use your own Telnet, FTP, Gopher, or other software instead of theirs. Using these resources, you can browse the World Wide Web, either with their own proprietary Web browser or by using any other Web browser. Finally, some services let you use special software required to access Internet resources such as Internet Relay Chat (IRC).

2 Each online service has a great deal of resources, content, and special areas that are not available to anyone except people who subscribe to that particular service. These areas use the online service's own proprietary software and interface; they do not work like the Internet. You dial directly into the online service to get at these resources.

3 When you use the resources of an online service, you don't go outside the service to the Internet—instead, you stay behind a firewall. A *firewall* is a security system of accepting or blocking packets as they are transmitted across a network. Individual online services establish firewalls that let only subscribers into the service.

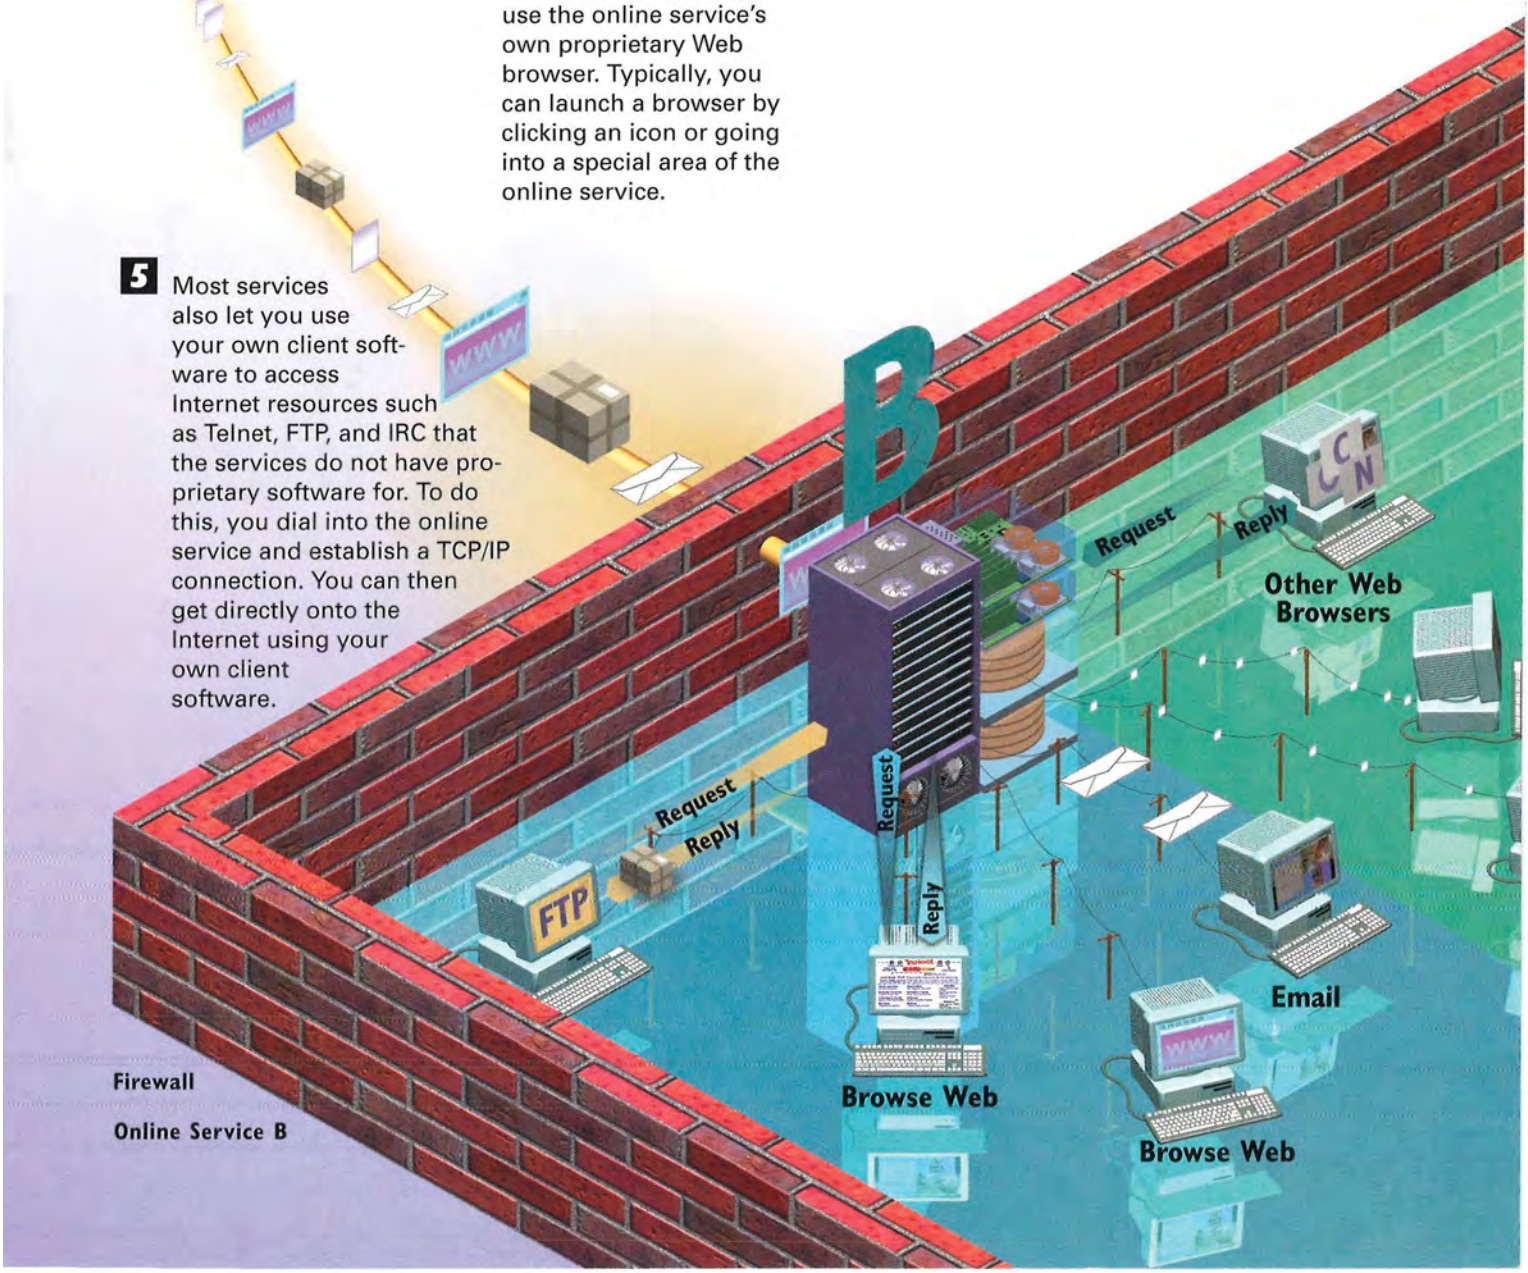
4 Online services use their own proprietary software and interfaces to give you access to many common Internet resources, such as Usenet newsgroups, Telnet, and FTP. You issue a command using the online service's own software. The online service then sends the command out over the Internet via a gateway to retrieve the information. The information is sent back via a gateway, through the online service's own software, and then to your computer.



7 Some online services let you use any Web browser you want to browse the World Wide Web. In that instance, you can click an icon on the online service, go into a special area of the online service, or simply launch your Web browser on your own computer, and then begin browsing the Web.

6 Online services also let you browse the World Wide Web. One way to browse the Web is to use the online service's own proprietary Web browser. Typically, you can launch a browser by clicking an icon or going into a special area of the online service.

5 Most services also let you use your own client software to access Internet resources such as Telnet, FTP, and IRC that the services do not have proprietary software for. To do this, you dial into the online service and establish a TCP/IP connection. You can then get directly onto the Internet using your own client software.



CHAPTER

14

How Wireless and Palmtop Connections Work



EVERY day, new, unique ways of connecting to the Internet and browsing the Web are being developed. While traditional PCs and modems are still the dominant way people access the Internet, increasingly many other types of devices can access the Web as well—and hardwares other than traditional modems are also becoming more and more popular ways to get onto the Internet, particularly wireless technologies.

Among the newest ways to connect to the Internet and browse the Web are via satellite connections; via Palmtop computers, such as the Palm and Windows CE devices; via cellular telephones; and via a wireless technology called Bluetooth.

When you connect to the Internet using a satellite connection, you can get access at a higher rate than normal 28.8 or 56 kilobits per second modems. You typically get data and Web sites transferred to you at 400 kilobits per second, more than seven times the speed of a 56Kbps modem.

Satellite-access systems are, in a way, odd hybrids. They still require that you use a modem. The modem is used to request information from the Internet, so when you send information through the Internet, it goes at normal modem speeds. However, that requested information is sent back to you at the high satellite transmission speeds, not the slower modem speeds. However, this means that if you're sending e-mail, or transferring a file to someone via FTP, that information is sent at modem speeds, not satellite speeds.

When you access the Internet via satellite, you still use a traditional PC. However, you can also get onto the Web in other ways using new types of digital devices, most notably Palmtop devices. These little computers, sometimes called personal digital assistants (PDAs) can literally put the Web in the palm of your hand. You connect a modem to them and can then get onto the Internet and send and receive e-mail and browse the Web with them.

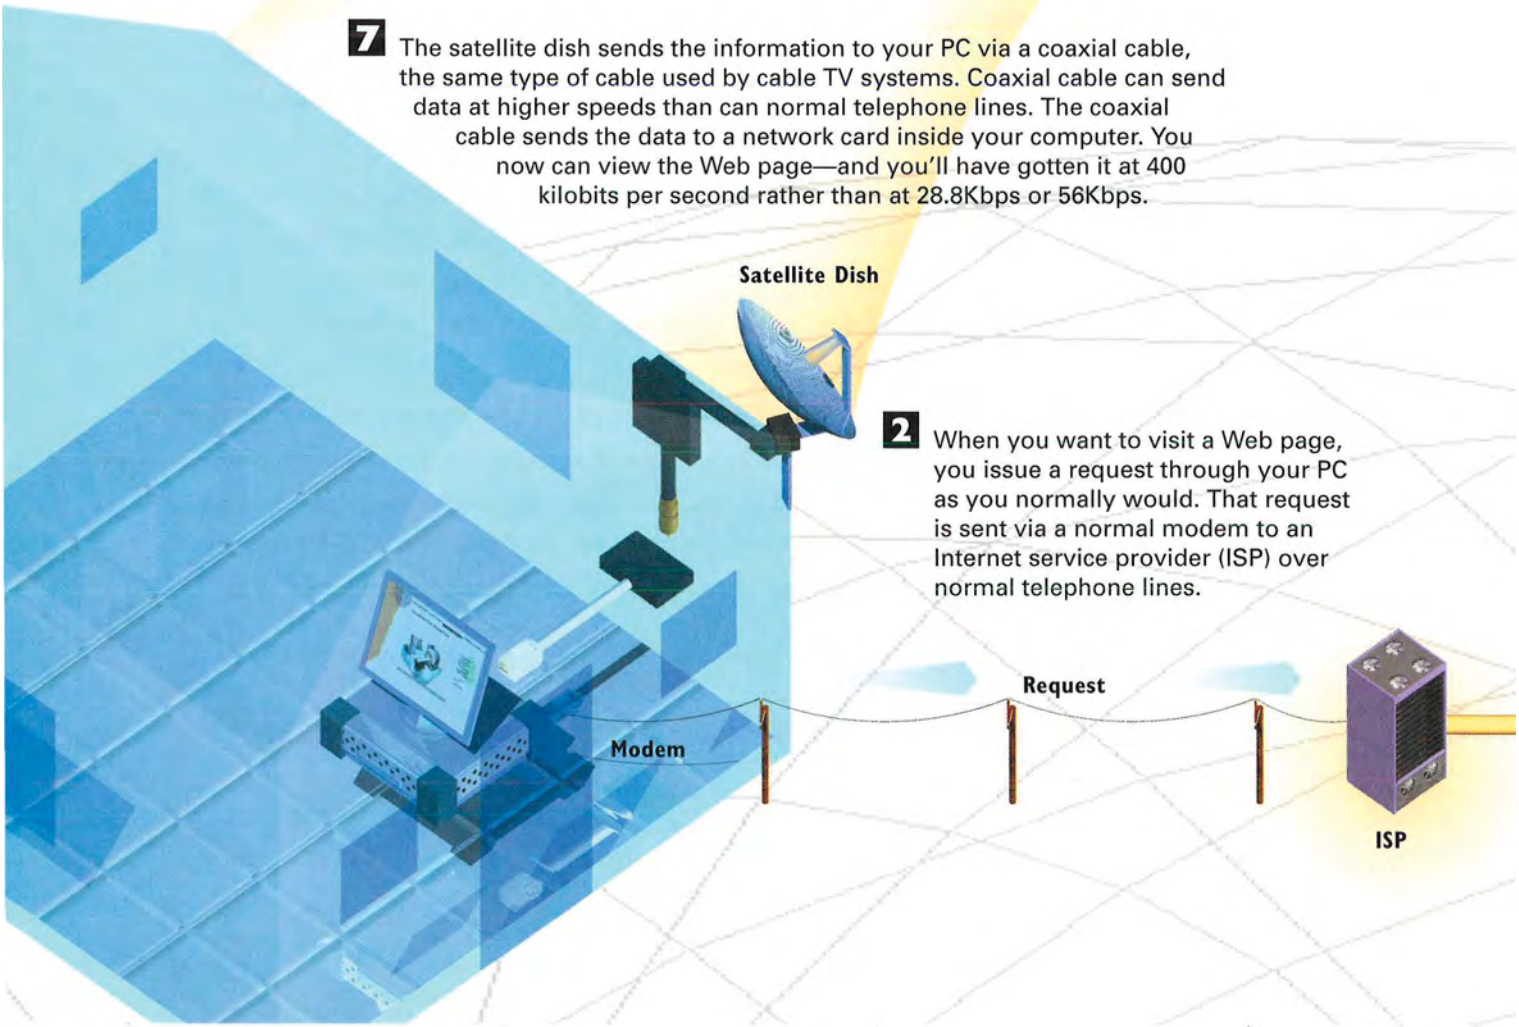
Some PDAs and cellular telephones can access the Internet without using normal telephone lines (called *landlines*). They use cellular technology to make a connection to the Internet. But because their screens are small and typically can't display graphics, Web pages must be reformatted into a format called Wireless Markup Language (WML) to be viewed properly on them. Some Web pages are built with WML, and those Web pages need not be specially formatted—they can be viewed as is.

How To Access the Internet via a Satellite Connection

1 Accessing the Internet via a satellite connection delivers information to your computer at much higher speeds than do normal modems—at 400 kilobits per second. To use it, you need a satellite dish on your home in the same way that you need a satellite dish to get satellite TV.

6 The satellite in turn sends the information at 400Kbps to the satellite dish in your home.

7 The satellite dish sends the information to your PC via a coaxial cable, the same type of cable used by cable TV systems. Coaxial cable can send data at higher speeds than can normal telephone lines. The coaxial cable sends the data to a network card inside your computer. You now can view the Web page—and you'll have gotten it at 400 kilobits per second rather than at 28.8Kbps or 56Kbps.



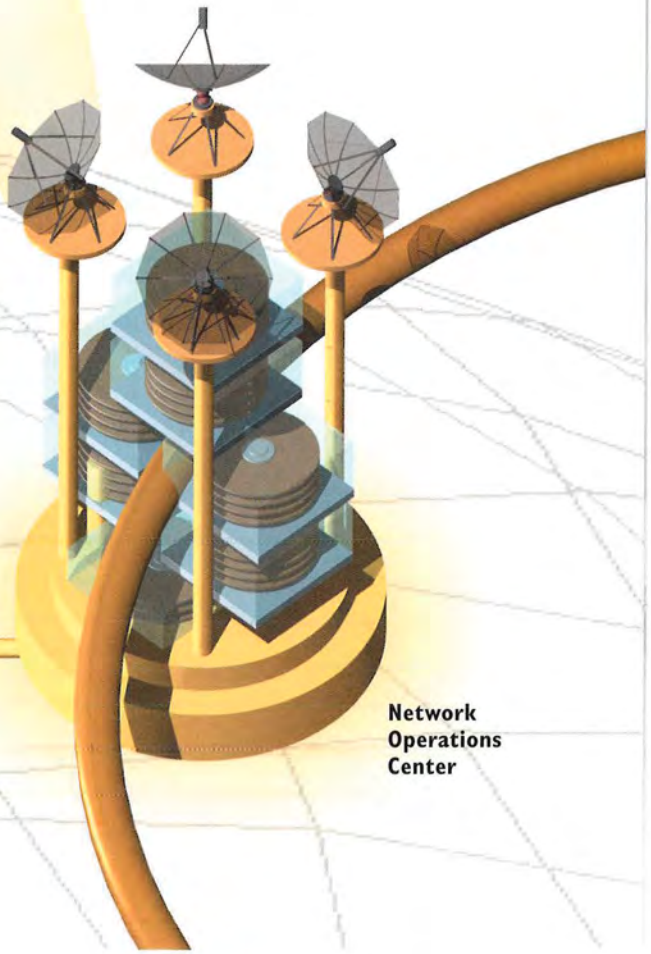
2 When you want to visit a Web page, you issue a request through your PC as you normally would. That request is sent via a normal modem to an Internet service provider (ISP) over normal telephone lines.



Satellite

5 The NOC sends the information to a satellite above the earth at 400Kbps.

4 Instead of transferring the information directly back to you over telephone lines, the Web server sends the information to the satellite company's network operations center (NOC) (also called a satellite ground station) via special high-speed links.



Network Operations Center

3 Your request to visit the page goes to the Web server that houses the site you want to visit.



Web Server

How Palmtop Computers Can Browse the Web



1 Palmtop computers such as 3Com's Palm series are small, handheld devices, whose small screens wouldn't seem to be able to allow you to browse the Web. But even these small devices coupled with special software, such as ProxiWeb, can be used for Web browsing. You install this software by first putting it on your PC and then transferring it over to the Palm via HotSync technology that transfers files and data between PCs and Palms.

2 After the browsing software is installed, you connect to the Internet using a modem and dialing into an Internet service provider. The modem can be a special one devised for the Palm, or it can be a PC modem.







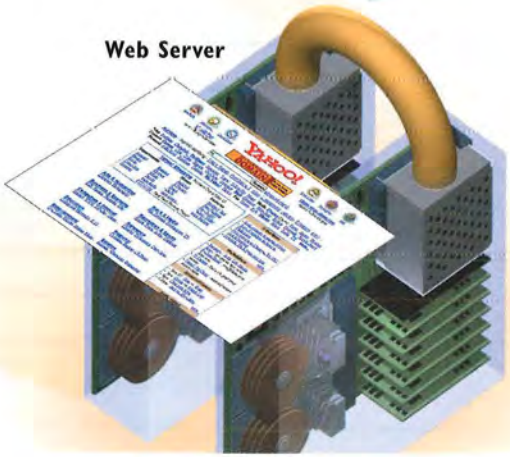
3 After you're connected to the Internet, when you want to visit a Web site you type its location into the browser. Instead of going to the site, though, the ProxiWeb browser contacts a special *proxy server* located on the Internet.

4 The proxy server contacts the Web server that hosts the page you want to visit. The page downloads to the proxy server. As the page is downloaded, the proxy server processes and reformats it, putting into a form that can be easily read by Palm device's small screen. It strips out the HTML coding and changes it into a special format the ProxiWeb browser on the Palm can read, and it changes the graphics into a format that can be read on the Palm as well.



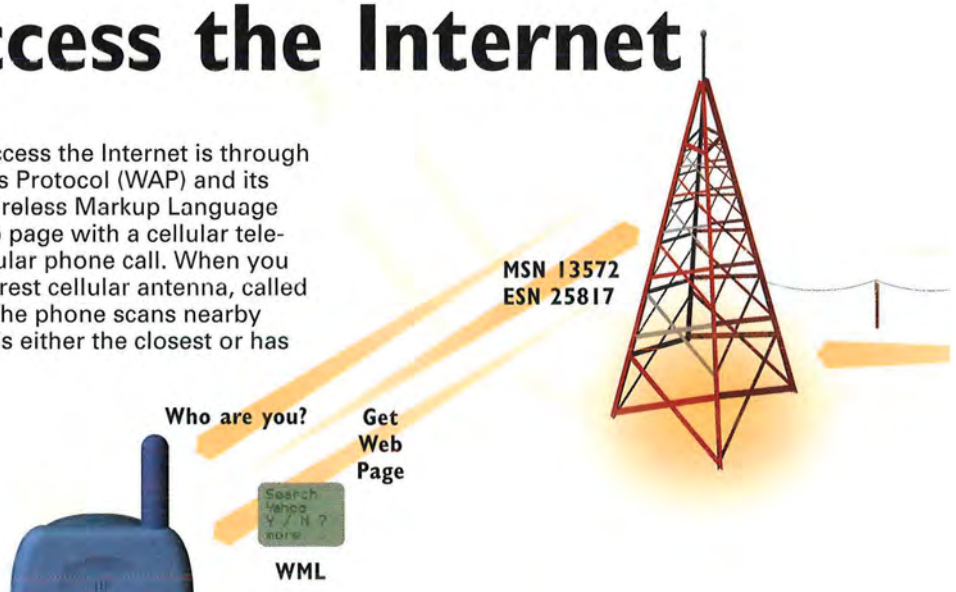
5 After the page is processed and reformatted, it is sent to the Palm device, where it can be read. The links are live just as if it were a normal Web page, so you can click them to view them. When you click a link, you contact the proxy server, which then gets you that page in the same way it sent you the initial page.





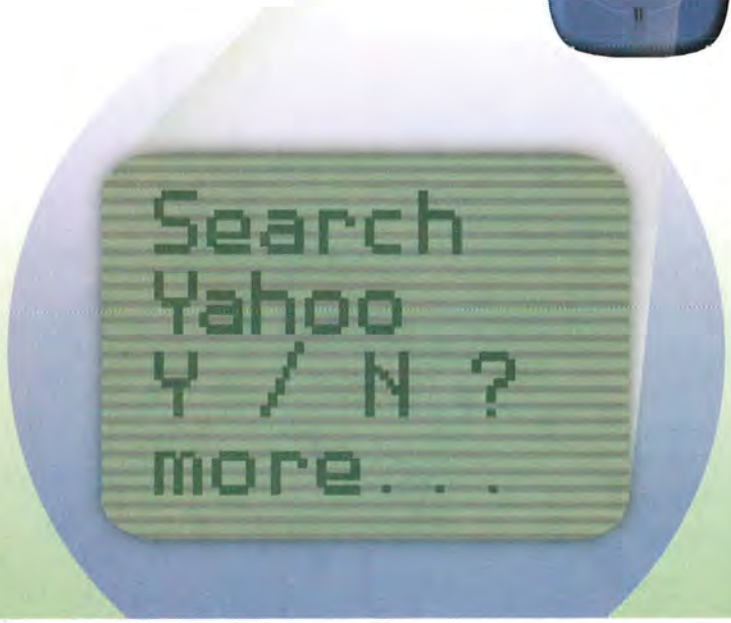
How Cellular Telephones and Palmtops Access the Internet

1 The main way cellular telephones access the Internet is through a protocol called the Wireless Access Protocol (WAP) and its associated markup language, the Wireless Markup Language (WML). To use WAP to access a Web page with a cellular telephone, you first need to make a cellular phone call. When you dial a cell phone, it looks for the nearest cellular antenna, called a *base station*, to transmit the call. The phone scans nearby base stations and locks on one that is either the closest or has the most powerful signal.

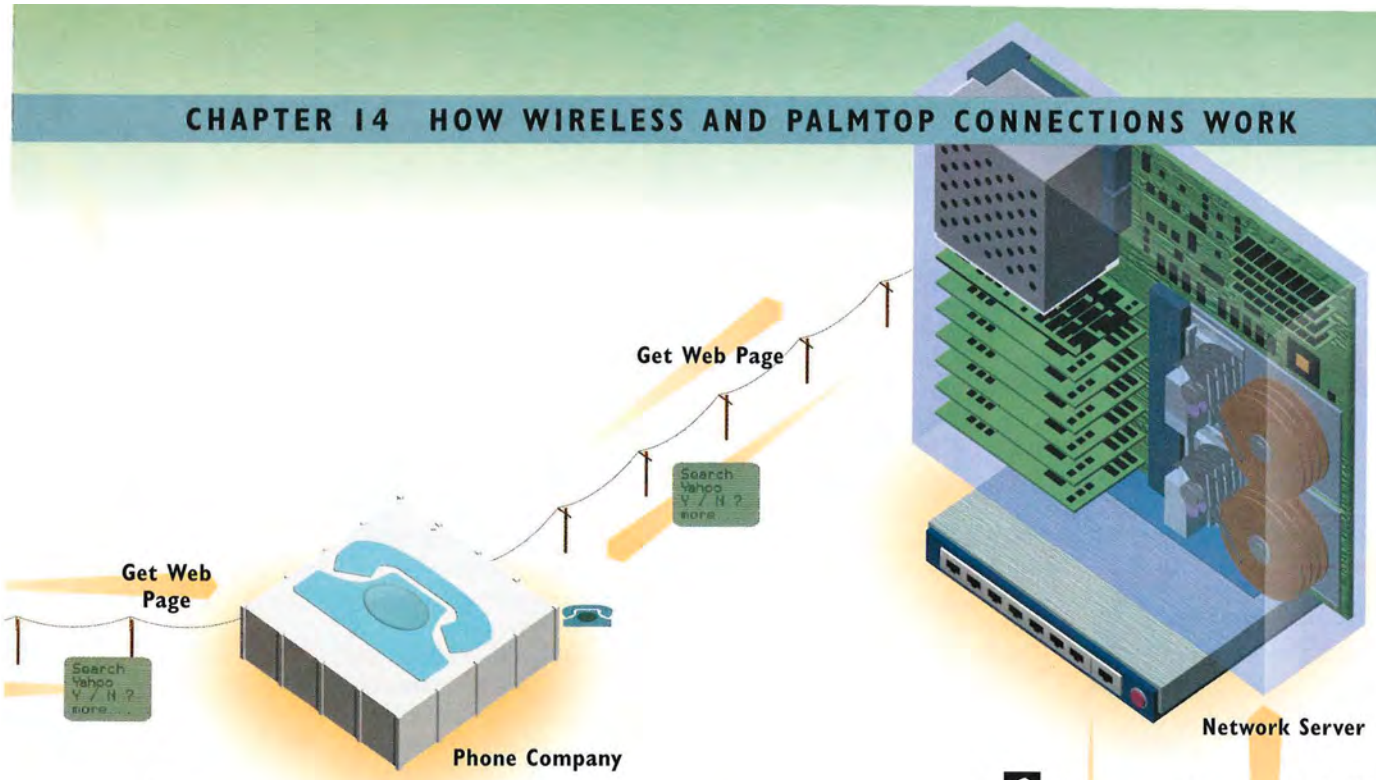


2 The phone asks for authorization to make the call. The base station looks at the telephone's mechanical serial number (MSN) and electronic serial number (ESN) to ensure that the phone is allowed to use the cellular network.

3 The base station sends the call to a telephone-wire–based network (also called a landline) and then to a network server and WAP gateway.



8 You can now read the page on your cellular telephone—it has been specially formatted for its display. However, cellular phones have difficulty handling graphics, so not all Web pages will display properly, even after they've been reformatted to WML.



7 The WML page is sent back through the landline to a base station. The base station sends the page to your cellular telephone.

6 If the page is in normal HTML format, the gateway reformats the page in WML, so your cell phone can read the page. If the page is already in WML format, the gateway doesn't need to reformat the page.

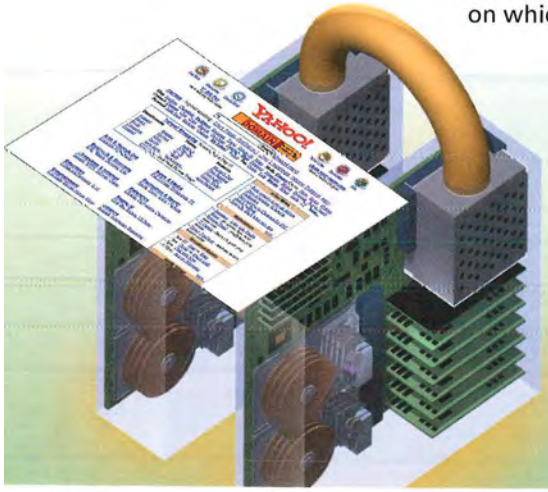


Get Web Page





4 The gateway sends the request for the Web page on to the Web server on which the page is located.



Get Web Page



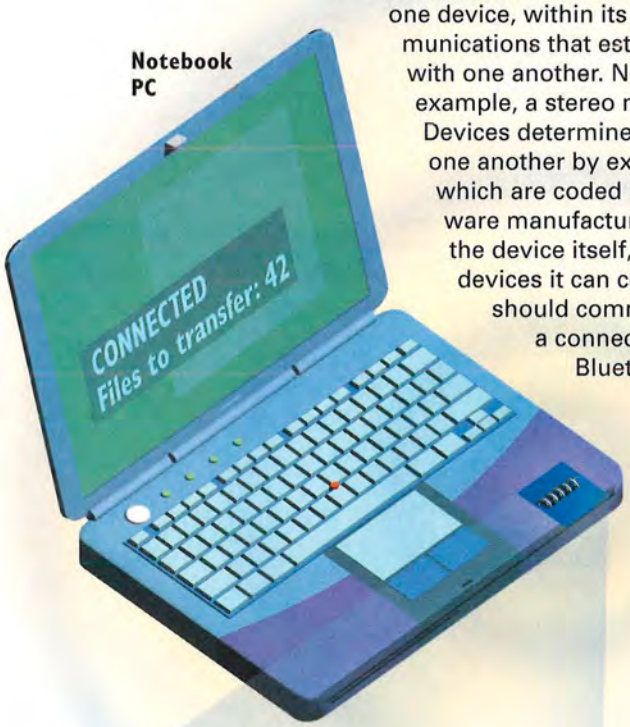
5 The page is sent back to the WAP gateway.

How Bluetooth Works



Router

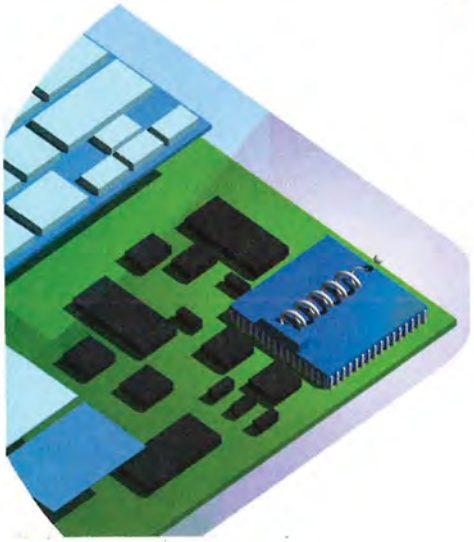
3 When a Bluetooth device finds another device, or more than one device, within its range, they go through a series of communications that establish whether they should communicate with one another. Not all devices will communicate—for example, a stereo might not communicate with a telephone. Devices determine whether they should communicate with one another by examining each other's Bluetooth profiles, which are coded into the devices' hardware by the hardware manufacturer. *Profiles* contain information about the device itself, what it is used for, and with which devices it can communicate. If devices determine they should communicate with one another, they establish a connection. The connection of two or more Bluetooth devices is called a *piconet*.



Notebook PC

2 The Bluetooth device constantly sends out a message, looking for other Bluetooth devices within its range.

1 Bluetooth is a wireless technology that allows computers, telephones, personal data assistants (PDAs), and even home devices such as stereos and TVs to communicate with one another. Each Bluetooth device has a microchip embedded in it that can send and receive radio signals. It can send both data and voice. The radio signals are sent and received in the 2.4GHz radio band, often referred to as the *industrial, scientific, and medical (ISM)* band. Inside the chip is software called a *link controller* that does the actual work of identifying other Bluetooth devices and sending and receiving data.

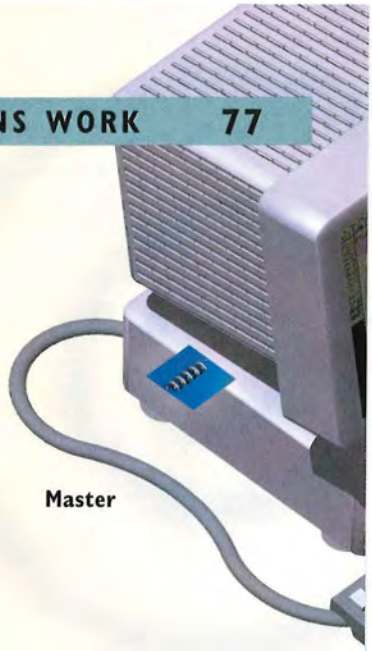


Bluetooth Chipset

Piconet I



4 After the connection is established, the devices can communicate with one another. You could use a Bluetooth device to access information from the Internet, if the device it's accessing the data from is connected to the Internet. For example, you could have a home network with Bluetooth capability and connect a palmtop computer to the Internet by connecting through the home network.



Master

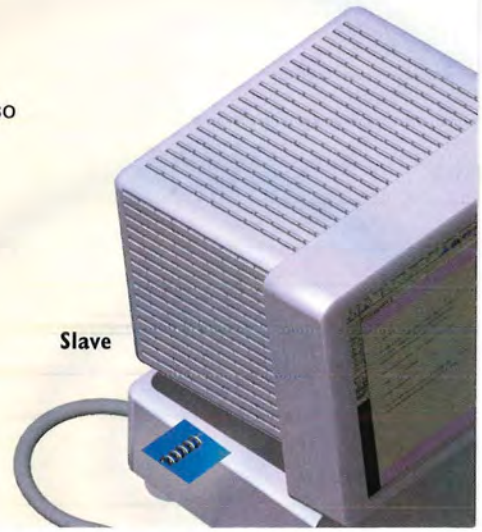


Slave

Handheld

5 If there are many Bluetooth devices or piconets near each other, their radio signals could conceivably interfere with one another. To ensure that doesn't happen, Bluetooth uses *spread-spectrum frequency hopping*. In this technique, the transmitters change their frequencies constantly—1,600 times per second. In this way, the chance of interference is very small—and if interference does occur, it occurs for only a tiny fraction of a second. When two or more devices are connected in a piconet, one device is the master and determines the frequencies to switch among. It instructs all the other devices on which frequencies to switch to, and when.

6 Piconets can link to one another, and any Bluetooth device can also be part of more than one piconet.

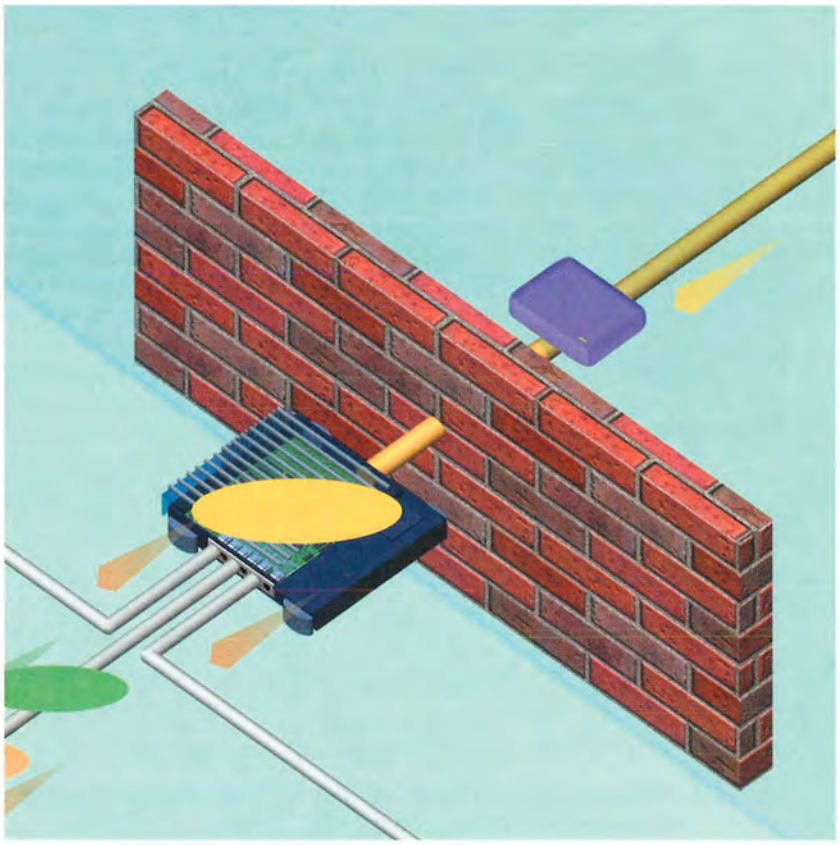


Slave

Piconet 2

CHAPTER

15 How Home Networks Work



THE more the Internet becomes a pervasive and necessary part of our lives, the more we need access to it—and not just at work, but at home as well.

Internet access at home has become a must-have for most people and families who can afford it. And because people have more than one computer at home, they want to be able to let each computer have Internet access—and they want to be able to share printers and send files back and forth between their computers as well.

This is especially important for people who have high-speed Internet connections, such as cable modems or DSL modems. Typically, these connections cost at least $30 a month or more. Rather than having to pay a separate connection fee for each computer, you can install a home network and allow two or more computers to share the same high-speed connection. That way, you pay for only one connection.

Several types of home networks can be used. By far the most popular right now are those that use a combination hub/router to connect the computers and allow them to share an Internet connection, printers, and files. The *hub* connects all the devices to one another, whereas the *router* does the work of connecting the devices to the Internet. Commonly, individual PCs are connected to the hub/router via the same type of Ethernet cables used in corporate local area networks (LANs). Each PC needs a network card to connect to the network—again, the same kind of network card used in corporate LANs.

Other types of home networks are becoming popular as well. Some home networks connect via the home's telephone wires. Still others connect using wireless technology—and wireless home networks might become the most popular because of the difficulty and expense of stringing Ethernet cable throughout a home. For now, though, the most popular networks are those that use the same Ethernet cables used in corporate LANs.

Increasingly, it's not only computers at home that connect to the Internet. Other types of devices and appliances are being connected as well, such as radio receivers; small, inexpensive e-mail devices that only send and receive e-mail; and even traditional home appliances such as refrigerators, microwave ovens, and alarm clocks.

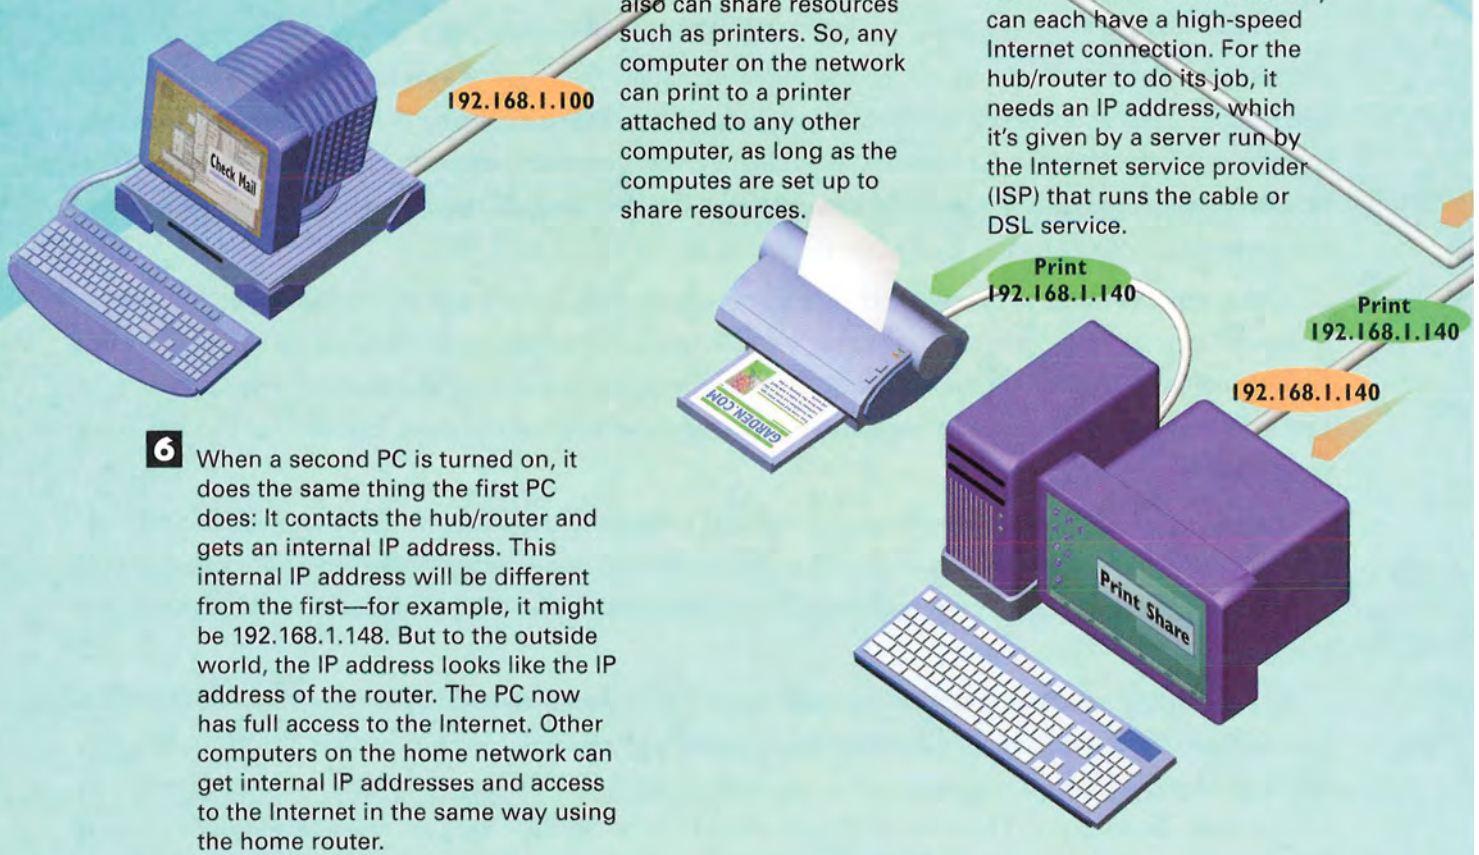
Connecting these types of devices and appliances will make life more convenient—you'll be able to use your refrigerator to automatically generate shopping lists, for example, and send orders directly to grocery stores. And you'll have an alarm clock that can change the time it awakens you based on traffic reports it garners from the Internet. These kinds of devices aren't mere fantasy—they are already being sold or tested and will soon be sold in a department store near you.

How Home Networks Let You Share Your Internet Connection

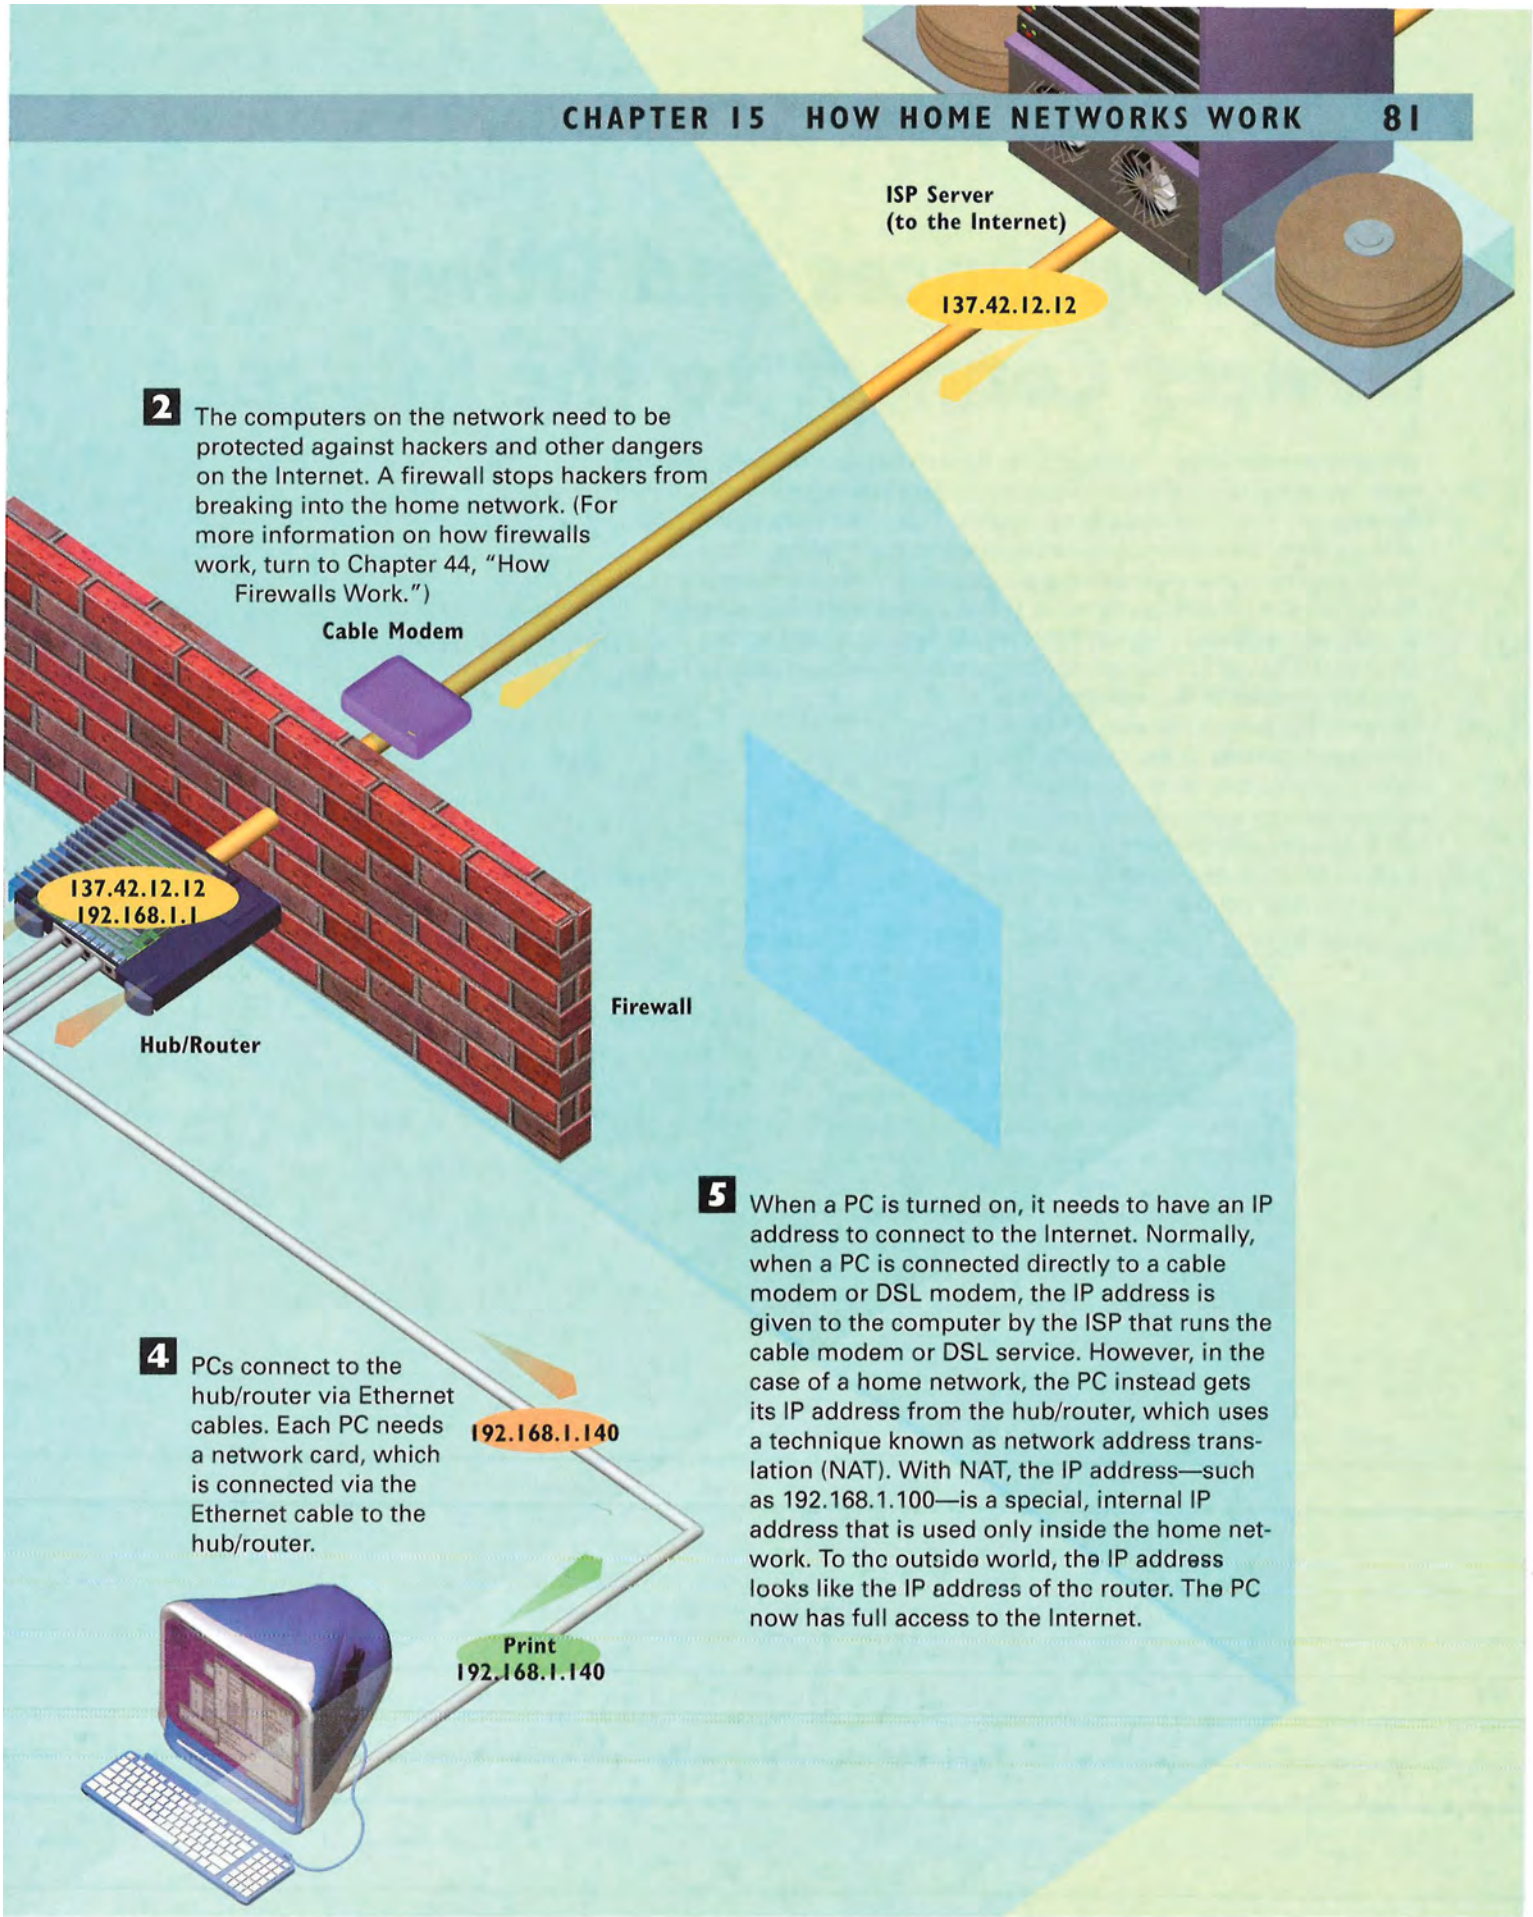
1 The main reason people set up home networks is to share high-speed Internet access, such as via a cable modem or DSL modem, among several computers. To give the home network access to the Internet, an Ethernet cable runs from the cable modem or DSL modem to the networking device, called a *hub/router*.

3 The hub/router does two primary jobs: It connects all the PCs to each other so they can share files and devices such as printers, and it connects all the PCs to the Internet so they can each have a high-speed Internet connection. For the hub/router to do its job, it needs an IP address, which it's given by a server run by the Internet service provider (ISP) that runs the cable or DSL service.

7 In addition to gaining Internet access through the hub/router, the PCs also can share resources such as printers. So, any computer on the network can print to a printer attached to any other computer, as long as the computes are set up to share resources.



6 When a second PC is turned on, it does the same thing the first PC does: It contacts the hub/router and gets an internal IP address. This internal IP address will be different from the first—for example, it might be 192.168.1.148. But to the outside world, the IP address looks like the IP address of the router. The PC now has full access to the Internet. Other computers on the home network can get internal IP addresses and access to the Internet in the same way using the home router.



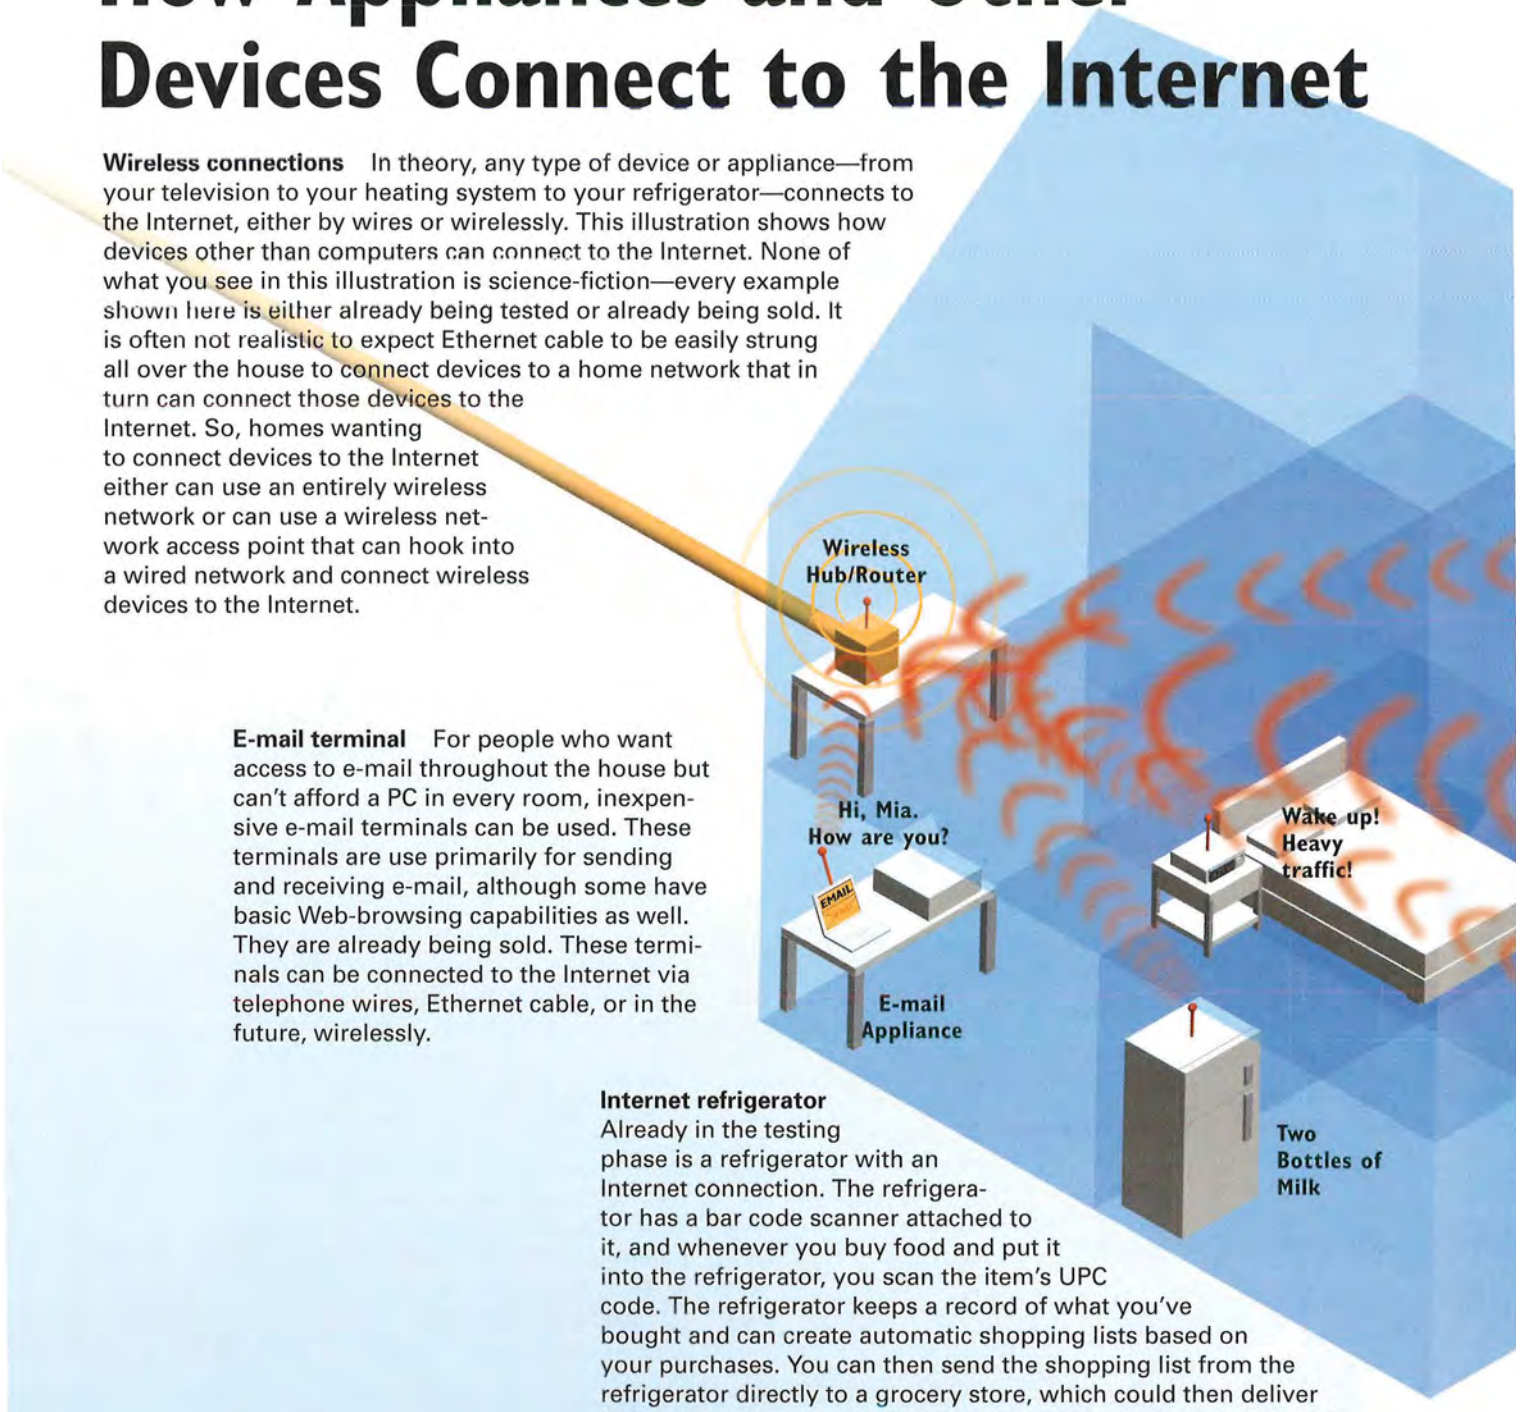
2 The computers on the network need to be protected against hackers and other dangers on the Internet. A firewall stops hackers from breaking into the home network. (For more information on how firewalls work, turn to Chapter 44, "How Firewalls Work.")

4 PCs connect to the hub/router via Ethernet cables. Each PC needs a network card, which is connected via the Ethernet cable to the hub/router.

5 When a PC is turned on, it needs to have an IP address to connect to the Internet. Normally, when a PC is connected directly to a cable modem or DSL modem, the IP address is given to the computer by the ISP that runs the cable modem or DSL service. However, in the case of a home network, the PC instead gets its IP address from the hub/router, which uses a technique known as network address translation (NAT). With NAT, the IP address—such as 192.168.1.100—is a special, internal IP address that is used only inside the home network. To the outside world, the IP address looks like the IP address of the router. The PC now has full access to the Internet.

How Appliances and Other Devices Connect to the Internet

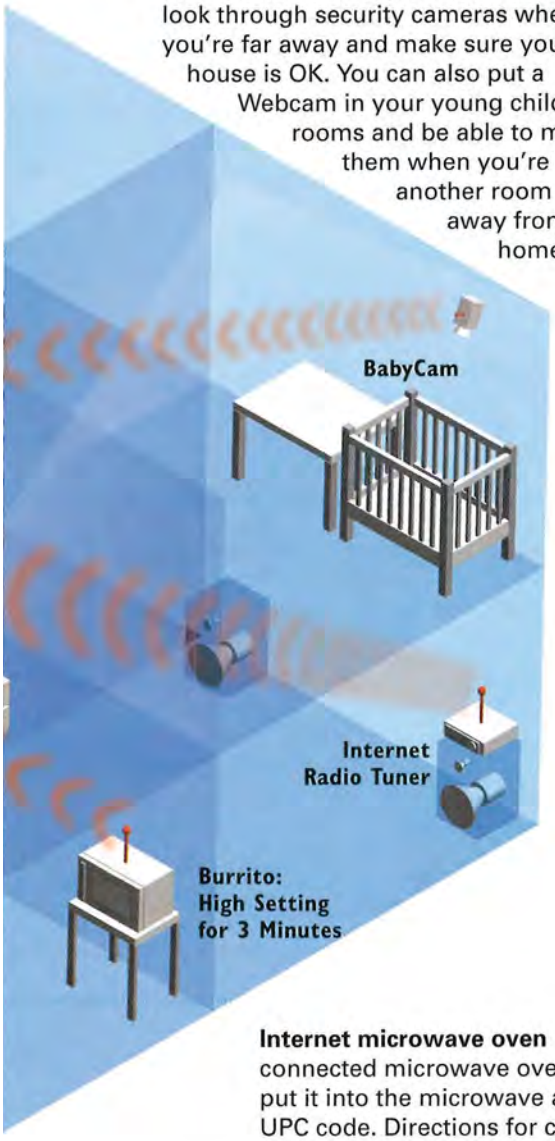
Wireless connections In theory, any type of device or appliance—from your television to your heating system to your refrigerator—connects to the Internet, either by wires or wirelessly. This illustration shows how devices other than computers can connect to the Internet. None of what you see in this illustration is science-fiction—every example shown here is either already being tested or already being sold. It is often not realistic to expect Ethernet cable to be easily strung all over the house to connect devices to a home network that in turn can connect those devices to the Internet. So, homes wanting to connect devices to the Internet either can use an entirely wireless network or can use a wireless network access point that can hook into a wired network and connect wireless devices to the Internet.



E-mail terminal For people who want access to e-mail throughout the house but can't afford a PC in every room, inexpensive e-mail terminals can be used. These terminals are use primarily for sending and receiving e-mail, although some have basic Web-browsing capabilities as well. They are already being sold. These terminals can be connected to the Internet via telephone wires, Ethernet cable, or in the future, wirelessly.

Internet refrigerator Already in the testing phase is a refrigerator with an Internet connection. The refrigerator has a bar code scanner attached to it, and whenever you buy food and put it into the refrigerator, you scan the item's UPC code. The refrigerator keeps a record of what you've bought and can create automatic shopping lists based on your purchases. You can then send the shopping list from the refrigerator directly to a grocery store, which could then deliver the goods to your home. Refrigerator makers Electrolux, GE, and Whirlpool are planning to sell these appliances.

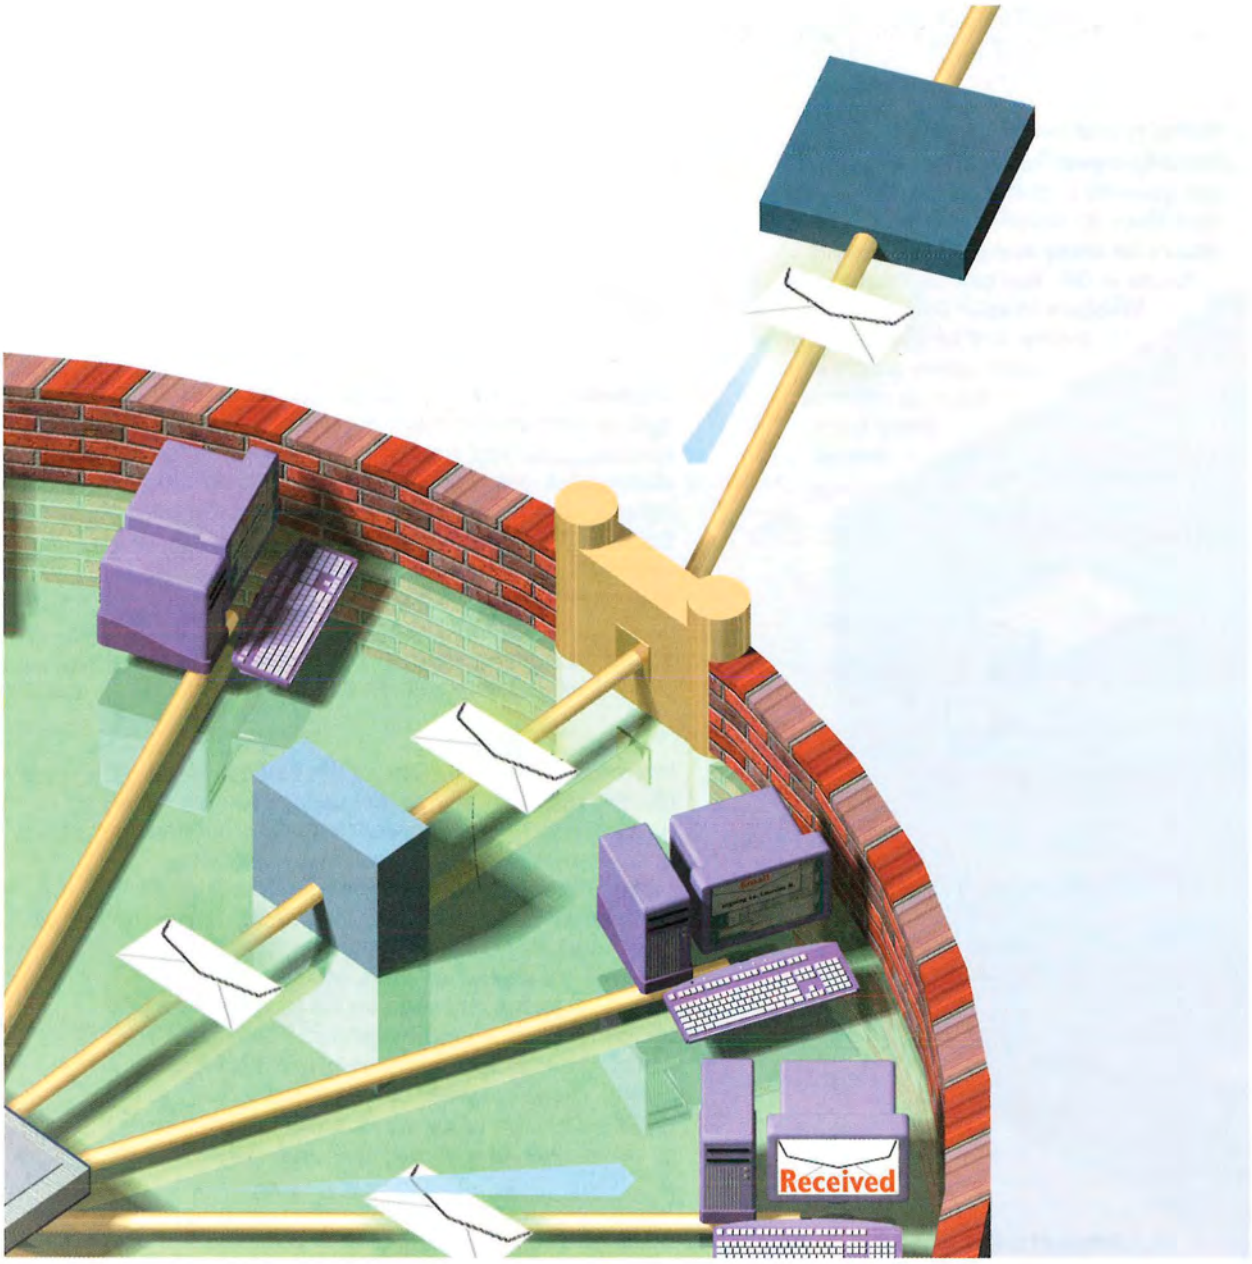
Security and monitoring system Security systems can be connected to the Internet so that you can, for example, look through security cameras when you're far away and make sure your house is OK. You can also put a Webcam in your young children's rooms and be able to monitor them when you're in another room or away from home.



Internet home control Ultimately, every part of your home will be able to be monitored and controlled via the Internet, from the lights to the heating system, security system, and beyond. You'll be able to monitor and control your home either through a central keypad and monitor at home, via a PC at home, or via a PC or some other Internet-connected device when you're away from home. Sunbeam, which makes the Mr. Coffee coffee maker, has announced that it will be selling a set of nine Internet-connected home devices, including a smart, Internet-connected alarm clock called the TimeHelper. The device can beep a warning when the Internet-connected Smart Coffeemaker runs low on water, check Internet traffic reports to see what traffic is like, set your alarm to the best time for getting into work, and even turn off or on your electric blanket based on weather reports. The Internet alarm clock is expected to sell for $40–$80.

Internet radio tuner One of the newest home devices that uses an Internet connection is a radio tuner. These tuners connect to the Internet and play Internet radio stations. They are already being sold. (For information about how Internet radio stations work, turn to Chapter 36, "How Music and Audio Work on the Internet.") They connect in a variety of ways, either through a PC or by themselves through an Ethernet connection. (In the future, they might connect wirelessly.) They can play the music in speakers throughout the house. Those speakers can be wired directly to the tuner in the traditional way, connected wirelessly, or connected via a home's telephone wires.

Internet microwave oven Also in the testing phase is an Internet-connected microwave oven. When you want to cook something, you put it into the microwave and pass a bar code scanner over the product's UPC code. Directions for cooking the item are sent directly to the microwave oven, which then cooks the item according to the directions. Samsung is planning on selling an Internet microwave oven. Some day, regular ovens might have this same capability as well.


Received

PART 3

COMMUNICATING ON THE INTERNET

Chapter 16: How E-mail Works
88

Chapter 17: How E-mail Spam Works
100

Chapter 18: How Newsgroups Work
106

Chapter 19: How Internet Chat and Instant Messaging Work
110

Chapter 20: How Internet Phone Calls Work
116

SINCE its very earliest days, the Internet has been concerned primarily with one task—making it easier for people to communicate with one another using computers. The Internet was created to enable university researchers to share their thoughts, work, and resources, and for military people to communicate with each other in case of war and even, theoretically, a nuclear attack.

Today, more than two decades after the inception of the first networks that grew into the Internet, it is still primarily a communications medium. Millions of people from all over the world share their thoughts, hopes, work, gossip, and comments on the wires and computers that make up the Internet. Many of the means of communication, such as electronic mail, have changed very little in the past 20 years. Yet other, entirely new ways of communicating have been devised, such as enabling you to use the Internet as your telephone, completely cutting out long-distance charges, even if you're calling to the other side of the world. Technologies enable people to communicate privately, one on one; others allow for vast discussion groups that span the globe, and still others enable both private communication with one person and public communication with large groups.

This section of the book looks at the main ways people communicate on the Internet.

Chapter 16, "How E-mail Works," takes a long look at what continues to be the most popular way for people to communicate on the Internet—electronic mail, or e-mail. E-mail remains possibly the greatest use of the Internet and is used for business and personal communication. You'll see how e-mail gets routed from your computer through the maze of wires that makes up the Internet and then ends up in the proper recipient's inbox. You'll look at all the elements of a mail message and learn how you can send binary files, such as pictures and sounds, through e-mail. You'll also explore mailing lists, where you can subscribe to any one of thousands of public discussions via e-mail or receive what are essentially electronic newsletters delivered to your e-mail inbox. You'll also look at how you can look up anyone's e-mail address using white page directories that use a technology called the Lightweight Access Directory Protocol (LADP). Finally in this chapter, you'll take a look at how e-mail can be encrypted so that snoopers and hackers can't read it as it makes its way across the Internet.

Chapter 17, "How E-mail Spam Works," covers one of the most controversial modes of communication on the Internet—the use of spam, the equivalent of junk mail via e-mail, that is sent to millions of people each day. Although spam is a problem on newsgroups as well as e-mail, e-mail is the greatest area of controversy. Spam annoys people and wastes their time while they clean out their mailboxes; it clogs the Internet so that other messages might be delivered late (or not at all); and it can be delivered by piggybacking onto other people's e-mail servers, costing them money. The chapter looks at how spam is sent, as well as ways that it can be blocked.

Chapter 18, "How Newsgroups Work," explores Usenet newsgroups—public discussion groups in which anyone can participate. Many thousands of these groups focus on every subject conceivable. You'll see how newsgroups work and how you can decipher their often arcane names.

Chapter 19, "How Internet Chat and Instant Messaging Work," covers the various ways people can *chat* on the Internet. When people chat on the Internet, they don't actually speak but type comments on their keyboards instead, and then people all over the world can read

and respond to them. This chapter closely examines how instant messaging works—a way in which people can chat one on one with others. Two of the most popular pieces of software for communicating on the Internet—America Online Instant Messenger (AIM) and ICQ—are chat software. This chapter also discusses the first type of Internet chat, called Internet Relay Chat (IRC). Although not as popular as it used to be, it's still a way many people chat with each other online.

Finally, Chapter 20, "How Internet Phone Calls Work," details one of the more intriguing new uses of the Internet—using it as your telephone. Today you can dial into your local Internet provider with your computer, and if you have the right hardware and software, you can talk with anyone similarly connected to the Internet anywhere in the world, all without paying long-distance telephone charges. This chapter also covers a new type of low-cost telephone service that has sprung up: IP Telephony. In IP Telephony, you talk over your telephone as you would normally—the phone call appears no different to you. But the Internet, rather than the normal phone system, is used to make the connection to the person you're calling.

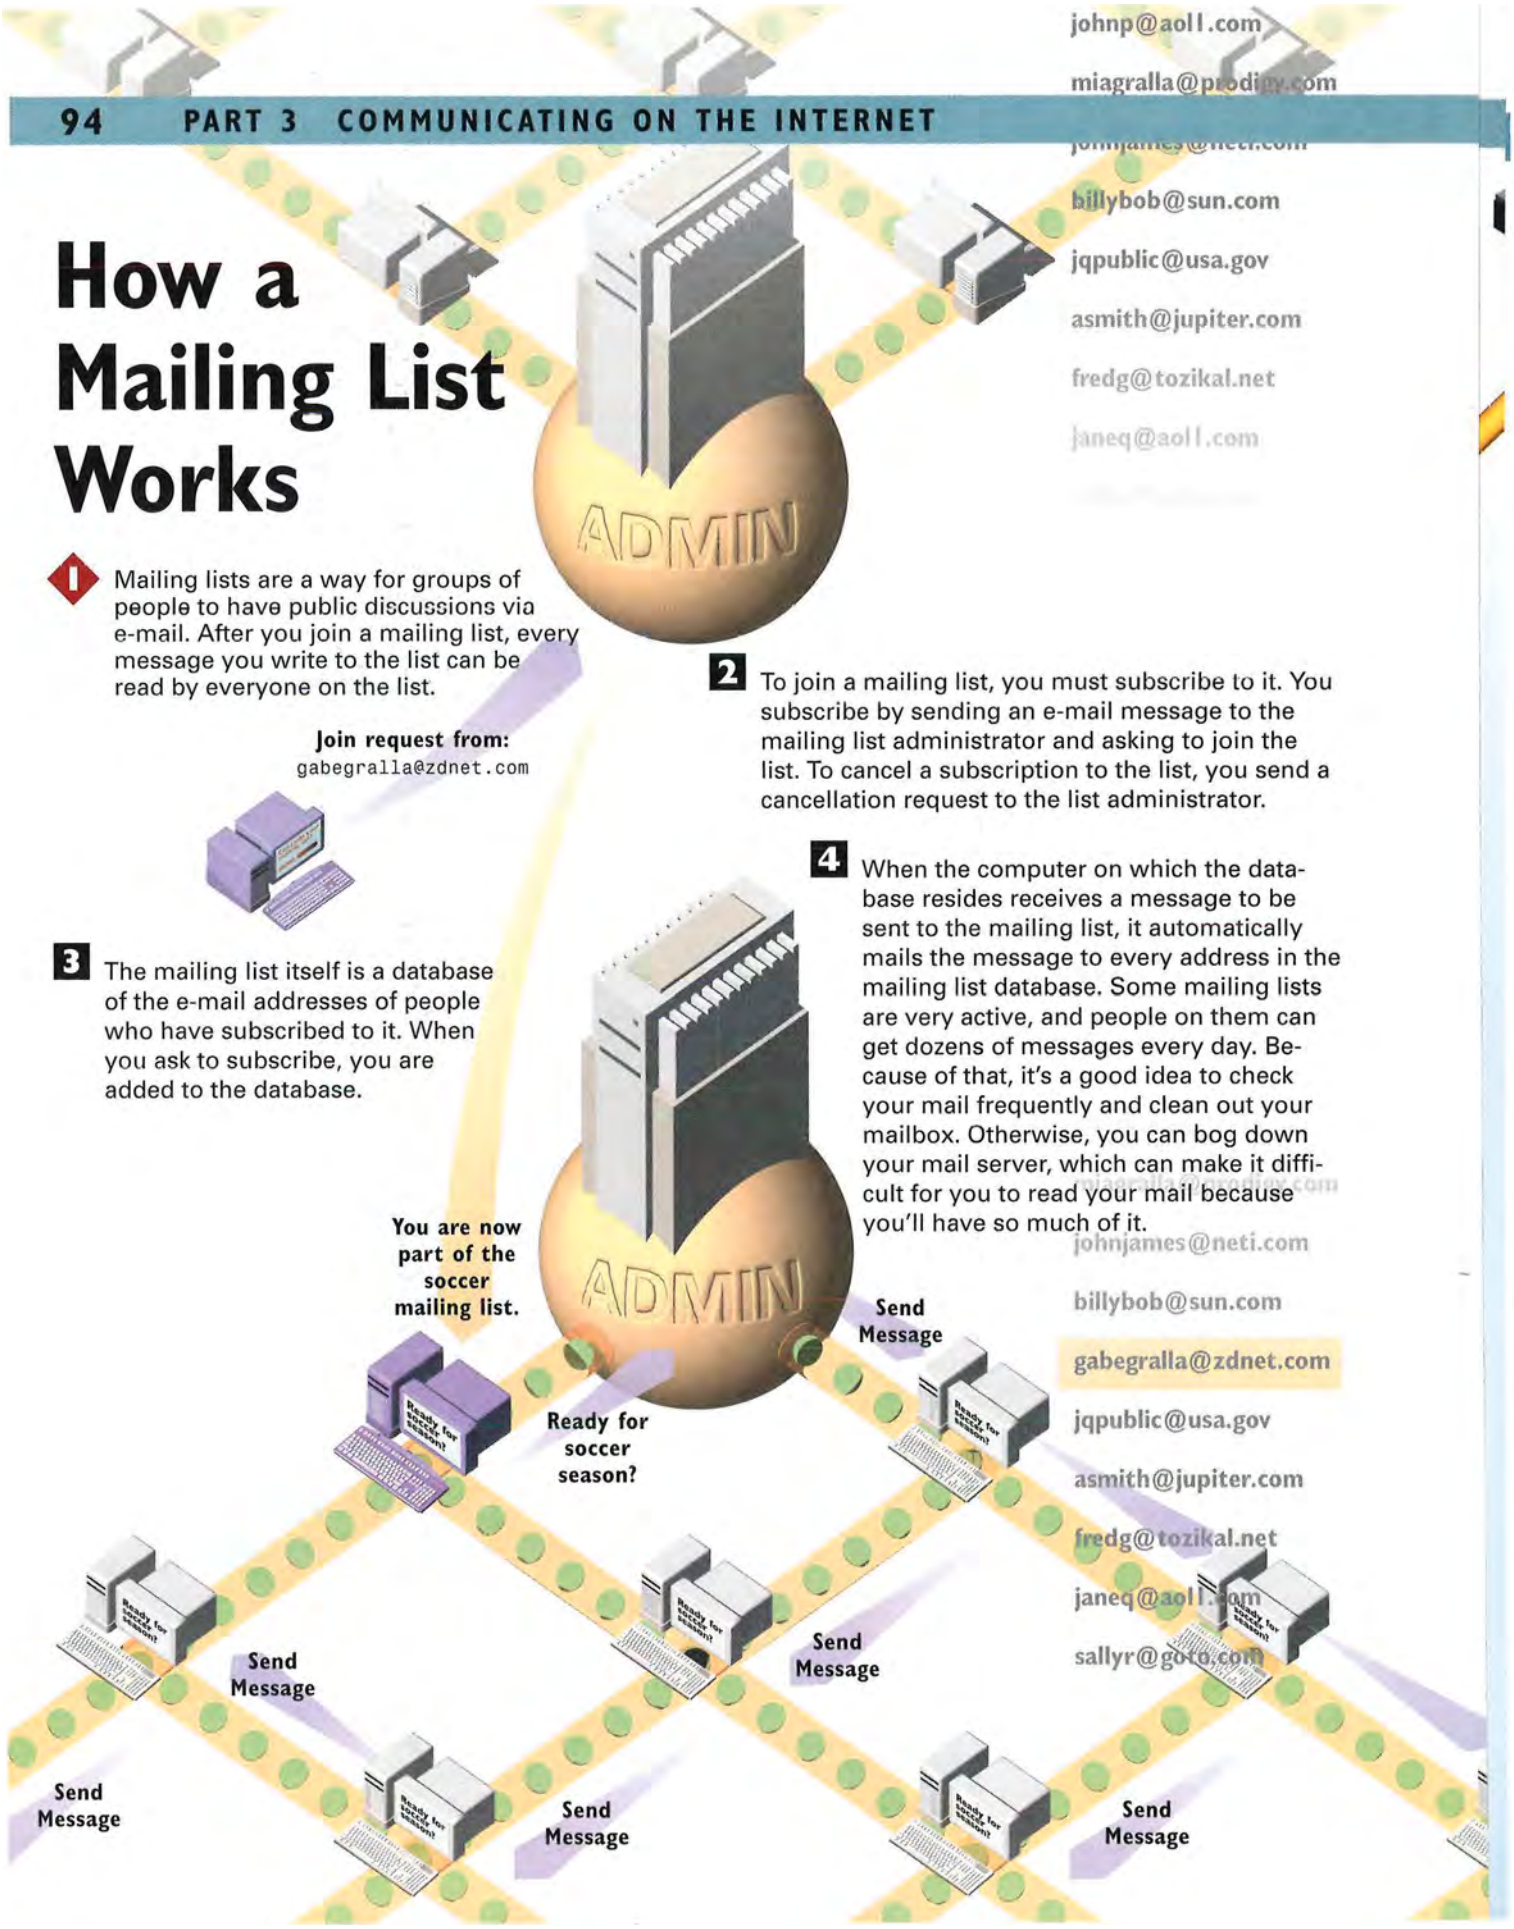
How a Mailing List Works

1 Mailing lists are a way for groups of people to have public discussions via e-mail. After you join a mailing list, every message you write to the list can be read by everyone on the list.

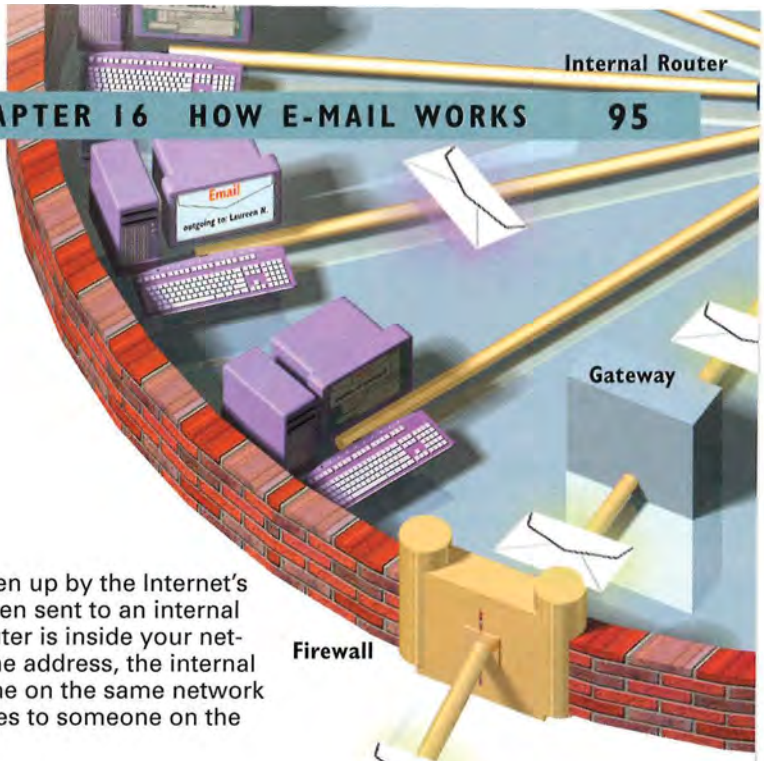
2 To join a mailing list, you must subscribe to it. You subscribe by sending an e-mail message to the mailing list administrator and asking to join the list. To cancel a subscription to the list, you send a cancellation request to the list administrator.

3 The mailing list itself is a database of the e-mail addresses of people who have subscribed to it. When you ask to subscribe, you are added to the database.

4 When the computer on which the database resides receives a message to be sent to the mailing list, it automatically mails the message to every address in the mailing list database. Some mailing lists are very active, and people on them can get dozens of messages every day. Because of that, it's a good idea to check your mail frequently and clean out your mailbox. Otherwise, you can bog down your mail server, which can make it difficult for you to read your mail because you'll have so much of it.



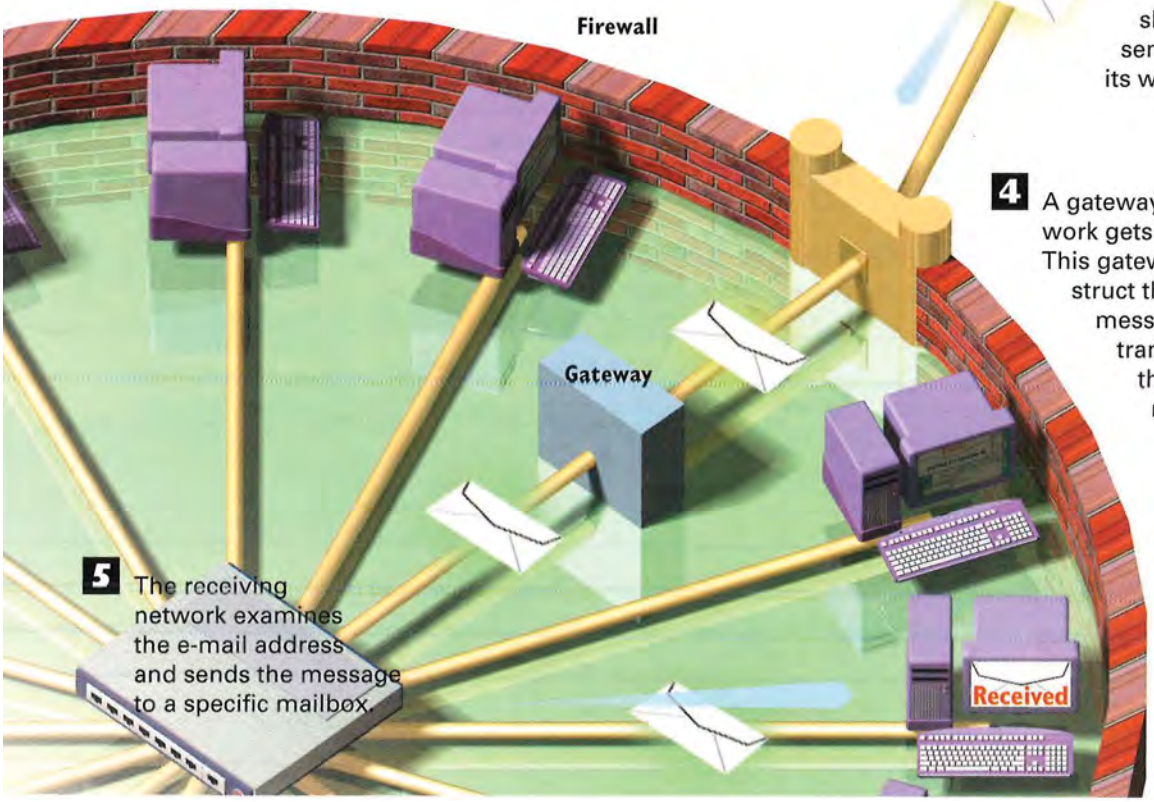
How E-mail Is Sent Between Networks



1 When you send an e-mail message, it is first broken up by the Internet's TCP protocol into IP packets. Those packets are then sent to an internal router that examines the address. (An internal router is inside your network, rather than out on the Internet.) Based on the address, the internal router decides whether the mail is sent to someone on the same network or to someone outside the network. If the mail goes to someone on the same network, the mail is delivered to them.

2 If the mail is addressed to someone outside the network, it might have to pass through a firewall—a computer that shields the network from the broader Internet so intruders can't break into the network. The firewall keeps track of messages and data going into and out of the network, to and from the Internet. It can also prevent certain packets from getting through it.

3 Once out on the Internet, the message is sent to an Internet router. The router examines the address, determines where the message should be sent, and then sends the message on its way.



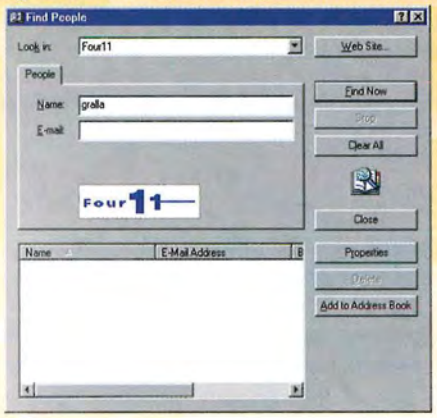
4 A gateway at the receiving network gets the e-mail message. This gateway uses TCP to reconstruct the IP packets into a full message. The gateway then translates the message into the protocol the target network uses and sends it on its way. The message also might be required to pass through a firewall on the receiving network.

5 The receiving network examines the e-mail address and sends the message to a specific mailbox.

How White Page Directories (LDAP) Work

1 Millions of people have e-mail addresses on the Internet, but finding an e-mail address has always been difficult if you know only a person's name. The Lightweight Directory Access Protocol (LDAP) makes creating white page–style directories easier, so you now can look up people's addresses if you know only their names. An LDAP directory is a database that follows LDAP rules and protocols for organizing its information. It is located on an Internet server. A company or business maintains the server and is responsible for keeping the information on it up-to-date by adding new names, deleting old names, and updating existing names.

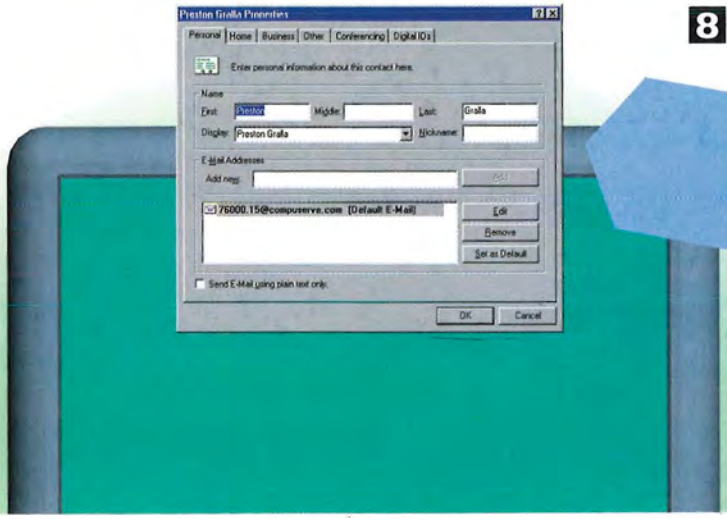




BINDRESPONSE

2 Like much other Internet technology, LDAP works on a client/server model. The client runs on a local computer. Many types of clients are available for many types of computers. To find someone's e-mail address, you run the client on your computer and type in the name of the person whose e-mail address you want to find.

Add Name to Address Book



8 Depending on the software you're using, you can now add the name and e-mail address to your local address book, or you can immediately send an e-mail to the person.

3 Before sending the request itself, the client sends an LDAP command, called the `BindRequest`, to the server. Basically, this is a command telling the server the client wants to make a request.

BINDREQUEST

4 The server answers with a command called the `BindResponse`, telling the client it can go ahead and make a request.

FIND "GRALLA"

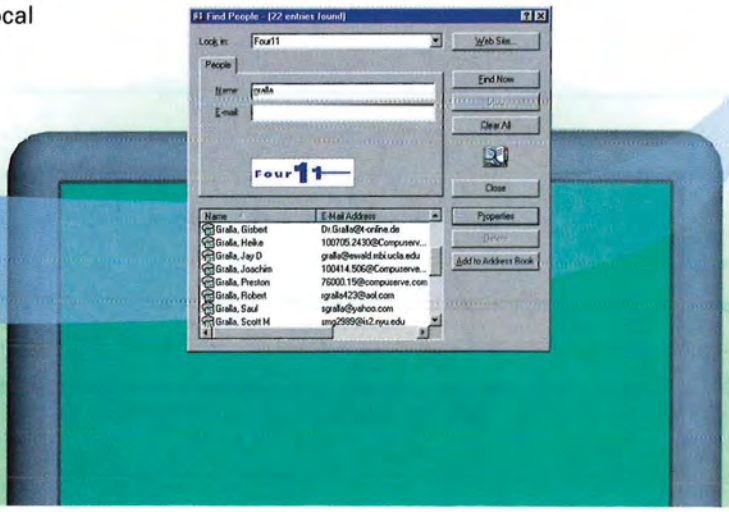
5 The client then sends the request to the server, sending a person's name and asking for his e-mail address.

LDAP
Lightweight Directory Access Protocol

Michael Troller	mt@m-troller.com
Noel Voskuil	noelv@pacbell.net
Alan Scott	a_scott@lantern.com
Bill Jones	boj@pol.net
Mary Peterson	mpters@rocketmail.com
Crash Corrigan	crash_98@excitemail.com
Stephan Adams	nanna_513@cincimail.net
Jenna MacCimbri	jim@zamora.com

HERE ARE THE RESULTS

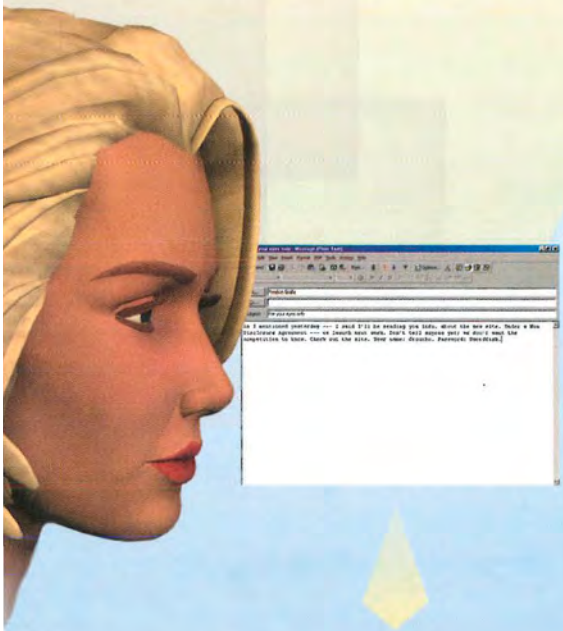
6 The LDAP directory performs a search and finds the e-mail address or addresses in its database. It sends the results to the client.



HERE ARE THE RESULTS

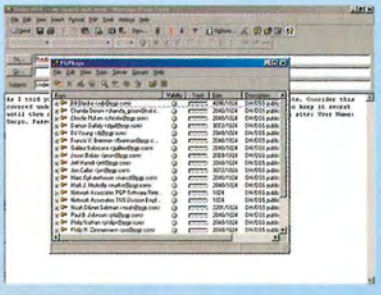
7 The client receives the results and displays them on the local computer.

How Encryption Can Keep E-mail Private



1 Encryption can be used to hide e-mail messages so snoopers and hackers can't read them as they are sent across the Internet. Many types of encryption exist, but in the most common one, keys are used. Everyone gets a public and private key. The *public key* is available for anyone to use to encrypt mail; the *private key* is used only by the recipient to decrypt it. This illustration shows how to encrypt e-mail with the popular encryption program Pretty Good Privacy (PGP). To start using PGP, someone uses her normal e-mail program to compose a piece of mail. After she composes the mail, she decides whether she wants to encrypt the message.

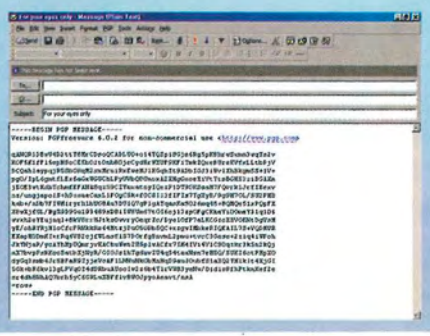


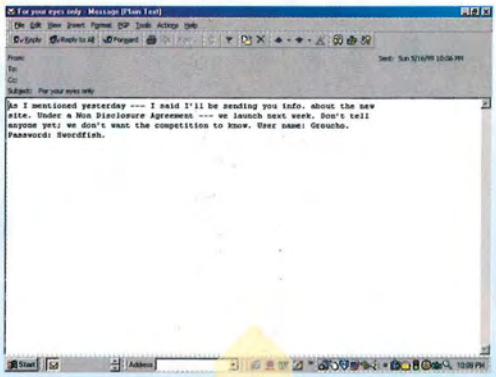


2 When someone decides to encrypt a piece of e-mail, she has to have a copy of the public key of the person to whom she is going to send the message. That key can be obtained a number of ways—from a public Internet site, or the recipient can send it via e-mail. After the person has the key, she stores it in a key ring on her computer and can call it up at any time.

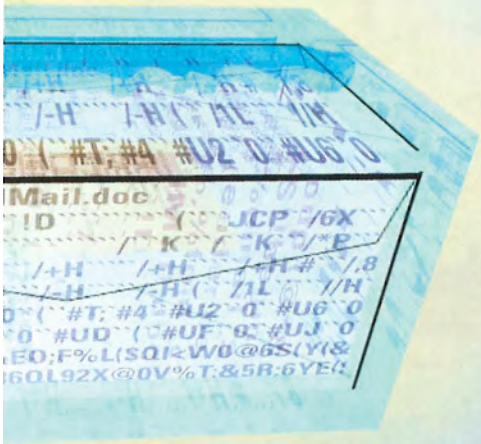
3 After choosing the key of the person who will receive the message, the message is encrypted using that person's public key.







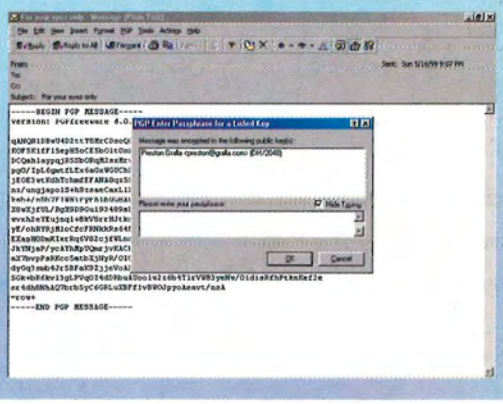
6 The private key decrypts the message, and the person can read it and use it like any other piece of e-mail.



5 The person receives the e-mail as he does any normal piece of e-mail. However, the e-mail is encrypted, and so can't yet be read. So, the person uses his private key to decrypt the message. Before the message can be decrypted, a password for the private key typically must be typed in. However, some e-mail programs and encryption software can be set so that the message is decrypted automatically, as soon as it's received.



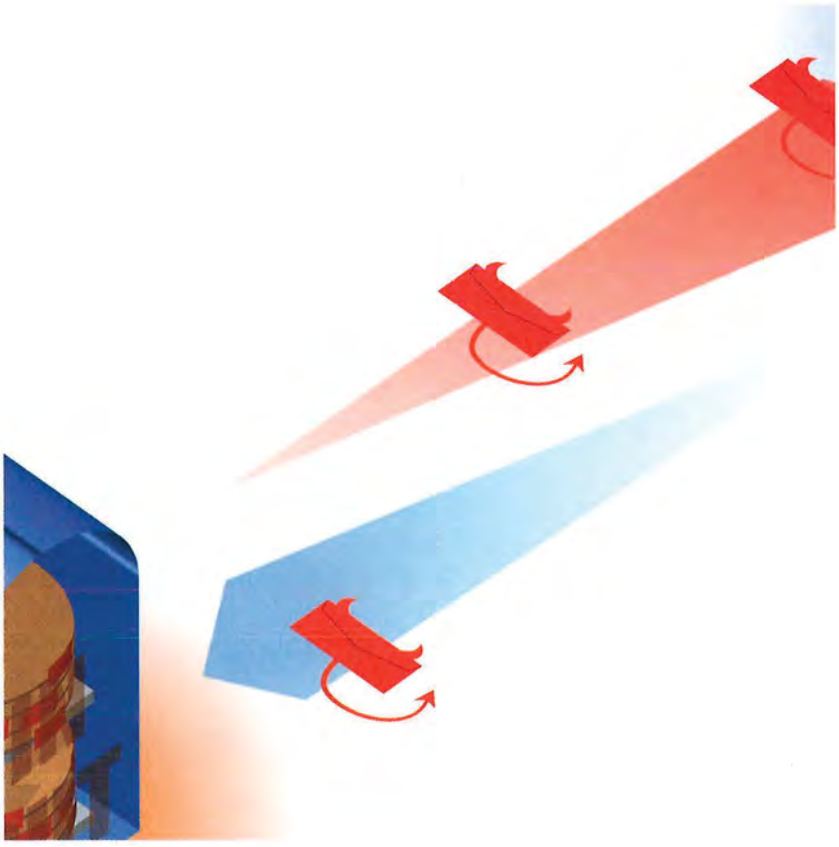
4 The encrypted e-mail is sent in the same way that any normal e-mail is. The difference is that the e-mail is encrypted so that anyone who reads the mail as it makes its way across the Internet won't be able to understand it—he'll see only random characters.





CHAPTER

17 How E-mail Spam Works



ONE of the most contentious issues to surface on the Internet in the last few years has to do with what Internet users call *spam*. Spam is unsolicited junk e-mail that commercial companies send out, asking you to buy their goods and services. At times it might contain inducements to visit the seller's site. The e-mail usually contains a phone number to call, an address to send money to, or a Web site to visit to buy the goods and services.

The term *spam* comes from a Monty Python skit in which every item on a menu contained Spam luncheon meat. It was originally used to refer to unsolicited postings for commercial products or services on Usenet, especially when they were cross-posted to several newsgroups.

Spam might seem like a minor annoyance, but the truth is, it can cause major problems. Spam floods the Internet with unwanted mail, which can lead to delayed or lost mail. It clogs the Internet pipeline, making other information slower to send. It wastes time for those who have to go through their e-mail boxes deleting unsolicited mail, especially when they pay for their e-mail service by the hour. Additionally, it's fairly common for spammers to hide their real e-mail addresses by forging other people's names onto the From or Sender header of an e-mail message. So, those people whose names were forged might be the target of angry mail. This makes it difficult for Webmasters and mail administrators to filter spam messages by From address or domain name. Sometimes spammers even use other people's servers to deliver their bulk e-mail; in essence, forcing someone else to pay the costs of the spammer's mail delivery.

In some ways, spam is not very different from traditional junk mail. Spammers buy or compile massive lists of e-mail addresses, in the same way that junk mailers buy or compile U.S. Postal addresses. The spammer then uses special software to send a solicitation to every person on the list—not uncommonly, tens of thousands of pieces of e-mail in a single spam mailing. To hide their true identities, spammers forge names onto the headers of e-mail messages and even "relay" their spam to another mail server on the Internet, so that finding out where the mail comes from is impossible. Often, a user will request to be taken off the list by replying to e-mail addresses the spammers provide. However, this verifies the user's address and he will get even more spam.

A variety of ways have been devised to block spam, including having e-mail filters on e-mail software ignore any mail from known spammers. This doesn't always work well, however, because spammers often change or forge their e-mail addresses. There are also calls for the courts or Congress to take action. Congress has considered several laws, including one that would ban spam entirely, in the same way that junk faxes were banned several years ago. Until then, try doing an Internet search for Web sites that will help you download and install software to help filter your e-mail and stop spam.

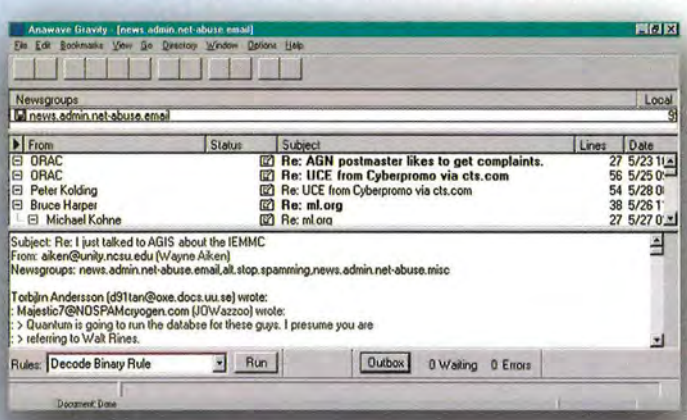
How E-mail Spam Is Sent



1 *Spam* is a term used to describe unsolicited e-mail sent to you, often by commercial firms that attempt to sell you goods and services. Spam is sent out as bulk e-mail, often to lists of 10,000 or more people at once. It's inexpensive to send, so its use has exploded on the Internet to the point where it is common for someone to get several dozen of these messages in a single day. Spam has become enough of an annoyance to warrant calls to ban it outright.

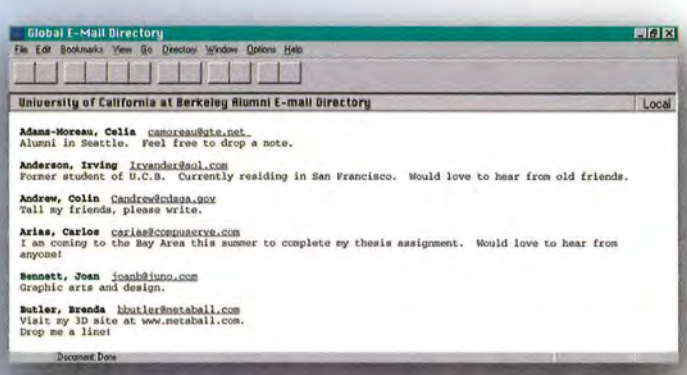
2 To send out unsolicited bulk e-mail, a spammer first needs to get a list of e-mail addresses. Often, spammers buy the lists from companies that compile them. These companies use automated software robots to get the e-mail addresses. The robots get the lists from a number of sources. One way is to go into Usenet newsgroups and harvest e-mail addresses by looking inside every message, which usually has in it the e-mail address of the person who posted it.

Bulk Mail List



Get e-mail address.

ADD: all addresses



3 E-mail addresses also can be harvested from e-mail directories on the Web sites that allow people to look up others' e-mail addresses. Software robots can go into the directory and grab every address in the directory. Robots also can go into chat areas, such as those on America Online, and gather e-mail addresses.

Get e-mail address.

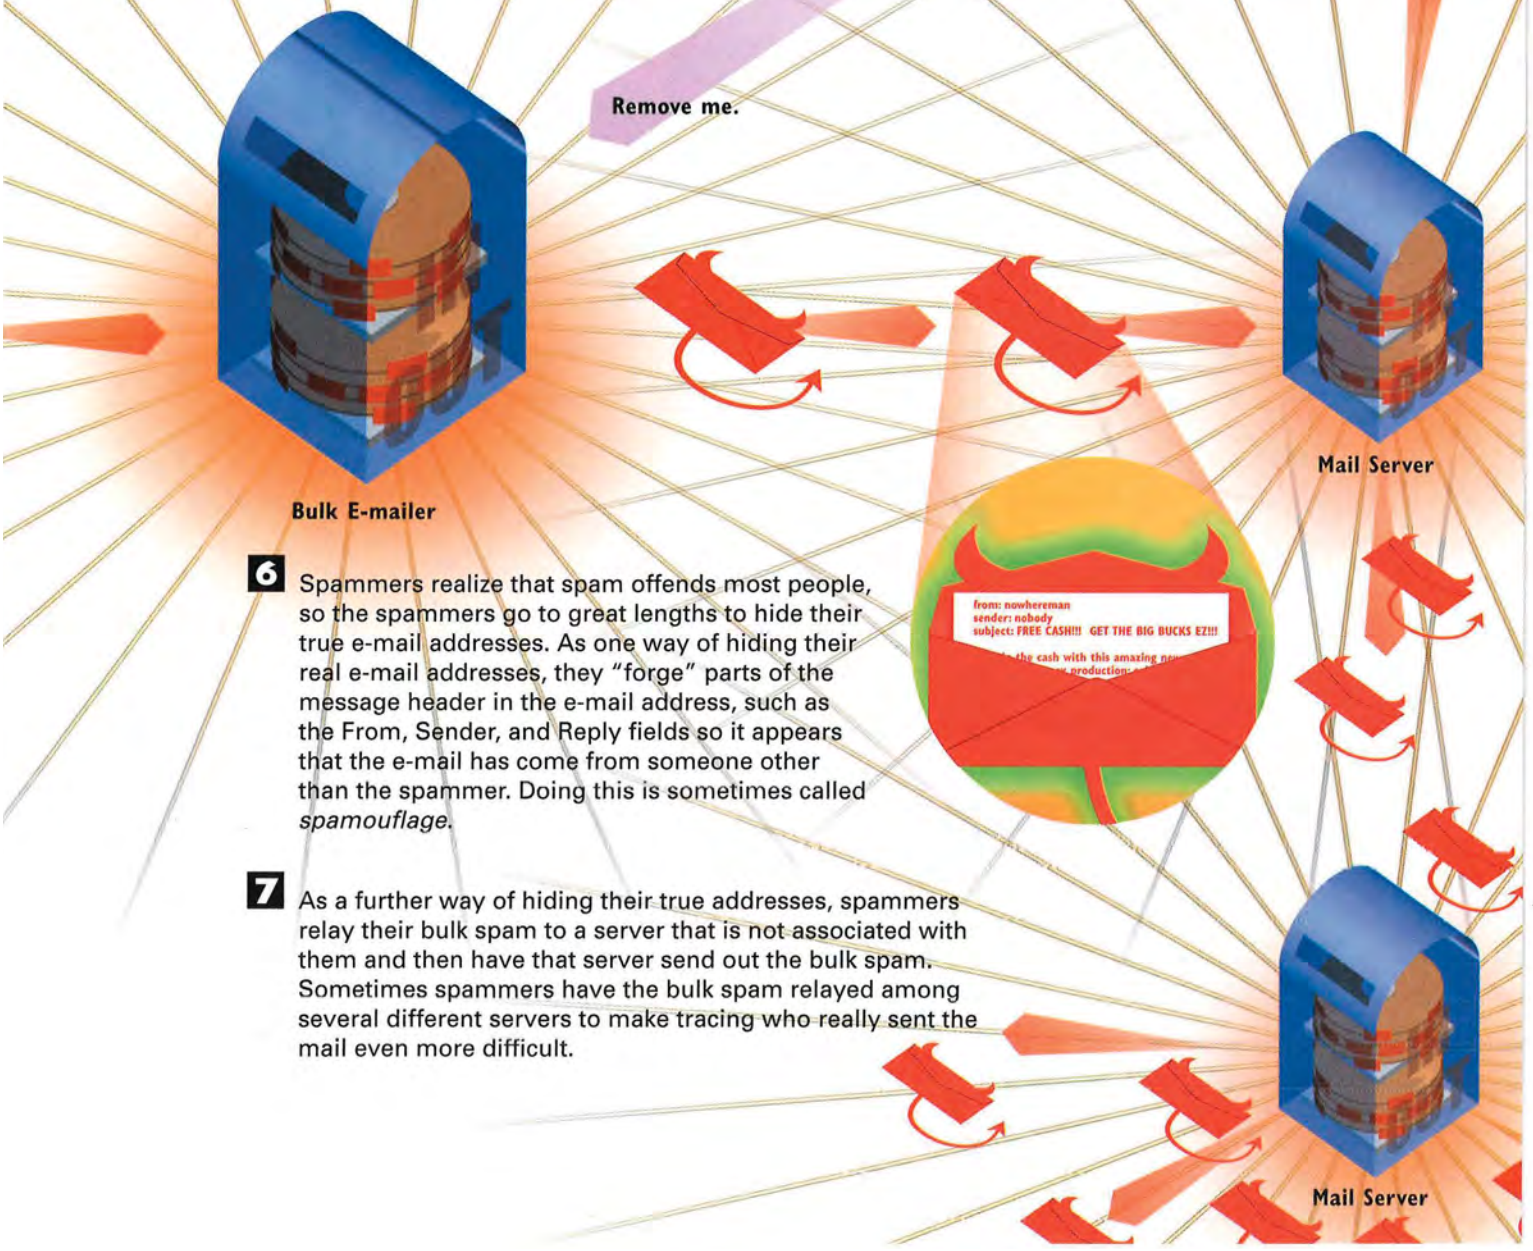
ADD: all addresses

welt@well.com
vegrebski@excite.com
ncsof@webtv.net
johns1739@aol.com
stevej@apple.com
johnh@liberty.net
alxopl@webtv.net
corpflak@msn.com
sillysam7@kidstoys.net
morm@cheerybar.com
johnqp@prodigy.com
eatright3@heath.com
qwertz@yahoo.com
aiken@unity.net
triweb@li.com
zgooby@znakz.com
joanh@juno.com
hotrod99@fastcar.com
richir@kerbanko.com
morefun@home.com
rebar@builder.org
freddie592@aol.com
strand@femtonet.com
xtheria@fast.com
george@whths.gov
bigboat@nautilus.com
travelr@trips.fun.com
boffo@snafu.com
yabbadabba@aol.com
morepower@scottie.com

4 Some spammers include in the e-mail a return address to which someone who no longer wants to receive spam can send a message and be taken off the spam list. When the remove message is received, a robot automatically takes the person off the list. However, spammers rarely do this because most people would opt not to be on the spam lists.

5 The spammer either buys the resulting e-mail list or compiles one of her own. The spammer uses the list, along with bulk mailing software, and sends a spam message to every person on the list. In the message might be a return address, Web site, or phone number where the receiver can get more information about the goods and services being sold.

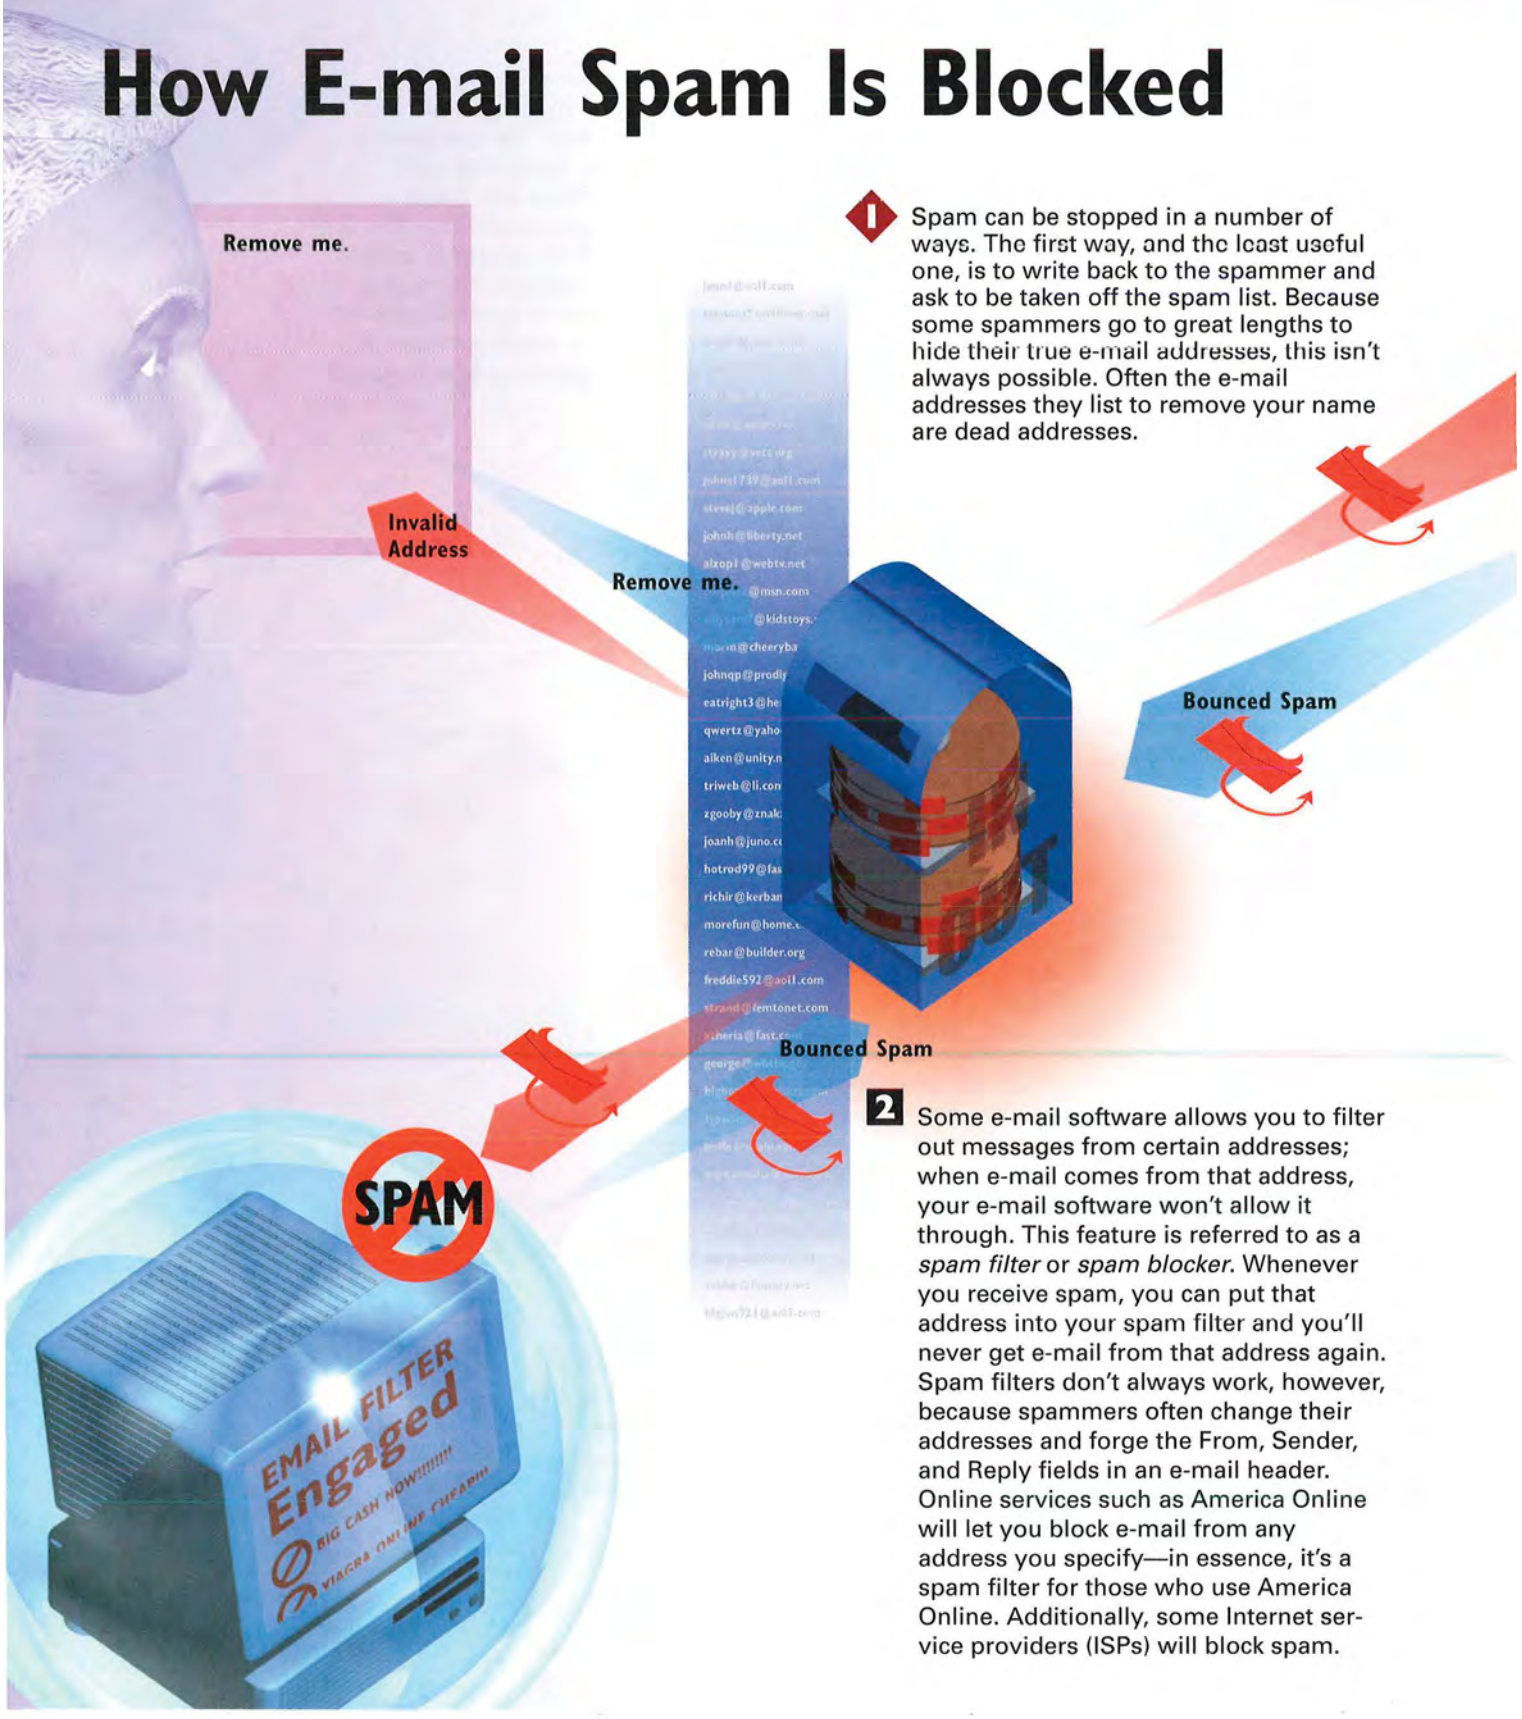




6 Spammers realize that spam offends most people, so the spammers go to great lengths to hide their true e-mail addresses. As one way of hiding their real e-mail addresses, they "forge" parts of the message header in the e-mail address, such as the From, Sender, and Reply fields so it appears that the e-mail has come from someone other than the spammer. Doing this is sometimes called *spamouflage*.

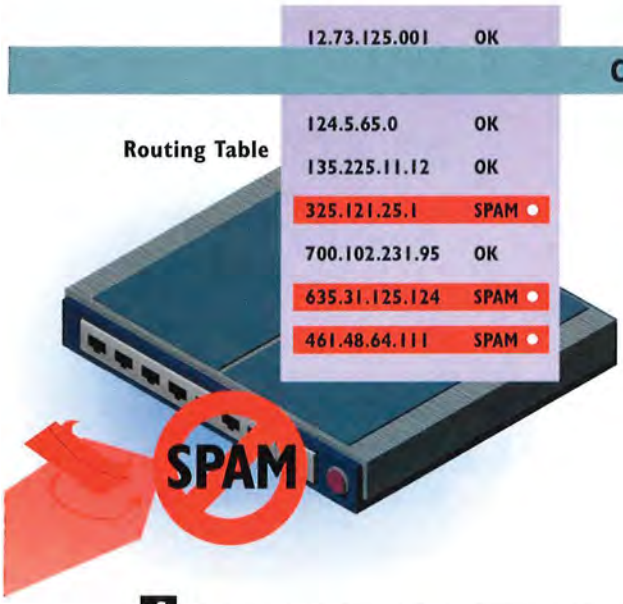
7 As a further way of hiding their true addresses, spammers relay their bulk spam to a server that is not associated with them and then have that server send out the bulk spam. Sometimes spammers have the bulk spam relayed among several different servers to make tracing who really sent the mail even more difficult.

How E-mail Spam Is Blocked



1 Spam can be stopped in a number of ways. The first way, and the least useful one, is to write back to the spammer and ask to be taken off the spam list. Because some spammers go to great lengths to hide their true e-mail addresses, this isn't always possible. Often the e-mail addresses they list to remove your name are dead addresses.

2 Some e-mail software allows you to filter out messages from certain addresses; when e-mail comes from that address, your e-mail software won't allow it through. This feature is referred to as a *spam filter* or *spam blocker*. Whenever you receive spam, you can put that address into your spam filter and you'll never get e-mail from that address again. Spam filters don't always work, however, because spammers often change their addresses and forge the From, Sender, and Reply fields in an e-mail header. Online services such as America Online will let you block e-mail from any address you specify—in essence, it's a spam filter for those who use America Online. Additionally, some Internet service providers (ISPs) will block spam.



3 ISPs and online service such as America Online can block spammers from sending bulk mail to their subscribers. A router examines all incoming mail to the ISP or online service. The router has been told that when e-mail comes from certain addresses, to block mail from getting into the network. These addresses are put in a routing table that can be changed whenever new spammers are found. ISPs have a variety of other ways for detecting spam as well.



4 Because spammers often change their addresses, using routing tables doesn't always work. Online services and ISPs have gone to court to ban spammers from sending e-mail to their customers. Although the law remains murky, in a number of instances, the courts have decided in favor of online services and the ISPs and have banned spammers from sending mail through them.



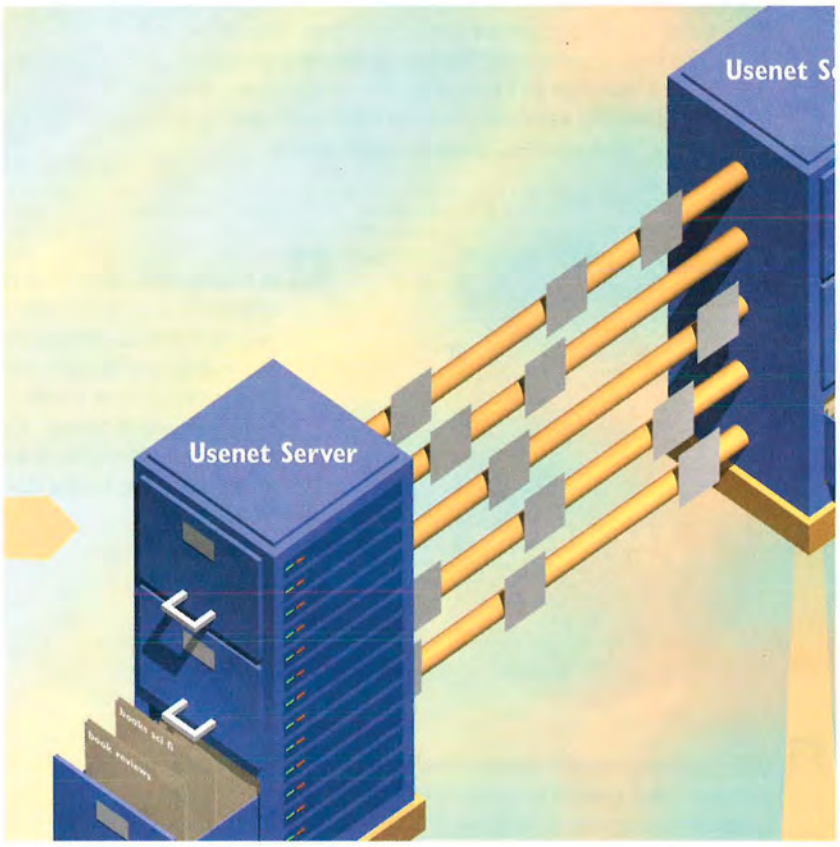
5 A number of laws and schemes have been proposed to regulate or outlaw spam. In one plan, every piece of spam would have to contain a specific piece of information in the message header, identifying it as unsolicited e-mail. In this way, people could set their spam filters to block every single piece of spam, filtering out that piece of information. Some laws have been proposed that would outlaw spam entirely, in the same way that junk faxes were banned.

6 One way to prevent being spammed is to be sure that your e-mail address isn't added to spam lists. To do this, when posting to Usenet newsgroups, edit your header so that it doesn't contain your e-mail address. You should also notify e-mail directories that you'd like to be taken off their lists. In this way, your e-mail address won't be harvested by robots, and you should get less spam.



CHAPTER

18 How Newsgroups Work



USENET, the world's biggest electronic discussion forum, provides a way for messages to be sent among computers across the entire Internet. People from all over the world participate in discussions on thousands of topics in specific areas of interest called *newsgroups*.

Hundreds of different major hierarchies of newsgroups exist, such as recreation (identified by the letters "rec") and computers (identified by the letters "comp"). Within these major hierarchies are subcategories (such as rec.arts) and further subcategories (such as rec.arts.books). Individual newsgroups can cover anything from movies to parenting, ecology, sports teams, clip art, and news about Usenet itself. Not all Internet sites carry all newsgroups. An administrator at each site decides which newsgroups to carry.

To participate in newsgroups, you need special software to read and respond to them. Readers are available for PC, Macintosh, and Unix computers. Online services such as CompuServe and America Online have their own proprietary software that lets you participate in newsgroups.

A good newsgroup reader lets you view the ongoing discussions as threads. *Threads* are ongoing conversations grouped by topic. So, for example, the rec.arts.books newsgroup might have many different threads going on at one time, each discussing a different book.

Many newsgroups have a list of Frequently Asked Questions, or FAQs (pronounced "facks"), associated with them. These FAQs answer common questions about the newsgroup. It's a good idea to read the FAQ before submitting questions to the newsgroup as a whole.

You participate in newsgroups by reading the messages and responding to them. There are moderated and unmoderated newsgroups. In a *moderated newsgroup*, each message goes to a human moderator. The moderator looks at the messages, ensuring they're appropriate for the group. If they are appropriate, the messages are posted. All messages sent to an *unmoderated newsgroup* are automatically posted.

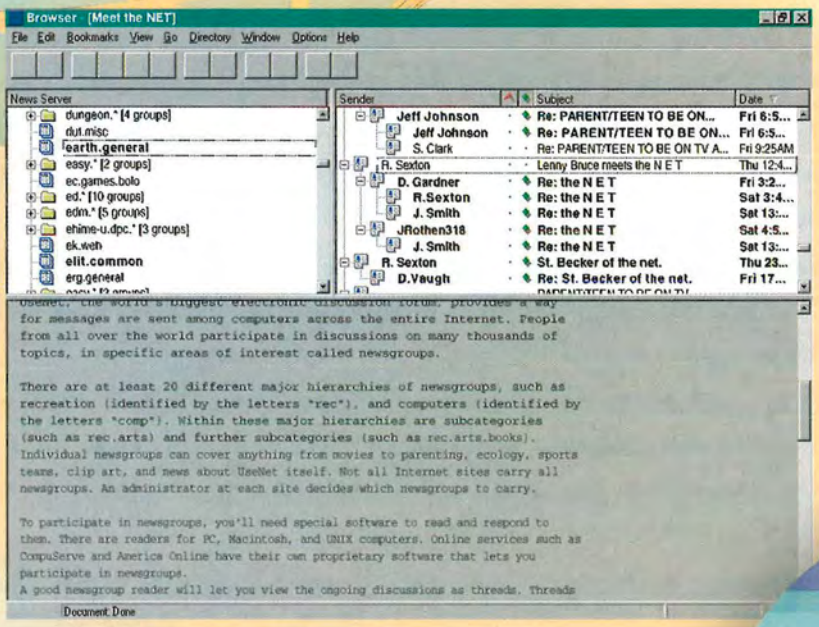
When messages are posted, Usenet servers distribute them to other sites that carry the newsgroup. A site usually carries only the most current messages; otherwise, they would soon run out of storage space. Some sites *archive*, or store, old discussions.

A convenient way to check newsgroups is to subscribe to those that interest you. That way, whenever you check the Usenet server, new messages in your subscribed newsgroup will be delivered to you. You can also cancel your subscription to a newsgroup if you are no longer interested in it. You can read newsgroups without subscribing to them; in that case, you must manually ask to read specific newsgroups instead of having it done automatically for you.

Binary files such as pictures and multimedia can be posted in newsgroups. These files must be specially encoded for them to be posted. To view them, you must transfer them to your computer and then unencode them with special software. A common encoding scheme used on newsgroups is called uuencode. Most Usenet software readers have unencoding built in.

How Usenet Works

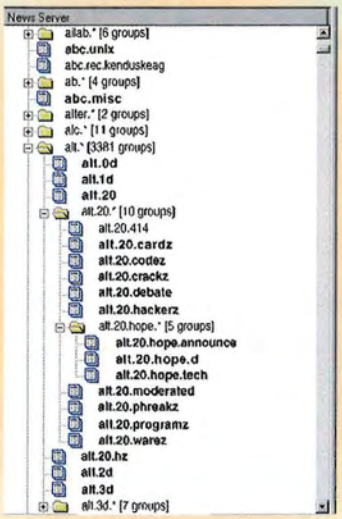
1 Usenet is a global bulletin board and discussion area. It collects messages about many thousands of topics into newsgroups, which are freewheeling discussion areas in which anyone can participate. Newsgroups can be found on many host computers across the Internet. Thousands of newsgroups exist that cover just about every topic you've ever imagined and many you probably haven't thought of.



2 To participate, people send and read messages that are posted to the newsgroup. Two kinds of newsgroups exist: moderated newsgroups and unmoderated newsgroups.

3 In a *moderated* newsgroup, a human moderator receives and reads all the messages for the newsgroup. The moderator then decides which messages should be posted. The messages to be posted are put on a newsgroup server. The other messages are discarded. In *unmoderated* newsgroups, all messages are put directly on the server. Newsgroups and all their messages are stored on a Usenet server. They are organized by broad categories and then broken down into specific topics.





6 Newsgroup reader software lets you read messages and respond to newsgroups. The software gives you ways to manage your newsgroups by also allowing you to subscribe to newsgroups, which means new messages will automatically be delivered to you when you check the server. You'll also be able to cancel your subscription.









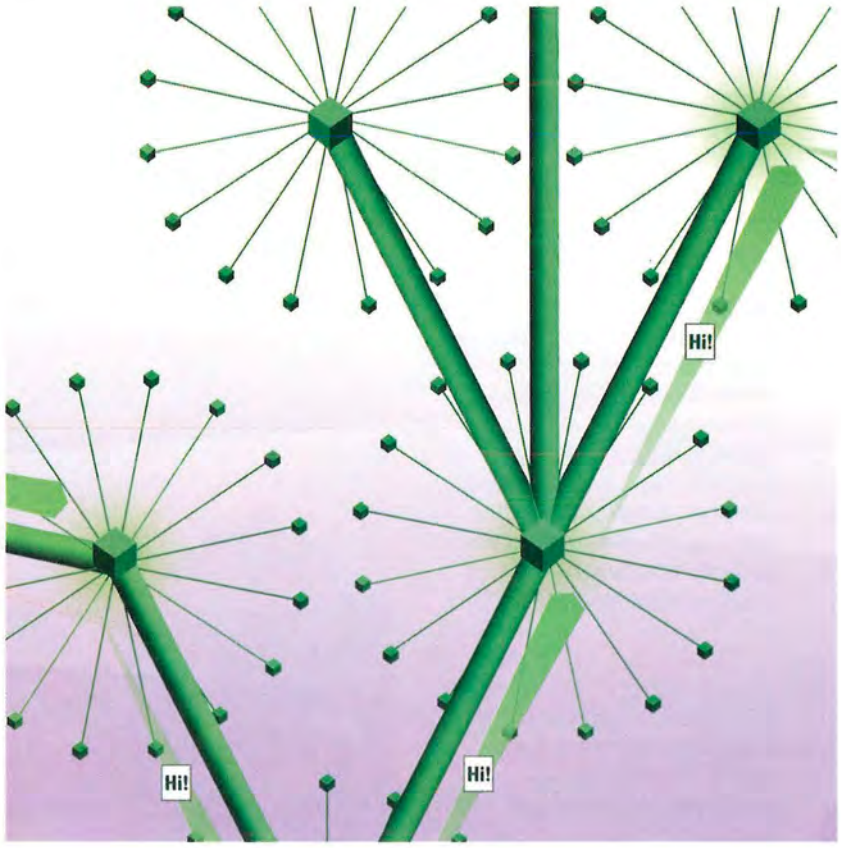
5 Pictures, multimedia files, and even executable programs can be posted in newsgroups for other people to see and use. However, because of the technology used in newsgroups, those files must be specially encoded to be posted. To view, play, or use the files, you first must transfer them to your own computer and then unencode them with special software. A common encoding scheme used on newsgroups is called uuencode. Versions of this encoding and decoding program are available that work on PC, Macintosh, and Unix computers. MIME is another encoding scheme.

4 Usenet servers communicate with one another so that all messages posted on one server are duplicated on the other servers. Although there are many Usenet servers, not all servers carry all newsgroups. Each site decides which newsgroups to carry.

CHAPTER

19

How Internet Chat and Instant Messaging Work



ONE of the most immediate ways to communicate with others via the Internet is to participate in live chat. Chat doesn't refer to people actually talking to each other and hearing each other's voices. Instead, it means that you hold live keyboard "conversations" with other people on the Internet—that is, you type words on your computer and other people on the Internet can see those words on their computers immediately, and vice versa. You can hold chats with many people simultaneously all over the world.

A number of ways are available to chat on the Internet, but one of the most popular ones is called *Internet relay chat (IRC)*. Every day, thousands of people all over the world hold conversations on many topics via IRC. Each topic is called a *channel*. When you join a channel, you can see what other people on the channel type on their keyboards. In turn, everyone in the channel can see whatever you type on your keyboard. You can also hold individual side conversations with someone. Channels live on various servers around the world. Some servers have only a few channels, and others have many of them.

IRC has facilitated communications during natural disasters, wars, and other crises. In 1993, for example, during the attempted Communist coup in Russia when Russian legislators barricaded themselves inside the Parliament building, an IRC "news channel" was set up for relaying real-time, first-person accounts of the events taking place.

IRC follows a client/server model, which means that both client and server software are required to use it. Many IRC clients are available for many types of computers, so whether you have a PC, Macintosh, or Unix workstation, you can use IRC.

Your IRC client communicates with an IRC server on the Internet. You log on to a server using the client and select a channel on which you want to chat. When you type words on your keyboard, they are sent to the server. The server is part of a global IRC server network. The server sends your message to other servers, which, in turn, send your messages to people who are part of your channel. They can then see and respond to your message.

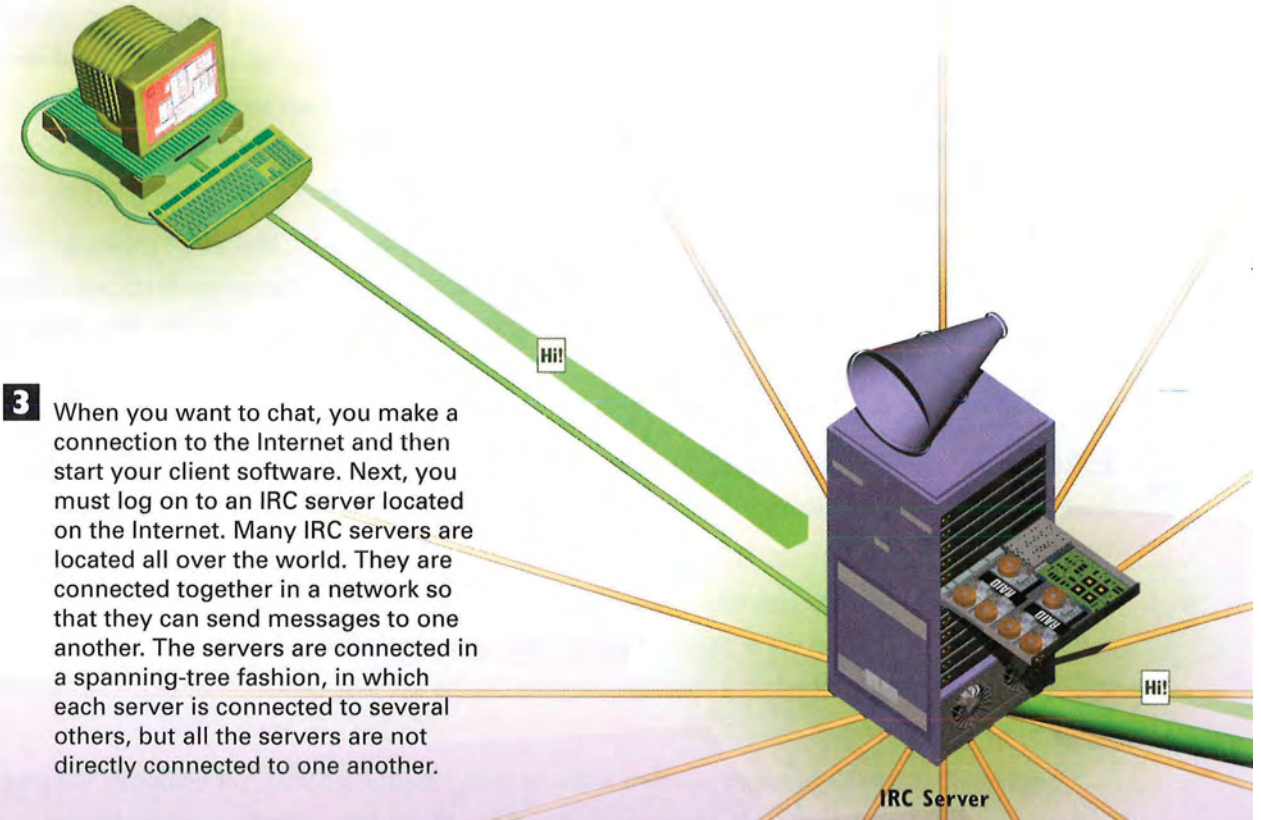
You can chat on the Internet in other ways, as well. Many Web sites, for example, use proprietary chat software that doesn't use the IRC protocol but enables you to chat when you're on the site.

Another type of chat is called *instant messaging*. In instant messaging, you communicate privately, one on one, with another person. You can create special lists so you're informed when your "buddies" come online ready to chat, and they're informed when you come online.

How IRC Works

1 Internet relay chat (IRC) is a way for people all over the world to "chat" with one another by using their keyboards. The typed words are instantly relayed to computers all over the world, where recipients can read them. This process occurs in real-time, so everyone sees the words as people type them.

2 IRC runs on a client/server model; therefore, to use it, you need client software on your computer. Many IRC clients are available for PCs, Macintoshes, Unix workstations, and other kinds of computers.

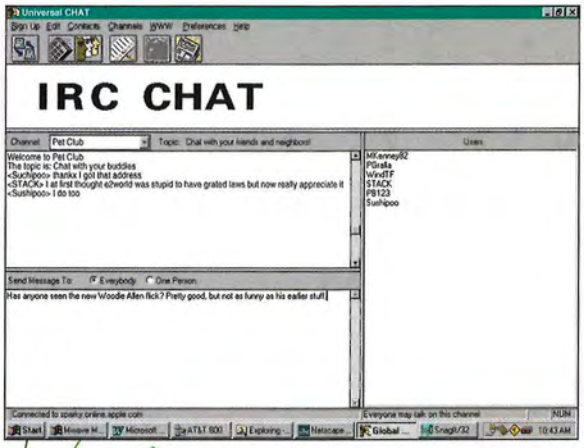


3 When you want to chat, you make a connection to the Internet and then start your client software. Next, you must log on to an IRC server located on the Internet. Many IRC servers are located all over the world. They are connected together in a network so that they can send messages to one another. The servers are connected in a spanning-tree fashion, in which each server is connected to several others, but all the servers are not directly connected to one another.



4 When you connect to a server, you choose a specific channel to join and choose a username to identify yourself when you chat. Many channels are available that cover various topics.

5 After you've joined a channel, you can see the conversations that are taking place. To join the conversation, type your message at your keyboard.



6 Your message is sent from the client software on your PC to the IRC server to which you're connected.

8 Each server sends your message to the client software of the people connected to the channel at each server. People can read and respond to your message at their computers.

7 The message is then sent from your server to other servers where people on your channel are logged on. The message gets sent from server to server. Under the spanning-tree server structure, the message always takes the shortest route through the network to reach its final destination.

How Instant Messaging Works

1 The Internet version of America Online's Instant Messenger (AIM), one of the most popular instant messaging systems, runs as a piece of client software on your computer. To use it, you must be connected to the Internet. When you run the software, it opens up a TCP connection to an Instant Messenger login server. The software sends your screen name and password over the connection to log you into the server.



3 Instant message software includes buddy list capabilities. That means you keep a list of people you want send instant messages to, and when they come online, you are notified so you can send instant messages to and receive instant messages from them. You create a buddy list in your AIM software by adding your buddies' screen names to it.

4 When you establish a connection with the AIM server, your client software sends a list of your buddies to the server. The server checks to see whether any of the buddies are online—and it continues to do that for as long as you run the software on your computer. If you change the list of buddies during your session, that information is sent to the server as well so it can keep track of new buddies or ignore buddies you've deleted from your list.



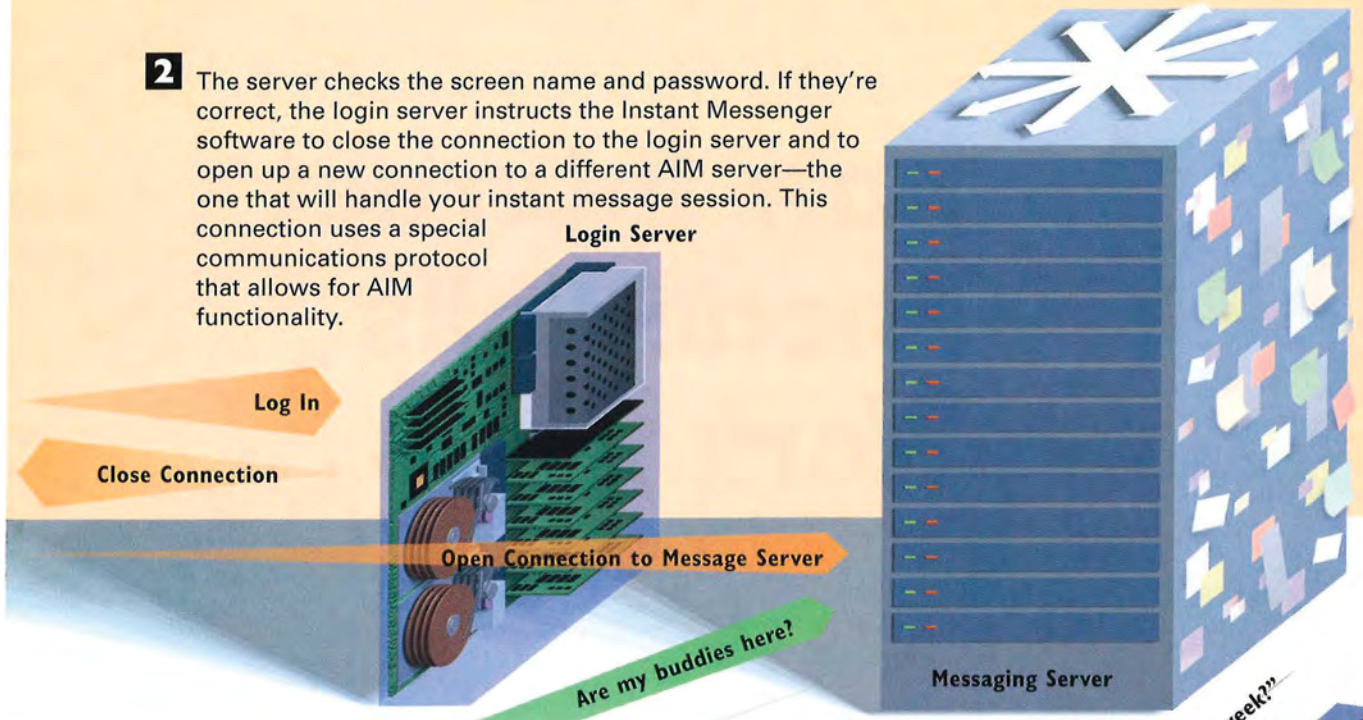
"Cool Lizard is online."

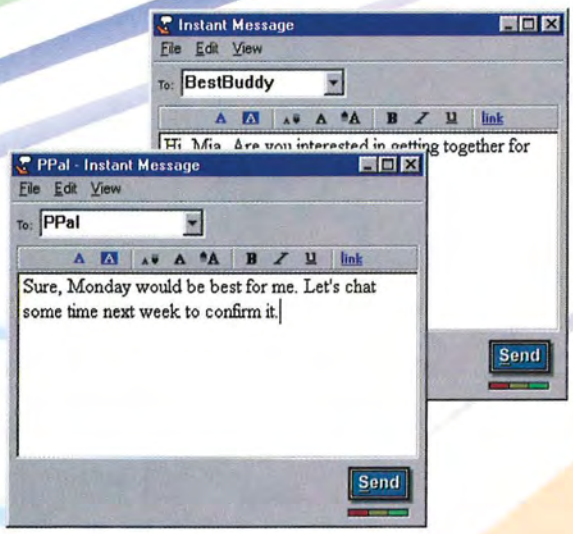
"Hi Mia. Are you interested in getting together next week?"

"Sure. Monday would be best for me. Let's chat later to confirm it."

5 When any of your buddies runs AIM and logs in, your client software is told that he's online, and you get a notice that he's online. You can now send messages to and receive instant messages from him.

2 The server checks the screen name and password. If they're correct, the login server instructs the Instant Messenger software to close the connection to the login server and to open up a new connection to a different AIM server—the one that will handle your instant message session. This connection uses a special communications protocol that allows for AIM functionality.





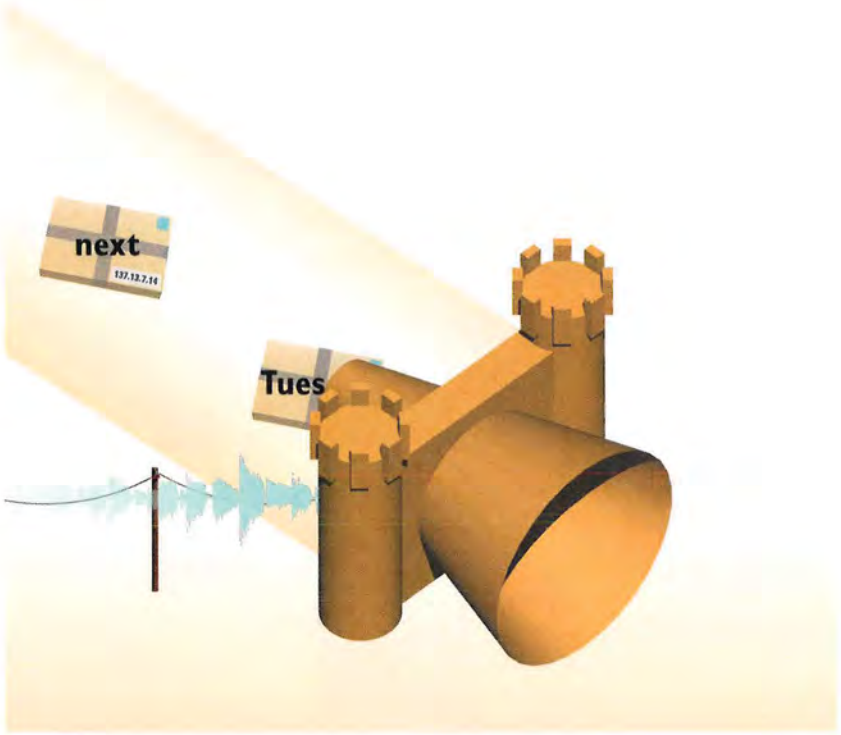
"Hi Mia. Are you interested in getting together next week?"

"Sure. Monday would be best for me. Let's chat later to confirm it."

6 When you send an instant message, your message goes to the AIM server. The AIM server then routes the message to the proper person. Similarly, when he responds, his message first goes to the server and then to you. The Yahoo! Pager software, in contrast, enables people to chat directly with each other without having to go through a server.

CHAPTER

20 How Internet Phone Calls Work



THE Internet has pioneered many new ways to communicate, such as e-mail, live chat, and newsgroups. But it can be used to enable some old-fashioned communications as well: You can make telephone calls using the Internet. When you do so, the sound of your voice and that of the person you're talking to are broken down into packets. Those packets are delivered using the Internet's TCP/IP protocol.

You can make telephone calls over the Internet in two ways. In the first way, you use your computer and special hardware and software to make calls, so you communicate through your PC. In the second way, often referred to as *Internet telephony*, or *voice over IP*, you make a phone call as you normally do on your telephone, except your call is routed over the Internet rather than through the normal phone service. This usually cuts the cost of making long-distance phone calls.

The revolutionary part of using the Internet to make phone calls from your PC is the price: It's free. You pay only for your Internet connection, as if you were browsing the Web or sending e-mail. You don't actually have to pay for the phone call itself. In addition, you can make calls anywhere in the world. However, you will be able to make telephone calls to and receive calls only from someone who has an Internet address, so you won't be able to replace your telephone by using the Internet. A number of competing products enable you to talk over the Internet, but as yet, they don't communicate with one another. No standard way of making Internet telephone calls exists, so you can talk only to people who use the exact same software you do for making and receiving phone calls.

Many schemes and software packages enable you to make phone calls from your PC over the Internet. You don't actually use your telephone with any of them. Instead, you speak into a microphone attached to your computer and listen through speakers and a sound card.

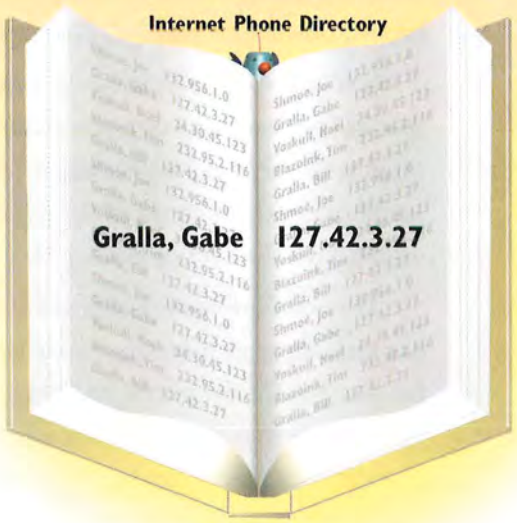
Although each company's software works somewhat differently, in general, they use similar ways of enabling people to make phone calls over the Internet. To make a phone call, first you must know someone's IP address. (You can consult directories that function as telephone books to get this information.) After you find the person's Internet address, you can connect to him by running special software and then double-clicking his name. (The person on the other end must be running the same software you are using.) A message goes out to him over the Internet to see whether he's available to talk.

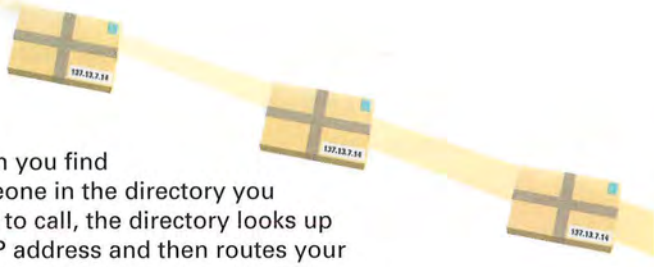
When you make phone calls over the Internet using your normal phone, things work differently. You make the call as you would any other phone call. After you make the call, however, your voice is digitized and converted into IP packets, and then it's delivered using the same TCP/IP technology the rest of the Internet uses. On the other end of the call, the digitized voice data is converted and can be listened to over the phone.

How PC Phone Calls on the Internet Work

1 When you connect to the Internet and want to send or receive telephone calls with special phone software, you connect to a special directory. The directory contains a list of all the people who can connect to the Internet using that particular phone software.

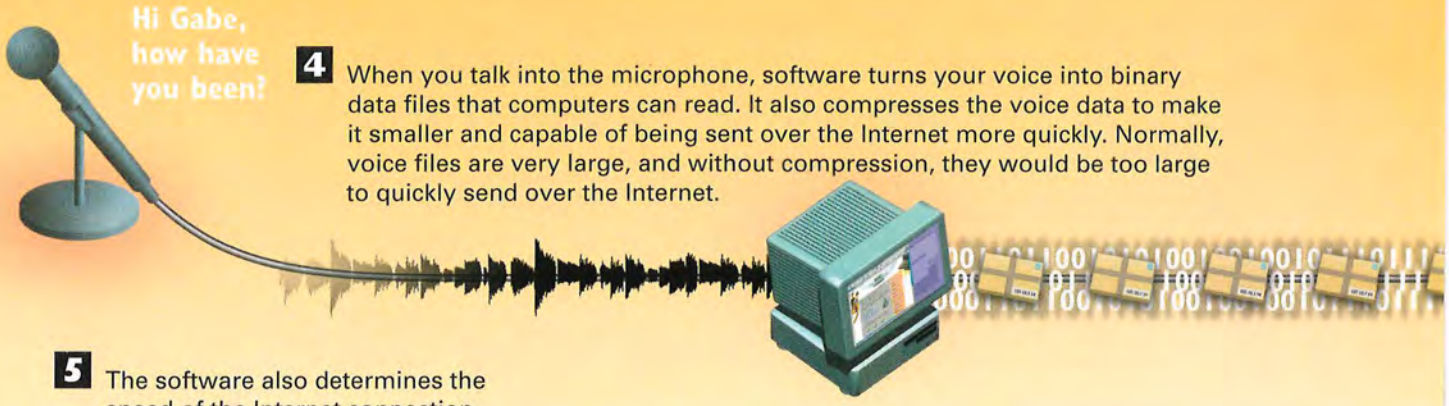






2 You can look up someone's name in the directory. The directory matches a person's name (such as Gabe Gralla) with his Internet IP address (such as `192.168.1.1`). It's important that the directory match a name with an IP address because someone potentially can have a different IP address every time he logs on to the Internet. The directory matches names to the most current IP addresses.

3 When you find someone in the directory you want to call, the directory looks up the IP address and then routes your call directly to that computer. The phone "rings" at the person's computer, and he picks up the receiver. The two of you can now talk, using the sound cards and microphones attached to your computers.

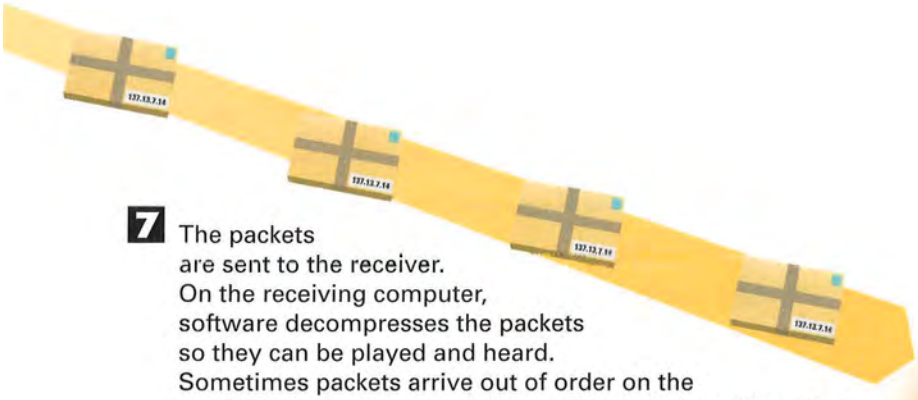


4 When you talk into the microphone, software turns your voice into binary data files that computers can read. It also compresses the voice data to make it smaller and capable of being sent over the Internet more quickly. Normally, voice files are very large, and without compression, they would be too large to quickly send over the Internet.

5 The software also determines the speed of the Internet connection. If the connection is a high-speed connection, it creates voice files of high quality. If it's a lower-speed connection, it creates files of a lesser sound quality because at lower speeds it would take too long to send voice files, even after they had been compressed.



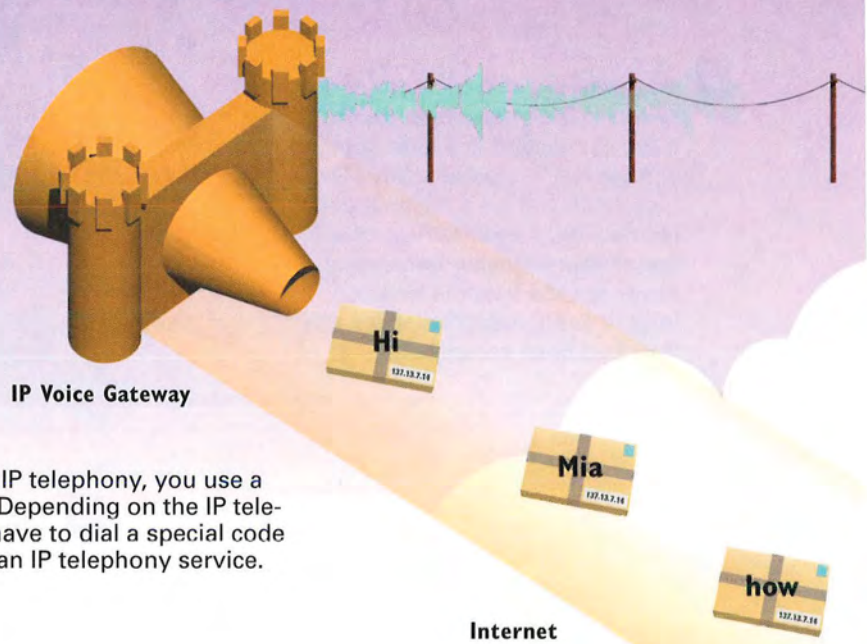
6 The software breaks up the voice files into a series of packets to be sent over the Internet using the Internet's TCP/IP protocol.



7 The packets are sent to the receiver. On the receiving computer, software decompresses the packets so they can be played and heard. Sometimes packets arrive out of order on the receiving end, so some of them will be missing. When that happens, the software, based on the received packets, is capable of estimating what the sounds are in the missing packet that has yet to be received. When the missing packet arrives, the software knows that it has already re-created what that packet would have sounded like. The software then discards the packet.

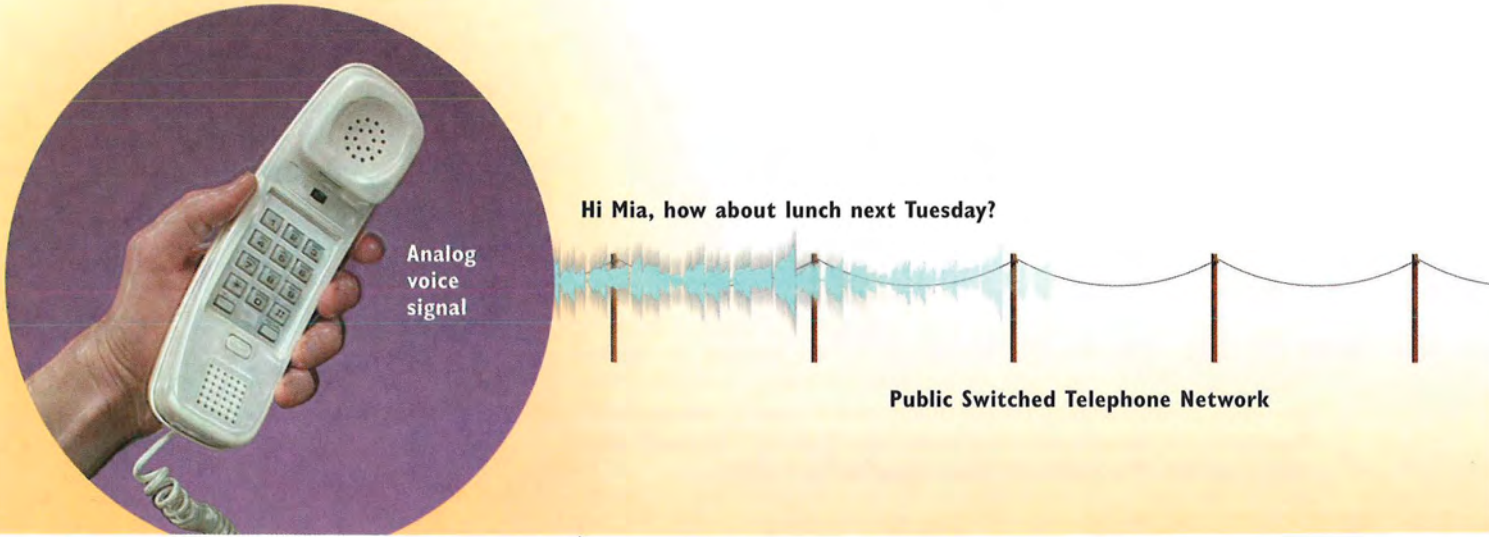


How Internet Telephony Works

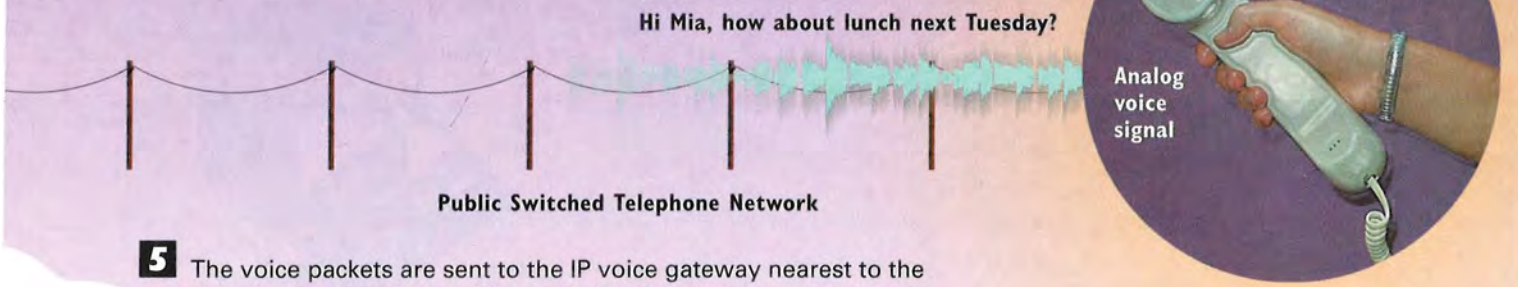


1 When you make a phone call using IP telephony, you use a normal telephone, not a computer. Depending on the IP telephony service you use, you might have to dial a special code first, as a way of routing the call to an IP telephony service.

2 The first part of the call goes over the normal Public Switched Telephone Network (PSTN), as does any other call; at this point, it is no different in any way.

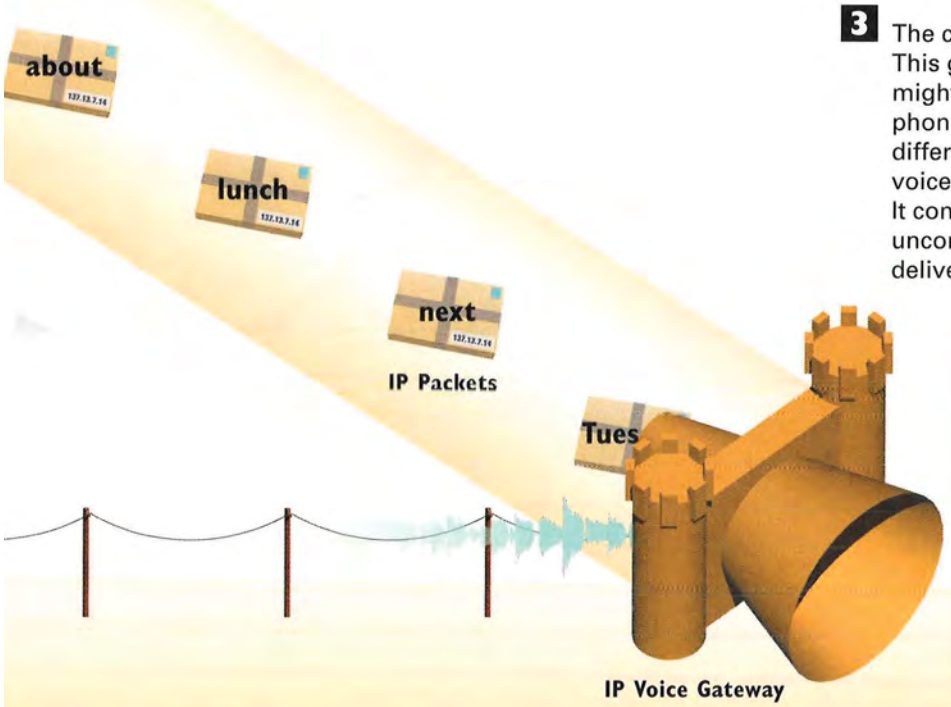


6 The call travels on the PSTN the same as any other telephone call. The person receiving the call picks up the phone and talks, the same as with any other telephone conversation. When the receiving person speaks into the phone, the voice is sent back through the PSTN to the IP voice gateway and across the Internet in the same way that the call was initially made.



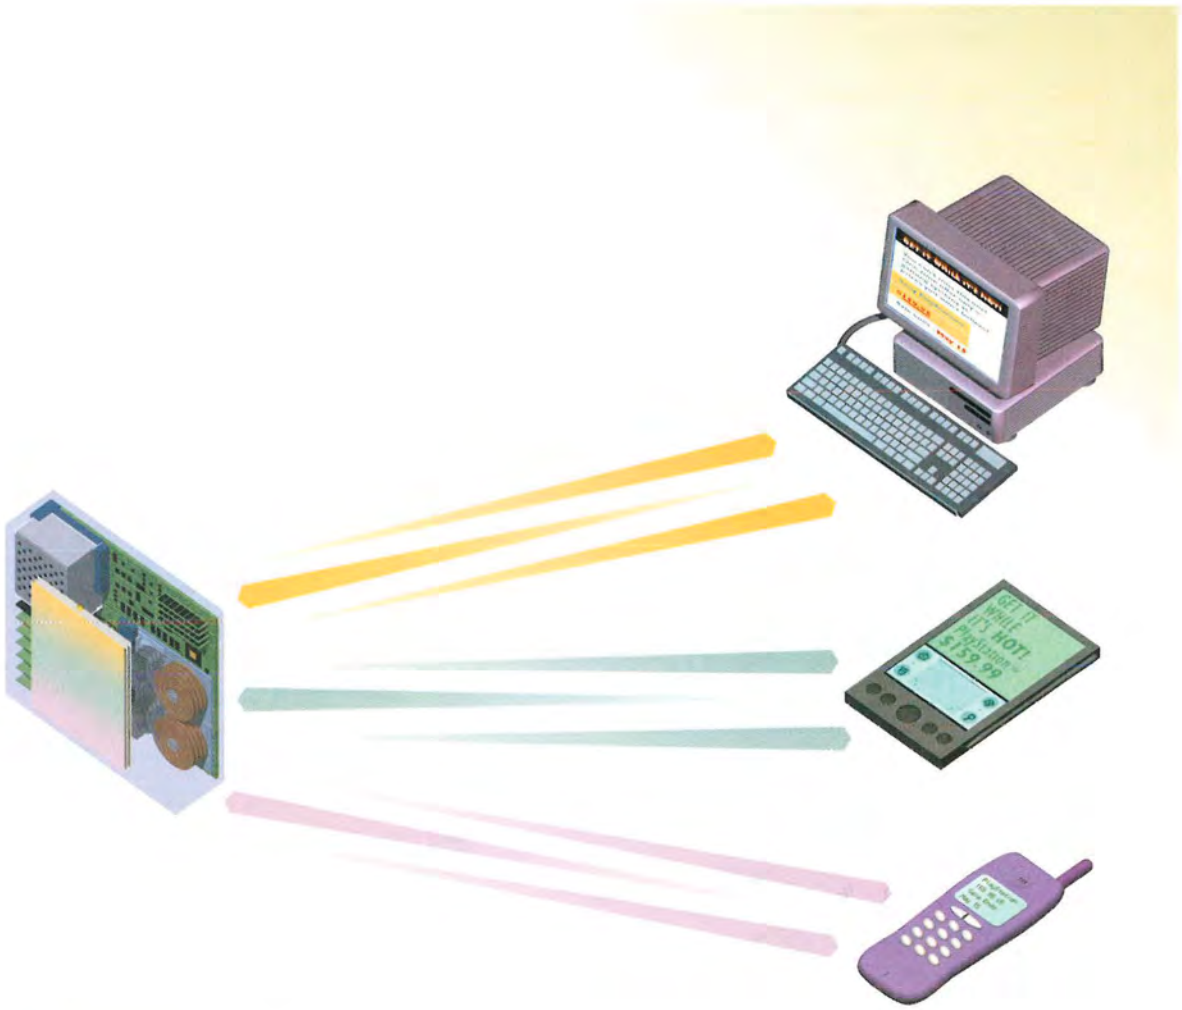
5 The voice packets are sent to the IP voice gateway nearest to the phone call's destination. The receiving IP voice gateway takes the voice packets, uncompresses them, converts them back into their original form, and then sends them through the PSTN.

4 The gateway breaks up the compressed, digitized voice signal into IP packets. These packets are sent out across the Internet the same as any other IP packets, using the Internet's TCP/IP protocol.



3 The call is sent to a special IP voice gateway. This gateway—a server and server software—might be located at the branch of the telephone company or might instead be in a different location. The gateway converts the voice signal to digital data and compresses it. It compresses it because the data files of an uncompressed voice call could be too large to deliver in a timely fashion across the Internet.

IP (Internet Protocol) Telephony refers to telephone calls that are made over normal telephones but are delivered, in part, through the Internet instead of solely through a voice network, commonly called the Public Switched Telephone Network, or PSTN. Because the calls are routed through the Internet, they can be delivered at less cost than normal phone calls. But the voice quality might not be as good as calls delivered over the normal phone network.


GET IT WHILE IT'S HOT!
You can't miss this one!
One-time offer only –
gaming systems at
prices you won't believe!
Sony PlayStations

PART 4

USING THE WORLD WIDE WEB

Chapter 21: How Web Pages Work
126

Chapter 22: How Web Browsers Work
134

Chapter 23: How Markup Languages Work
140

Chapter 24: How Hypertext Works
148

Chapter 25: How URLs Work
152

Chapter 26: How Image Maps and Interactive Forms Work
156

Chapter 27: How Web Host Servers Work
160

Chapter 28: How Web Sites Work with Databases
164

Chapter 29: How Internet-Based Software Works
168

WHEN many people use the word "Internet," they are really talking about the World Wide Web. The Web is the most interesting, most innovative, most visible, and fastest growing part of the Internet. To a great degree, the explosive growth of the Web has been what has fueled the enormous amount of interest in the Internet in the past several years. When people refer to "surfing the Net," they're usually talking about using the World Wide Web.

This section of the book looks in great detail at how the World Wide Web works, from such basic technologies as how Web pages work to more advanced features, such as image-maps and the way the Web is becoming increasingly integrated directly into your computer. You'll learn what the Web comprises and how a Web browser works and thoroughly investigate URLs (uniform resource locators), as well as many other aspects of the Web and Web browsers.

Chapter 21, "How Web Pages Work," examines the most basic part of the Web. It starts by covering the general technology of the Web and describes how the Web works. Web pages are, in essence, multimedia publications that can include music, audio, video, and animation, as well as graphics and text. Web pages are connected via hypertext that enables you to jump from any page to any other page, and to graphics, binary files, multimedia files, and other Internet resources. To jump from one page to another, or to another resource, you merely click a hypertext link.

The chapter details how all that works and looks at the client/server model of the Web. It shows what happens behind the scenes when you type a URL into your browser—how that information gets routed properly so you visit the Web site that interests you. The chapter also delves into how Web pages are organized on a site and how an entire Web site works as a unit to deliver its information to you.

Chapter 22, "How Web Browsers Work," examines browsers. Web browsers are pieces of software that interpret the language of the Web, Hypertext Markup Language (HTML), and then display those results on your computer. The chapter begins by delving into how browsers do that job.

Browsers, however, now go far beyond merely allowing you to visit the Web. Both Netscape Communicator and Internet Explorer have grown to become full-blown application suites. Communicator, for example, includes the Navigator browser along with many other programs for communicating with others on the Internet. It includes a news reader, an e-mail package, conferencing and whiteboard software, and even a powerful editing tool that enables you to create and publish your own Web pages.

This chapter discusses one of the most common, and annoying, facts about the Web—the sometimes incomprehensible error messages you receive when you can't get to a Web site or Web page. You'll see what those messages mean and understand why you can't visit a Web page or site when you get one of them.

As mentioned earlier, HTML is the language of the Web. Chapter 23, "How Markup Languages Work," explains how HTML forms the building blocks for creating Web pages. The language is an essential set of directions that tells your browser how to display and manage a Web document. The chapter shows in detail how all that works.

Chapter 24, "How Hypertext Works," looks at hyperlinks, which set the Web apart from most other portions of the Internet. This chapter describes the various ways documents can link to each other on the Web and explains concepts such as relative links and absolute links.

Chapter 25, "How URLs Work," looks in great detail at URLs, the addresses you type into your browser to visit a Web site. You'll look at the underlying structure of a URL and gain a better understanding of how a URL is put together and what it can tell you about the site you're visiting. You'll also see precisely how URLs help retrieve documents from the Web.

Chapter 26, "How Image Maps and Interactive Forms Work," looks at two technologies you use every day on the Web without realizing it: imagemaps and interactive forms. Image maps aren't maps in a traditional sense. Instead, they're graphics with URLs embedded inside them. When you click one part of the graphic, you're sent to one site, and when you click another, you're sent to a different site. An imagemap might be a picture of a house, for example, and when you click the living room, you're sent to an entertainment site; when you click a home office, you're sent to a business site.

Interactive forms are the forms you fill out on the Web for doing things such as registering at a site or sending information about yourself before you're allowed to download a particular piece of software free of charge.

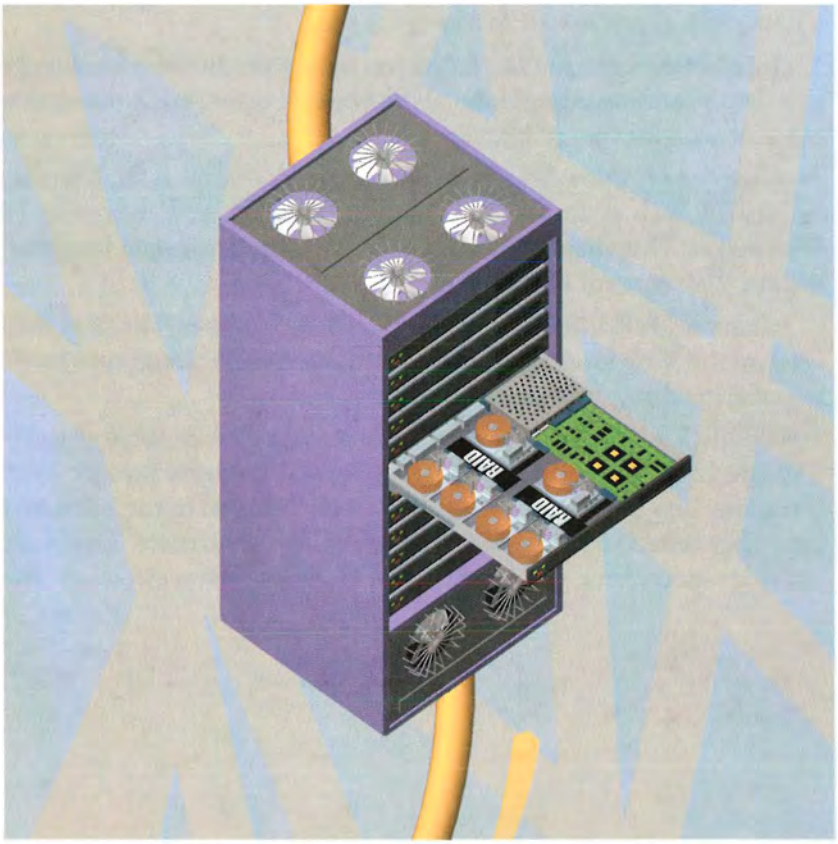
Chapter 27, "How Web Host Servers Work," looks at Web server software. As mentioned before, the Web works on a client/server model. Your Web browser (the client) contacts the Web server. This chapter explains how Web server software interacts with your browser to deliver Web content to you.

Chapter 28, "How Web Sites Work with Databases," looks at databases. Databases are used on the Web for many things. Web indexes and search sites such as Yahoo! are, in essence, databases that interact with the Web.

Finally, Chapter 29, "How Internet-Based Software Works," shows you how Web-based software (called ASPs for Application Service Providers) works. This is one of the more intriguing uses of the Internet because it enables you to run software not from your own computer, but instead from a distant computer on the Internet. That way, you never need to upgrade—every time you connect to that software, you always get the newest version.

CHAPTER

21 How Web Pages Work



THE World Wide Web is the fastest growing and, in many ways, the most exciting and intriguing part of the Internet. When people refer to "surfing the Net," more often than not they're talking about using the World Wide Web.

As its name implies, the World Wide Web is a globally connected network. The Web contains many things, but what makes it so fascinating to so many are the Web "pages" that incorporate text, graphics, sound, animation, and other multimedia elements. In essence, each page is an interactive multimedia publication that can include videos and music as well as graphics and text.

Pages are connected to one another using *hypertext* that allows you to move from any page to any other page, and to graphics, binary files, multimedia files, as well as any Internet resource. To jump from one page to another, you click a hypertext link—a link that connects Web pages and resources.

The Web operates on a client/server model. You run Web *client* browser software, such as Netscape Navigator or Microsoft's Internet Explorer, on your computer. That client contacts a Web *server* and requests information or resources. The Web server locates and then sends the information to the Web browser, which displays the results.

Pages on the Web are built using a markup language called Hypertext Markup Language (HTML). The language contains commands that tell your browser how to display text, graphics, and multimedia files. It also contains commands for linking the page to other pages and to other Internet resources.

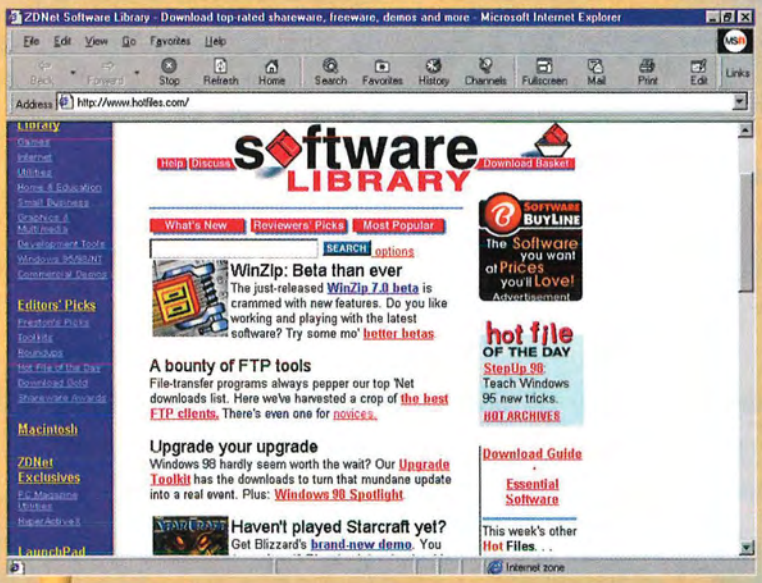
The term *home page* often is used to refer to the first, or top, page in a collection of pages that make up a Web site. This is to distinguish it from the many pages collected together as a single "package" of sorts that often make up Web sites. The home page is similar to a magazine cover or the front page of a newspaper. Usually, it acts as an introduction to the site, explaining its purpose and describing the information found on other pages throughout the site. In this way, the home page often acts as the table of contents for the rest of the site.

In general, Web sites use three types of organizational structures to organize their pages. In a *tree structure*, a pyramid or an outline format enables users to easily navigate through the site and find the information they want. In a *linear structure*, one page leads to the next, which then leads to the next, and so on, in a straight line. Finally, in a *random structure*, pages are connected to one another seemingly at random.

The last illustration in this chapter shows how someone might build her own pages using an HTML editor. After you create your pages, you use FTP software to post them onto a small portion of a Web server. You can either rent the server space from a local Internet service provider (ISP) or set up your own Web server. (Some sites also let you build Web pages by using the site's built-in tools. In that case, you wouldn't build a page from scratch and send it via FTP—instead, you'd build the page using online tools provided by the site.

How the World Wide Web Works

1 The World Wide Web is the fastest growing and most innovative part of the Internet. When you browse the Web, you view multimedia pages composed of text, graphics, sound, and video. The Web uses hypertext links that allow you to jump from one place to another on the Web. The language that enables you to use hypertext links and to view Web pages is called Hypertext Markup Language, more commonly known as HTML.



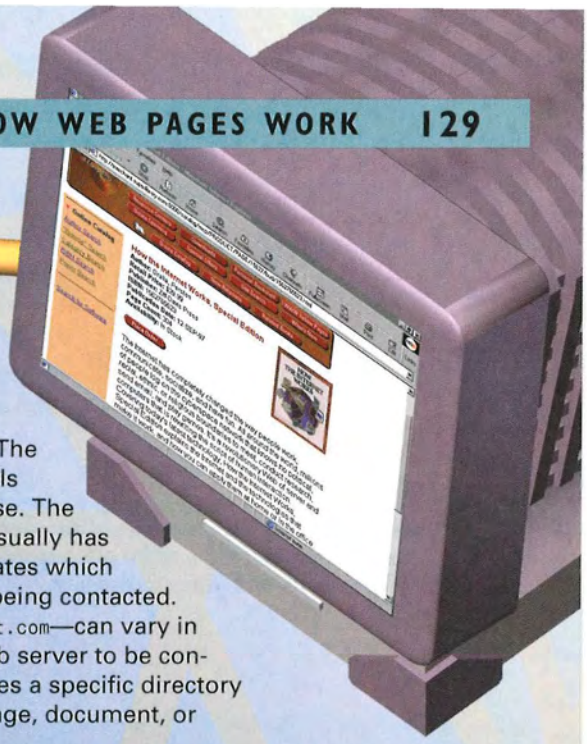
2 The Web works on a client/server model in which client software—known as a *Web browser*—runs on a local computer. The server software runs on a Web host. To use the Web, you first make an Internet connection and then launch your Web browser.

3 In a Web browser, you type the URL for a location you want to visit or click a link that sends you to the desired location. The names for Web locations are uniform resource locators (URLs). Your Web browser sends the URL request using Hypertext Transfer Protocol (HTTP), which defines the way the Web browser and Web server communicate with each other.



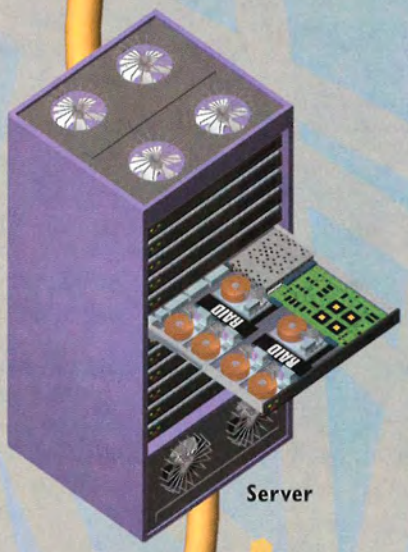
4 When the server finds the requested home page, document, or object, it sends that home page, document, or object back to the Web browser client. The information is then displayed on the computer screen in the Web browser. When the page is sent from the server, the HTTP connection is closed and can be reopened.

`http://www.zdnet.com/downloads`



http://www.zdnet.com/downloads

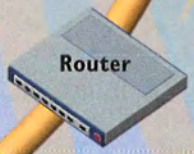
7 URLs contain several parts. The first part—the `http://`—details which Internet protocol to use. The second part—the part that usually has a `www` in it—sometimes indicates which type of Internet resource is being contacted. The third part—such as `zdnet.com`—can vary in length and identifies the Web server to be contacted. The final part identifies a specific directory on the server and a home page, document, or other Internet object.



Server

http://www.zdnet.com/downloads

6 The request is sent to the Internet. Internet routers examine the request to determine to which server to send the request. The information just to the right of the `http://` in the URL tells the Internet on which Web server the requested information can be found. Routers send the request to that Web server.

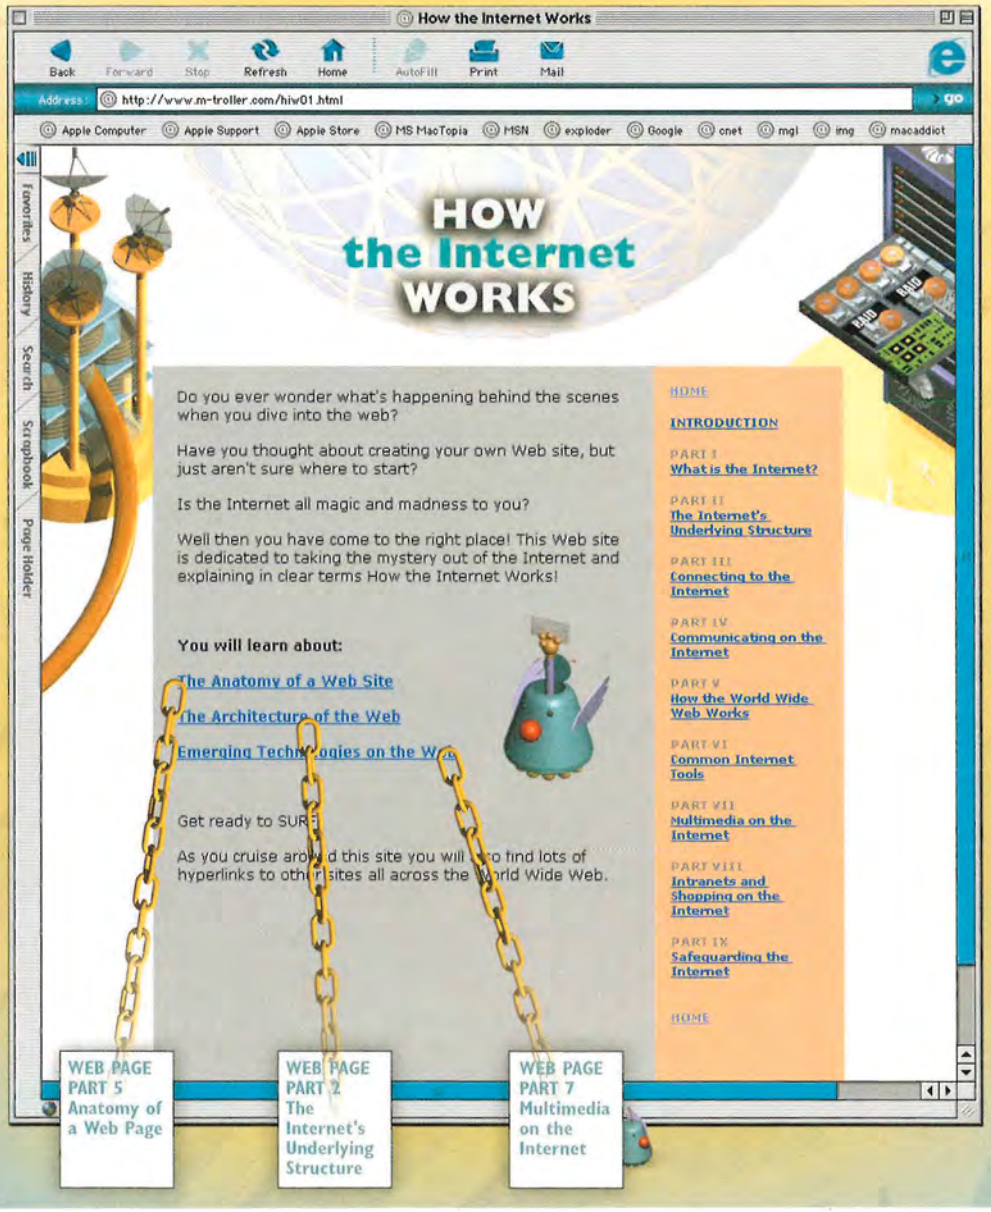


Router

5 The Web server receives the request using the HTTP protocol. The request tells the server which specific document is being requested and where that document is located.

How Web Pages Are Organized on a Web Site

1 The *home page* is the first or top page of any Web site. A site can be just one page or can comprise dozens or even hundreds of pages. In the latter case, the home page acts as a table of contents to organize the site and help users find information available on the site.



2 Underlined or highlighted hyperlink text is often embedded in the home page. The hyperlinks serve to connect the top page with other pages throughout the site.

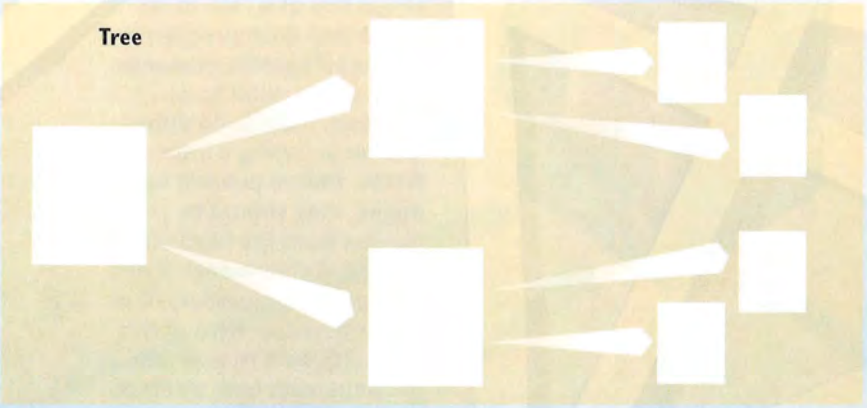
3 Related documents residing together on a Web host computer make up a Web site. However, a single server can host multiple Web sites, each contained in a separate area or directory, much like a hard drive can accommodate multiple directories. Some Web sites are so large and heavily trafficked that they can't fit on a single server and require multiple servers.

4 Good Web design principles suggest that pages throughout a site link back to the home page. This approach enables users to always find their way back to the top of a site to navigate in other directions.

5 Documents within a site can be linked to any other document in the site—and even to documents on other sites. Most Web sites, however, are designed in a pyramid or an outline structure that gives users a visual model to understand how information is arranged and indicates how to find and navigate through the site's documents.



6 Web sites typically are organized in one of three ways. The first is an outline or a tree structure, which arranges information hierarchically, moving from general information to more specific data.



7 The second organizational method is linear, in which one page leads to the next, which leads to the next one, and so on.



8 The third organizational structure is really a lack of structure, in which pages are connected to one another seemingly at random. (It's this last structure, though, that makes it clear why the Web is called the Web.)

How Web Sites Are Constructed

1 First, the raw materials are collected for the Web page. Content can be any number of things: family photos, poems, personal opinions, text of any kind, and even sound clips and movies. Text can easily be compiled with any word processor, such as Microsoft Word.



Text



Graphics



Multimedia



HTML Page

2 A page can come to life with a few well-chosen images. In designing a family home page, for example, a scanned photo might work well on the Web page. Icons or colored bullets spruce up a text list, and navigational icons, such as arrows or pointers, let a user move logically around the site. A whole range of shareware image editors is available on the Internet. *Shareware* is software you can try out for free, and for which you must pay only if you use it after a certain length of time.



3 Software is available that facilitates the arrangement of text, images, and other media without requiring a single line of HTML to be typed; one such program is Adobe's PageMill. However, if the pages need to be tweaked, there is no substitute for knowing a little HTML. Before posting new pages, they should be previewed from the hard disk using a Web browser. If any changes are necessary, it is easier to make them at this stage. It's best to preview the page with both Internet Explorer and Netscape Navigator because the same page can look different in each browser.



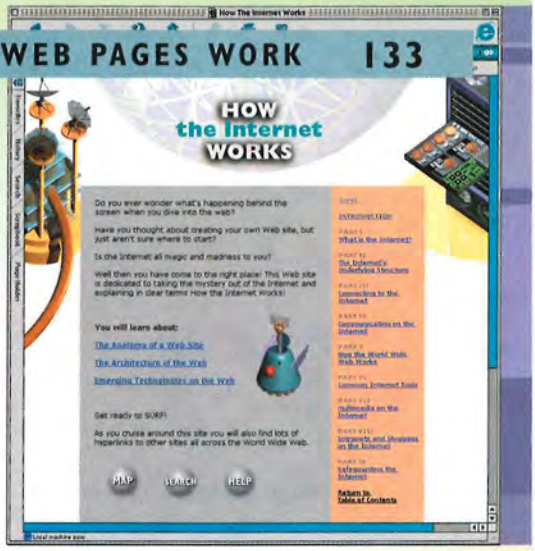
Text Editor



Image Editor

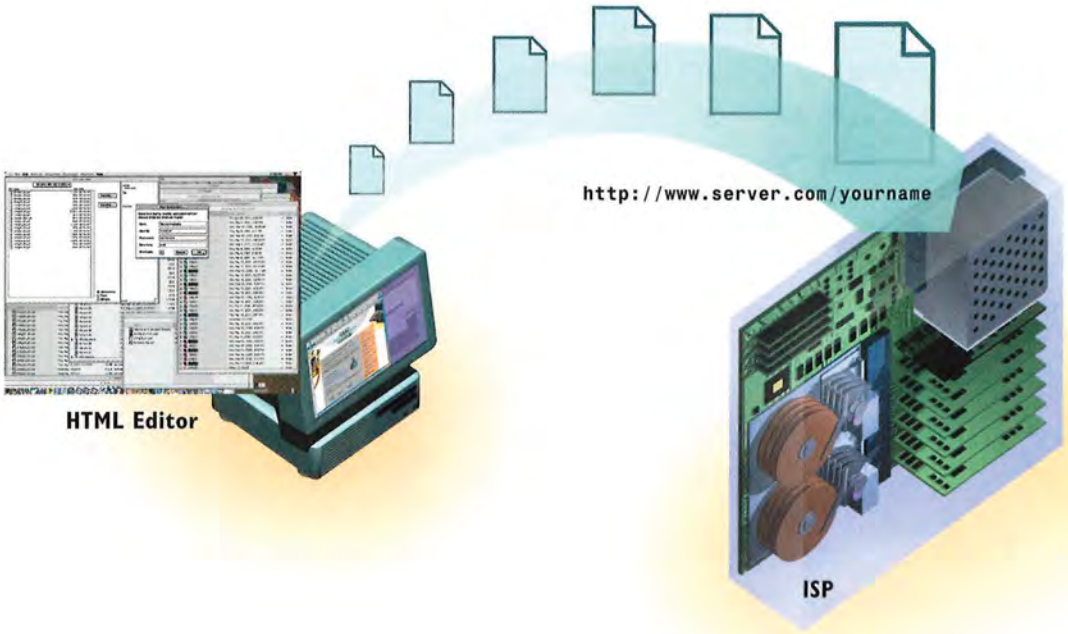


Media Editor



6 An FTP program is used to log on to the ISP's Web server; many good ones are available as shareware on the Internet. The ISP issues a username and password at sign-up time, which are used to log on to the directory where the HTML page, graphics, and other media will be placed. After the various elements are put in the correct folder, the "live" pages can be checked on the Web using the URL provided by the ISP.

5 The Web page can typically be posted by means of a dial-up connection to the Internet (via a local ISP). Many ISPs provide the space for doing this free of charge as part of their monthly fee.



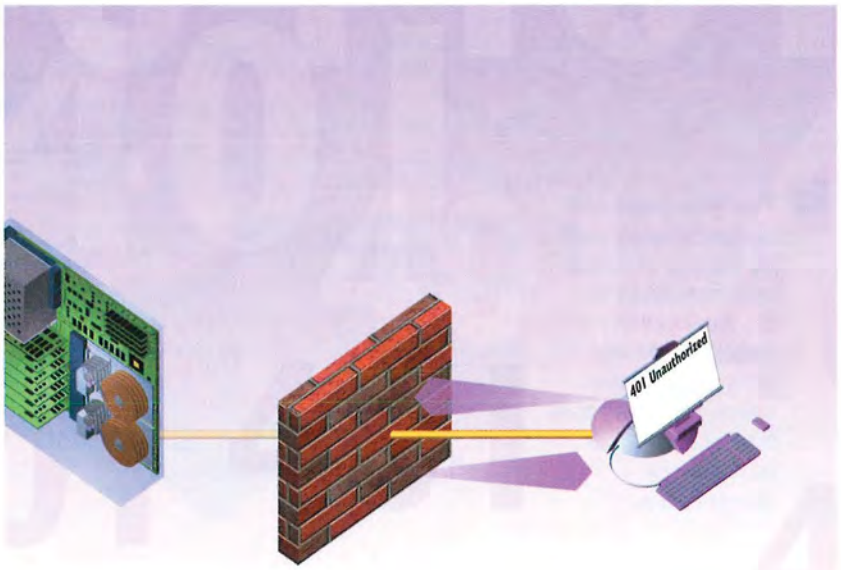


HTML Editor

4 If you want to add other media to your pages, such as digitized sound, music, or video, you need access to the digital files, or you can digitize them yourself with additional computer hardware. When you are dealing with media with large file sizes—such as a 4MB movie—it will take users a very long time to download the file. During editing, try to trim all the fat out of video and audio clips. Also try to scale back the length and resolution of audio and video to decrease file size.

CHAPTER

22 How Web Browsers Work



LIKE much of the Internet, the World Wide Web operates on a client/server model. You run a Web client on your computer—called a Web browser—such as Netscape Communicator or Microsoft's Internet Explorer. That client contacts a Web server and requests information or resources. The Web server locates and then sends the information to the Web browser, which displays the results.

When Web browsers contact servers, they're asking to be sent pages built with Hypertext Markup Language (HTML). They interpret those pages and display them on your computer. They also can display applications, programs, animations, and similar material created with programming languages such as Java and ActiveX, and scripting languages such as JavaScript.

Sometimes, home pages contain links to files the Web browser can't play or display, such as sound or animation files. In that case, you need a plug-in or a helper application. You configure your Web browser or operating system to use the helper application or plug-in whenever it encounters a sound or animation file the browser can't run or play.

Over the years, Web browsers have become increasingly sophisticated. Browsers have now become full-blown software suites that can do everything from videoconferencing to letting you create and publish HTML pages. They have also begun to blur the line between your local computer and the Internet—in essence, they can make your computer and the Internet function as a single computer system.

Toward this end of bringing the Internet more directly into computers, Microsoft has integrated Web browsing and the Internet directly into the operating system. For example, with Internet Explorer 4.0 and above, and with Windows 98 and above, the Windows desktop can be HTML-based. This means Web links can be directly embedded into the desktop. So, for example, you can have links to your favorite Web pages right on the desktop. And even applications such as word processors now have Web capabilities built into them, such as being capable of browsing the Web or building home pages. Even more significantly, using technology that Microsoft calls Active Desktop, Internet-based *desktop components* can live on the desktop. These components can be things such as stock tickers, which deliver live Web content directly to the desktop. You don't need to go fire up your Web browser to get the information; it's delivered straight to your Windows desktop without your having to do anything.

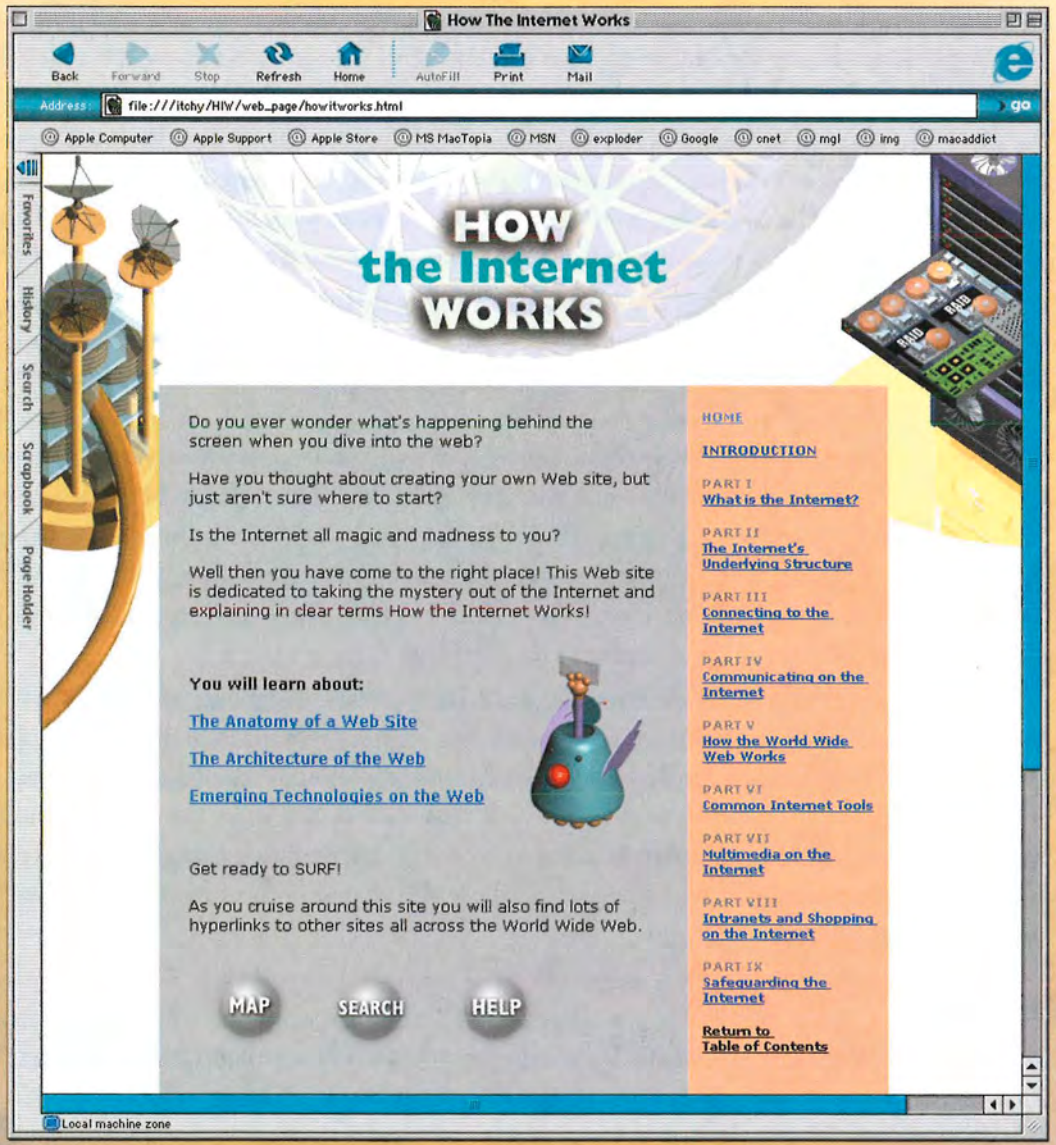
Both Microsoft and Netscape have also built entire suites of software around their browsers. Netscape, for example, calls its suite Netscape Communicator. Communicator includes modules for reading newsgroups; for reading, sending, and managing Internet mail; for audio conferencing; for collaborative work on *whiteboard applications* in which people can view and mark up the same documents simultaneously; and more. These enhancements will help usher in an era of collaborative computing. Not only will the Internet be used to transmit and receive information, but it will also alter the way we do business and help us communicate more effectively.

When browsing the Internet, one of the most frustrating experiences is the error messages browsers display when they're having trouble contacting a Web site. Depending on which browser you use, and which version of the browser you're using, those messages might differ. Version 5.0 and above of Internet Explorer, for example, sometimes displays error messages in plain English. But most browser don't. The final spread in this chapter lists the most common browser error messages—and what they mean.

How a Web Browser Works

1 Web browsers consist of client software that runs on your computer and displays home pages on the Web. There are clients for PC, Macintosh, and Unix computers.

2 A Web browser displays information on your computer by interpreting the Hypertext Markup Language (HTML) that is used to build home pages on the Web. Home pages usually display graphics, sound, and multimedia files, as well as links to other pages, files that can be downloaded, and other Internet resources.



```
<html>
<head>
<title>How The Internet Works</title>
<meta http-equiv="Content-Type" content="
</head>
<body bgcolor="#FFFFFF" link="#6666CC"
<div id="Layer1" style="position:absolute
<div id="Layer2" style="position:absolute
<div id="Layer3" style="position:absolute
<div id="Layer4" style="position:absolute
<div id="Layer5" style="position:absolute
<div id="Layer6" style="position:absolute
  <p><font size="3" face="Verdana, Arial,
    what's happening behind the screens
  <p><font size="3" face="Verdana, Arial,
    about creating your own Web site, bu
  <p><font size="3" face="Verdana, Arial,
    all magic and madness to you?</font
  <p><font size="3" face="Verdana, Arial,
    have come to the right place! This We
    out of the Internet and explaining in
  <p> </p>
  <p><font face="Verdana, Arial, Helvetica
    about:</b></font></p>
  <p><b><font face="Verdana, Arial, Helve
    Anatomy of a Web Site</a></font></
  <p><b><font face="Verdana, Arial, Helve
    Architecture of the Web</font></a></
  <p><b><font face="Verdana, Arial, Helve
    Technologies on the Web</font></a><
  <p> </p>
  <p><font face="Verdana, Arial, Helvetica
  <p><font face="Verdana, Arial, Helvetica
    site you will also find lots of hyper
    Wide Web.</font></p>
</div>
<div id="Layer7" style="position:absolute
  <p><font color="#999999"><b><font size=
  <p><font color="#9999CC"><a href="info
  <p><font color="#999999"><b><font size=
    I<br>
    <a href="pt01start.html">What is the
  <p><font color="#999999"><b><font size=
    II<br>
    <a href="part02start.html">The Interne
  <p><font color="#999999"><b><font size=
    III<br>
    <a href="part03start.html">Connecting
  <p><font color="#999999"><b><font size=
    IV<br>
    <a href="part04start.html">Communicat
  <p><font color="#999999"><b><font size=
    V<br>
    <font color="#6666CC"><a href="part0
  <p><font color="#999999"><b><font size=
    VI<br>
    <font color="#6666CC"><a href="part0
  <p><font color="#999999"><b><font size=
    VII<br>
    <font color="#6666CC"><a href="part0
  <p><font color="#999999"><b><font size=
    VIII<br>
    <font color="#6666CC"><a href="part0
    the Internet</a></font></font></b>
  <p><font color="#999999"><b><font size=
    IX<br>
    <font color="#6666CC"><a href="part0
  <p><font color="#999999"><b><font size=
    to <br>
    Table of Contents</font></a></font
  <p> </p>
  <p> </p>
  </div>
<div id="Layer8" style="position:absolute
<img src="bot.gif" width="151" height="
<div id="Layer9" style="position:absolute
<div id="Layer10" style="position:absolu
<div id="Layer11" style="position:absolu
</body>
</html>
```

Page Title

Background Color

Layer1

Hyperlink

3 The coding in the HTML files tells your browser how to display the text, graphics, links, and multimedia files on the home page. The HTML file your browser loads to display the home page doesn't actually have the graphics, sound, multimedia files, and other resources on it. Instead, it contains HTML references to those graphics and files. Your browser uses those references to find the files on the server and then display them on the home page.

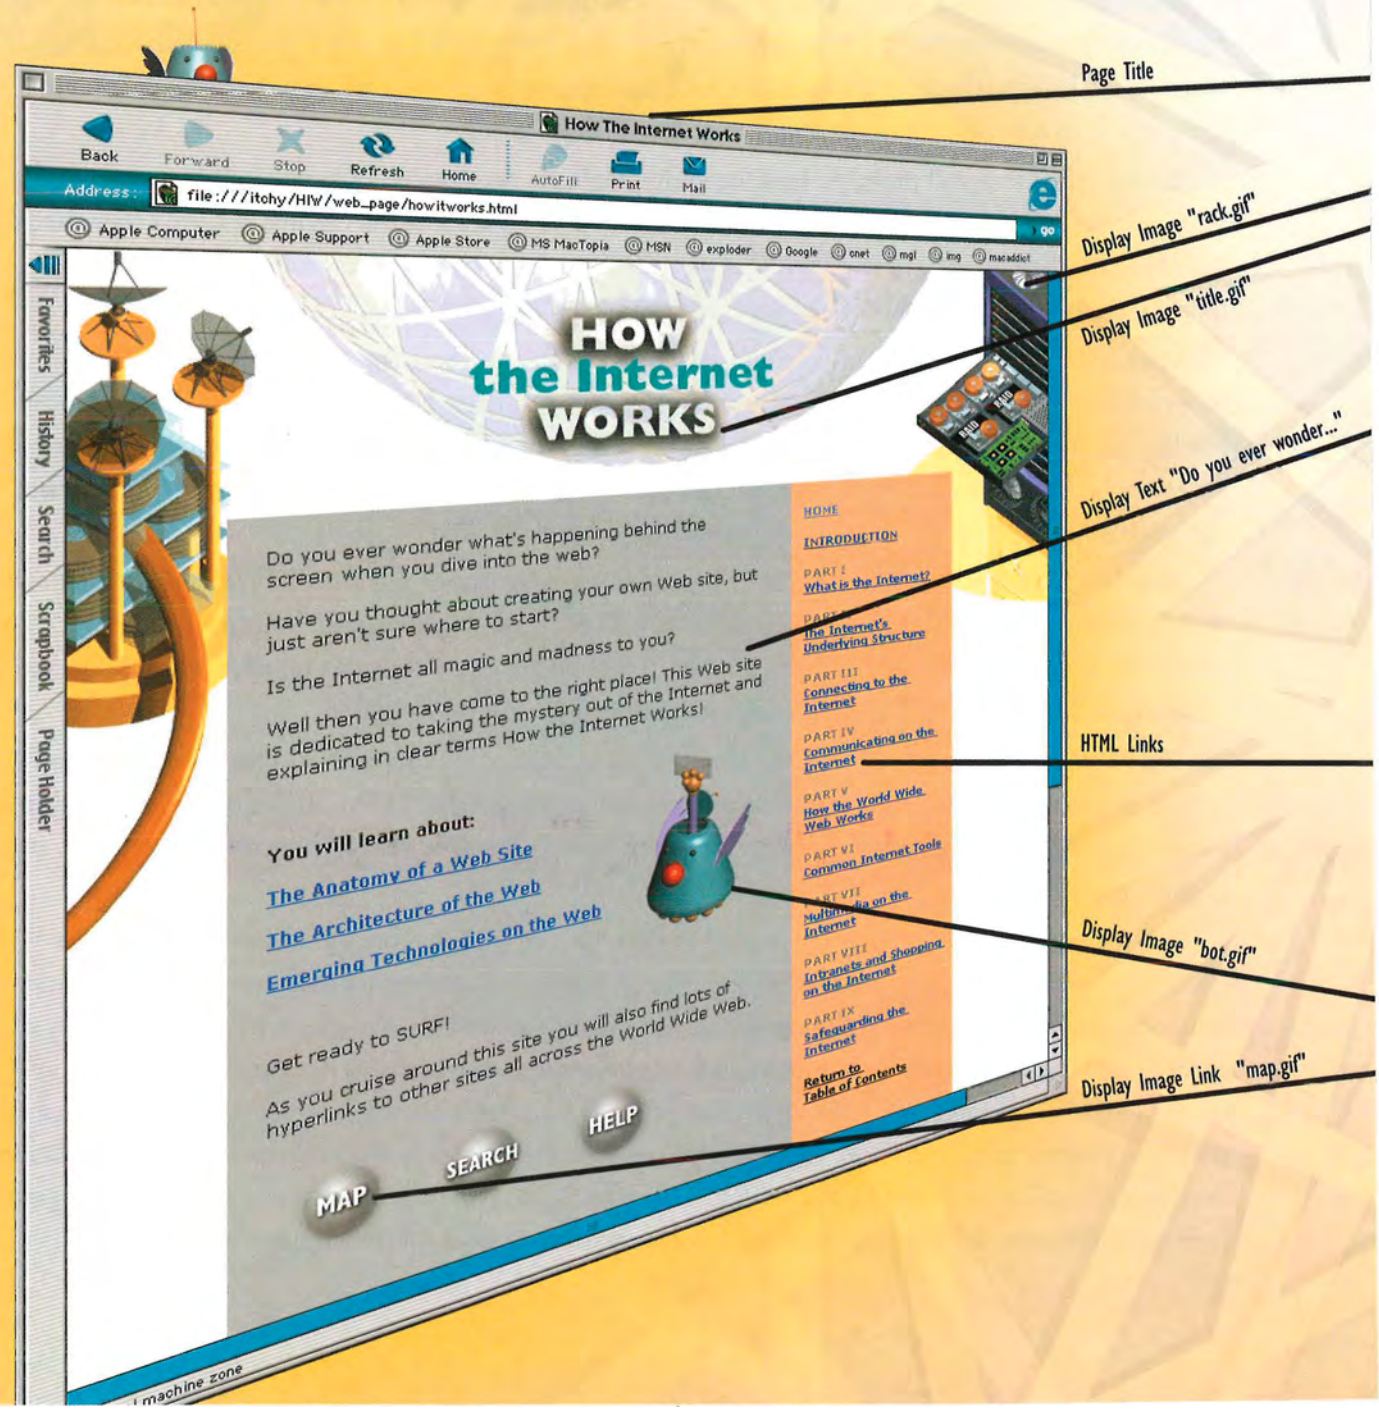
4 The Web browser also interprets HTML tags as links to other Web sites, or to other Web resources, such as graphics, multimedia files, newsgroups, or files to download. Depending on the link, it performs different actions. For example, if the HTML code specifies the link as another home page, the browser retrieves the URL specified in the HTML file when the user clicks the underlined link on the page. If the HTML code specifies a file to be downloaded, the browser downloads the file to your computer.

Many types of files exist on the Internet that Web browsers can't display. In particular, this pertains to many kinds of multimedia files, such as sound, video, and animation files. Still, there are often references to these types of files on Web pages. To view or play these files, you need what are called helper applications and plug-ins. You must configure your Web browser or operating system to launch these helper applications and plug-ins whenever you click an object that needs them in order to be viewed. Helper applications and plug-ins also can be used for displaying virtual reality pages, for chatting on the Internet, and for performing other Internet tasks.

The meanings of tags are easily decipherable. Every HTML tag, or instruction, is surrounded by a less-than and a greater-than sign—<P>. Often tags appear in pairs, the beginning tag and the ending tag. They are identical except for a simple slash in the end tag. So, a paragraph of text frequently is surrounded by tags such as this: <P> Paragraph of text.</P>. Also, tags are not case sensitive. <P> equals <p>.

How HTML Works

1 To display Web pages in any browser, you must add HTML tags to your original text. This process is called *tagging*.



2 Use HTML to give your text structure. All HTML files begin and end with the HTML tags. Headings are marked as such, as are paragraphs, line breaks, block quotes, and special character emphasis. Any carriage returns or indentations within the source text do not affect the browser's display of the page. HTML tags need to be put in if they are to be displayed in a browser.

3 The finished HTML document will be the source page for any browser on any computer. This simplicity of HTML makes cross-platform compatibility easy and reliable. The more complex and specialized the HTML tagging, the longer it will take to download and display the document.

Title

Graphic

Graphic

Text

Text

Hyperlinks

Graphic

Graphic
Hyperlink

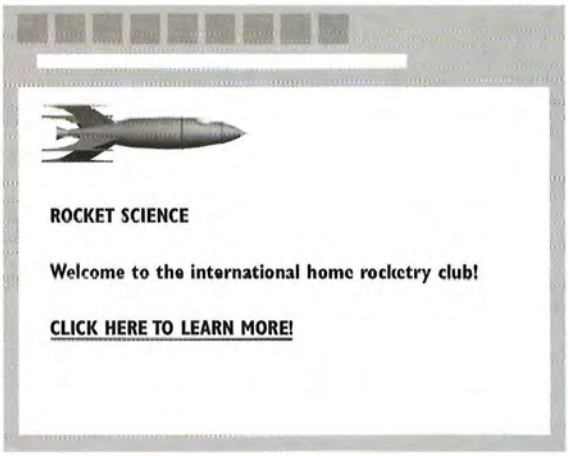
```
<html>
<head>
<title>How The Internet Works</title>
<meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1">
</head>

<body bgcolor="#FFFFFF" link="#6666CC" vlink="#666666">
<div id="Layer1" style="position:absolute; left:501px; top:-66px; width:180px; height:320px; z-index:1"><img src="rack.gif" width="217" height="372"></div>
<div id="Layer2" style="position:absolute; left:0; top:0; width:203px; height:352px; z-index:2"><img src="rtside.gif" width="125" height="450"></div>
<div id="Layer3" style="position:absolute; left:186px; top:0px; width:494px; height:162px; z-index:3"><img src="title.gif" width="545" height="173"></div>
<div id="Layer4" style="position:absolute; left:91px; top:180px; width:502px; height:875px; z-index:4; background-color: #CCCCCC; layer-background-color: #CCCCCC; border: 1px none #0
<div id="Layer5" style="position:absolute; left:506px; top:180px; width:156; height:877px; z-index:5; background-color: #FFCC99; layer-background-color: #FFCC99; border: 1px none #00
<div id="Layer6" style="position:absolute; left:115px; top:201px; width:369px; height:604px; z-index:6">
  <p><font size="3" face="Verdana, Arial, Helvetica, sans-serif">Do you ever wonder
    what's happening behind the screen when you dive into the web?</font></p>
  <p><font size="3" face="Verdana, Arial, Helvetica, sans-serif">Have you thought
    about creating your own Web site, but just aren't sure where to start?</font></p>
  <p><font size="3" face="Verdana, Arial, Helvetica, sans-serif">Is the Internet
    all magic and madness to you?</font></p>
  <p><font size="3" face="Verdana, Arial, Helvetica, sans-serif">Well then you
    have come to the right place! This Web site is dedicated to taking the mystery
    out of the Internet and explaining in clear terms How the Internet Works!</font></p>
  <p> </p>
  <p><font face="Verdana, Arial, Helvetica, sans-serif" size="3"><b>You will learn
    about:</b></font></p>
  <p><b><font face="Verdana, Arial, Helvetica, sans-serif" size="3"><a href="pt05_ch2102.html">The
    Anatomy of a Web Site</a></font></b></p>
  <p><b><font face="Verdana, Arial, Helvetica, sans-serif"><a href="pt01_ch0101.html"><font size="3">The
    Architecture of the Web</font></a></font></b></p>
  <p><b><font face="Verdana, Arial, Helvetica, sans-serif"><a href="pt07_ch3205.html"><font size="3">Emerging
    Technologies on the Web</font></a></font></b></p>
  <p> </p>
  <p><font face="Verdana, Arial, Helvetica, sans-serif">Get ready to SURF!</font></p>
  <p><font face="Verdana, Arial, Helvetica, sans-serif">As you cruise around this
    site you will also find lots of hyperlinks to other sites all across the World
    Wide Web.</font></p>
</div>
<div id="Layer7" style="position:absolute; left:517px; top:198px; width:139px; height:631px; z-index:7">
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif"><a href="index.html"><font color="#9999CC">HOME</font></a></font></b></font></p>
  <p><font color="#9999CC"><a href="intro.html"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">INTRODUCTION</font></b></a></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    I<br>
    <a href="pt01start.html">What is the Internet?</a></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    II<br>
    <a href="part02start.html">The Internet's Underlying Structure</a></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    III<br>
    <a href="part03start.html">Connecting to the Internet</a></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    IV<br>
    <a href="part04start.html">Communicating on the Internet</a></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    V<br>
    <font color="#6666CC"><a href="part05start.html">How the World Wide Web Works</a></font></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    VI<br>
    <font color="#6666CC"><a href="part06start.html">Common Internet Tools</a></font></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    VII<br>
    <font color="#6666CC"><a href="part07start.html">Multimedia on the Internet</a></font></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    VIII<br>
    <font color="#6666CC"><a href="part08start.html">Intranets and Shopping on
    the Internet</a></font></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif">PART
    IX<br>
    <font color="#6666CC"><a href="part09start.html">Safeguarding the Internet</a></font></font></b></font></p>
  <p><font color="#999999"><b><font size="2" face="Verdana, Arial, Helvetica, sans-serif"><a href="index.html"><font color="#000000">Return
    to <br>
    Table of Contents</font></a></font></b></font></p>
  <p> </p>
  <p> </p>
  </div>
<div id="Layer8" style="position:absolute; left:351px; top:388px; width:129px; height:132px; z-index:8"><img src="bot.gif" width="151" height="154"></div>
<div id="LayerY" style="position:absolute; left:120px; top:622px; width:77px; height:66px; z-index:9"><a href="map.html"><img src="map.gif" width="83" height="83" border="0"></a></di
<div id="Layer10" style="position:absolute; left:215px; top:623px; width:65px; height:76px; z-index:10"><a href="searchHW.html"><img src="search.gif" width="83" height="83" border=
<div id="Layer11" style="position:absolute; left:312px; top:622px; width:67px; height:101px; z-index:11"><a href="helpme.html"><img src="help.gif" width="83" height="83" border="0">
</body>
</html>
```

4 Most Web browsers enable your document to retain its structural integrity when you display, or *parse*, it. Headings will appear in a larger font size than text within paragraphs, for example, and block quotes will be uniformly indented. However, the look might vary from browser to browser. Note that browsers determine the exact font, size, and color. Also be aware that the relative importance of the elements is always kept intact.

How Dynamic HTML Works

1 Dynamic HTML (DHTML) differs from traditional HTML in that it enables Web pages to be changed on-the-fly, after they've been downloaded. In plain HTML, after a page is downloaded, it is static and can be changed only when a user takes an action of some kind. But DHTML, for example, could cause an animation of a rocket to fly across your browser window several seconds after the page has been downloaded—without your doing anything.



HTML



DHTML



Web Server

2 DHTML does its work without having to contact the server after the page downloads, so it can perform some interactive functions more quickly than other technologies that have to contact the server. The instructions for performing the commands are in the HTML commands that are in the page itself.

3 Although DHTML is often referred to as if it were a single technology, it is, in fact, a general term used for a group of technologies that can work together or by themselves to change a Web page after the page has been downloaded to your computer. These technologies are the Document Object Model (DOM), Cascading Style Sheets (CSS), and client-side scripting languages, such as JavaScript.



4 The DOM defines every object and element on a Web page and enables those objects to be manipulated or accessed. This includes fonts, graphics, tables, and visual elements, as well as elements you can't necessarily see, such as the browser's version number and the current date and time. Without DOM, all the elements on a page are static. So on the simplest level, DHTML could use the DOM to change the font of every letter, individually, on a Web page.



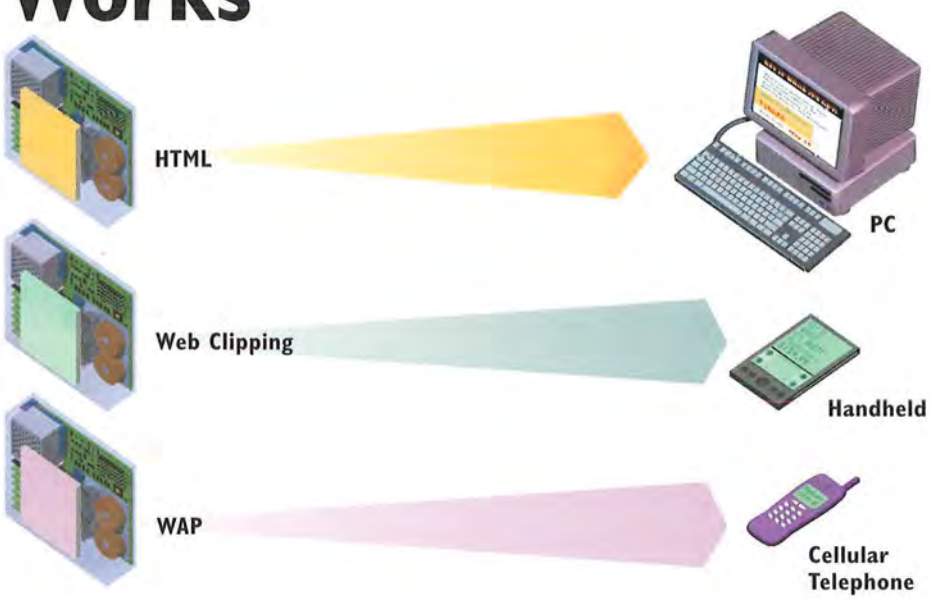
5 Cascading Style Sheets are, in essence, templates that apply formatting and style information to the elements of a Web page. They're called *cascading* because any single page can have more than one style sheet associated with it. Additionally, Cascading Style Sheets enable images to overlap one another. This enables animations to be created easily on a page.



6 Client-side scripting languages perform much of the work of DHTML. These languages access the DOM and manipulate its elements, and they do the same to Cascading Style Sheets. They perform the actions of DHTML. So a script, for example, could turn a word a different color when a mouse moves across it, or it could create easy-to-use collapsible navigation on every page on a Web site.

How XML Works

1 XML solves several a major problems for Web developers. Without it, to deliver Web pages to different devices such as computers, cell phones, and wireless Palm devices, a developer must create and maintain three separate Web sites, with special coding for each device—HTML for the computer, WAP for the cell phone, and what's called Web clipping for the wireless Palm devices. It's an expensive, difficult, and time-consuming proposition. And even if a developer is building a site only for computers, every time the design changes, all the pages must be recoded—again, an expensive and time-consuming proposition.



2 With XML, a developer can create the Web site just once. Then it can be automatically formatted to several types of devices, such as Internet-connected computers, wireless Palm devices, and cell phones using WAP. And even if the site is being built for only computers, when there is a redesign, with XML, all the pages need not be rebuilt.

```
<Sale Flyer>
<Offer>Get It While It's Hot!</Offer>
<Promotional Copy>
You can't miss this one!
One-time offer only –
gaming systems at prices
you won't believe! </Promotional Copy>
<Product>Sony PlayStations</Product>
<Price> $159.95</Price>
Sale ends <End Date> May 15
</End Date>
</Sale Flyer>
```

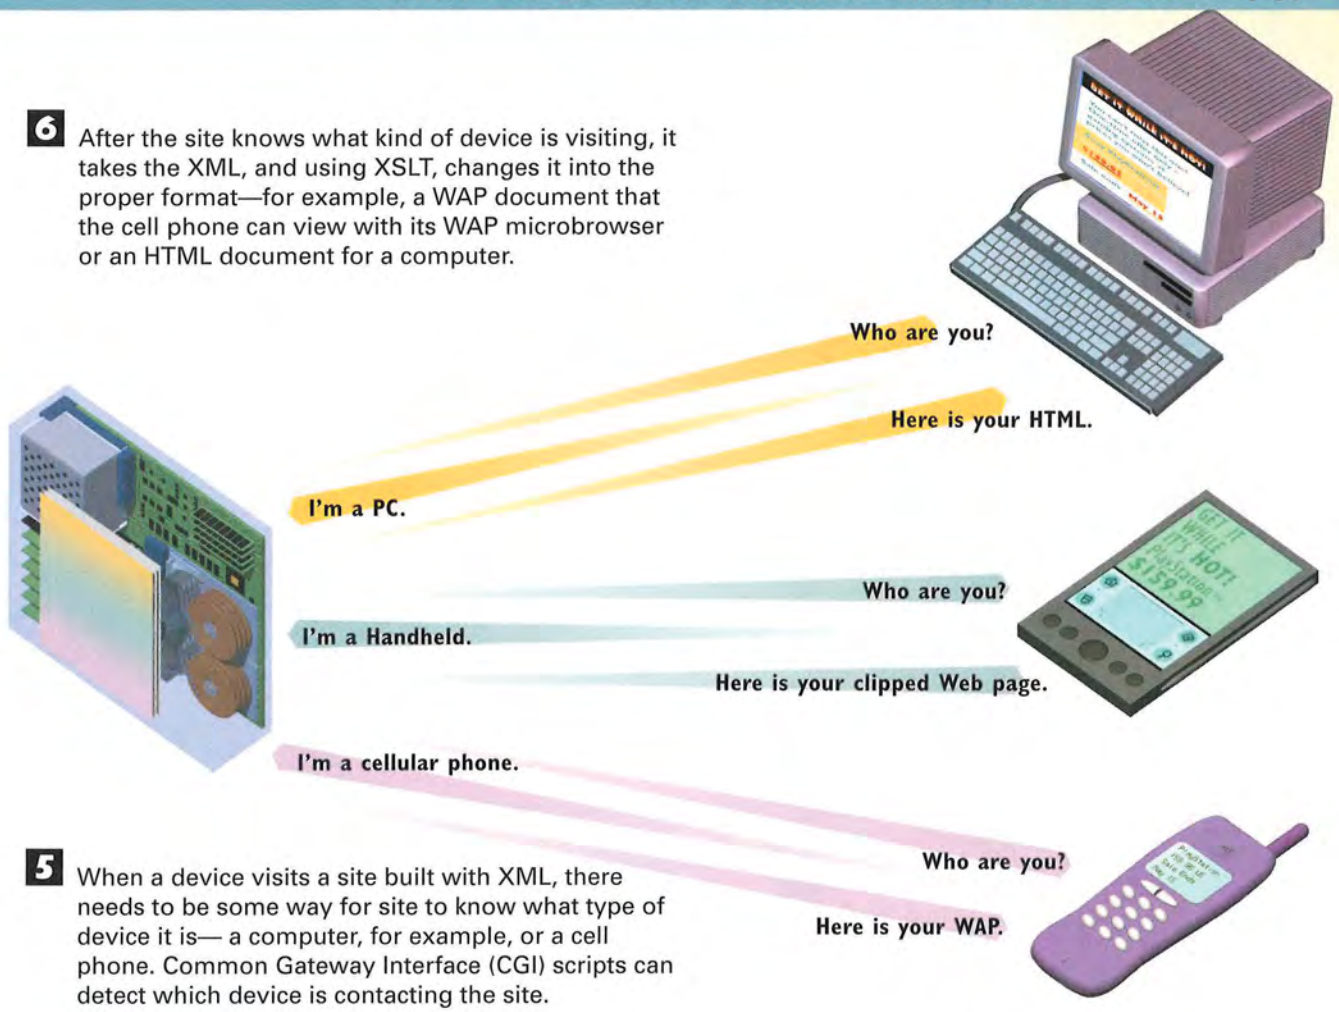
XML



XSLT

3 The most important concept to understand about XML is that the language is used only to convey information about content, not about the presentation of the content. So, for example, it doesn't give instructions on what size text should be. But it uses tags to define the type of content on the page. Then it uses other techniques, as you'll see in the next steps, to display those pages. In that way, a single page can be displayed many different ways, without having to go back and alter the original page—only the designs, which are separate from the content, need to be changed.

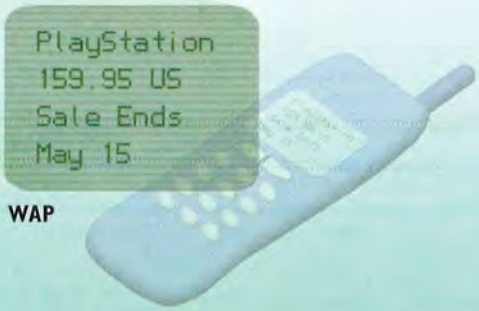
6 After the site knows what kind of device is visiting, it takes the XML, and using XSLT, changes it into the proper format—for example, a WAP document that the cell phone can view with its WAP microbrowser or an HTML document for a computer.



5 When a device visits a site built with XML, there needs to be some way for site to know what type of device it is— a computer, for example, or a cell phone. Common Gateway Interface (CGI) scripts can detect which device is contacting the site.



HTML

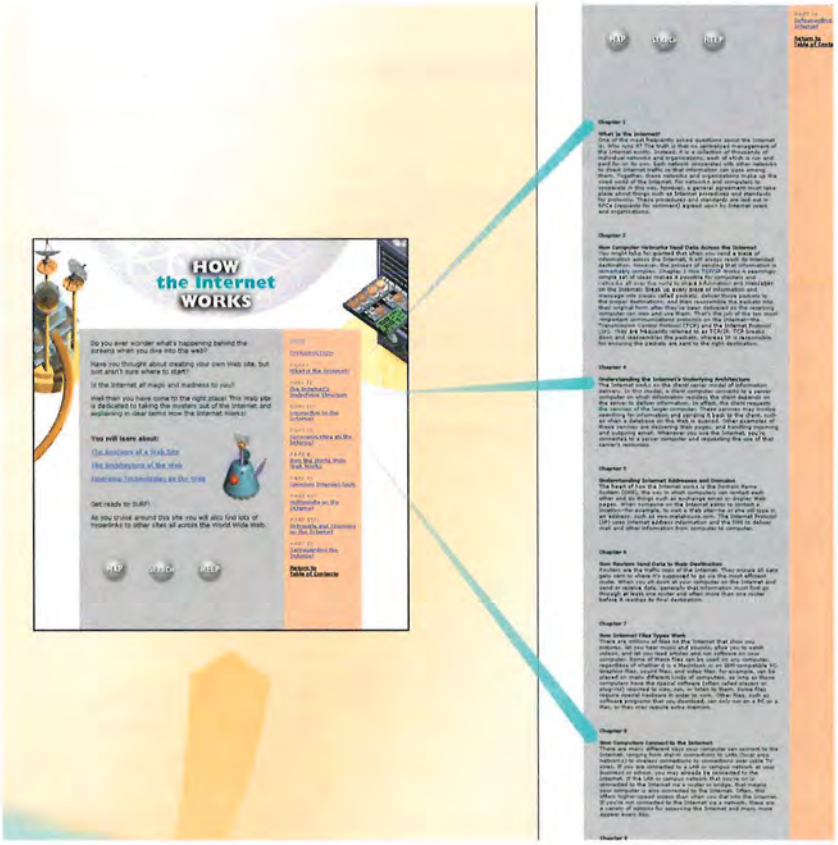


WAP

4 When XML content is posted on a Web site, different designs need to be applied to that content so that it can be viewable by devices connecting to it—for example, cell phones. eXtensible Style Language Transformations (XSLT) can be applied to the XML. XSLT can take XML and apply different designs to it or change it to other forms of XML—for example, it can take the XML and turn it into a WAP page that can be viewed by a cell phone and take the same XML and turn it into a HTML document with a different design.

CHAPTER

24 How Hypertext Works



IN the late 1960s, a computer scientist named Ted Nelson introduced hypertext, a concept that lays the foundation for the World Wide Web and its connections between documents, or pages.

Nelson wanted to create a new way of exploring information. He wanted to provide the reader with a spontaneous means of accessing more and more in-depth information about something that sparked the reader's interest when reading text on the page. Rather than reading a document from beginning to end, digesting the material in a sequential order, the reader could highlight a word and receive more information on the meaning of that word, for example.

Nelson imagined that you could read the United States Constitution, come across the term "electoral college," and then open another document that explains how the electoral college works. From that document, you might open another document that lists the votes of the electoral college from its founding to the present. From there, you might choose to open a document about William Jefferson Clinton, then another about the First Cat, Socks, then another about the care and feeding of felines, and so forth. Ultimately, you could find a subject that wasn't even remotely connected to the Constitution, but which you would find interesting or entertaining.

This hypertext concept was obviously on Tim Berners-Lee's mind when he began thinking about how researchers could share their work across the Internet. He envisioned a system in which a document could be linked to other documents, enabling researchers to easily find more and related information simply by following a link from one document on the network to another.

Typically, hypertext consists of a hyperlink that appears onscreen as a highlighted word, icon, or graphic. By moving a mouse cursor over the item, or object, and clicking it, you easily navigate to additional information. On the Web, that information can be located at any other place on the Web, be it on the same host server or one across the globe. A linked object can be various media, such as text (linking from one character to a whole document, for example), a graphical button (such as direction arrows that move from page to page), or still images (photos, icons, or a comic strip), for example. The documents and objects that are being linked to can be on the same site as the original document or on an entirely different site.

Hypertext links are embedded into a Web document using Hypertext Markup Language (HTML). A text link usually appears onscreen as an underlined word or phrase and is sometimes rendered in a different color from other text, depending on how your Web browser interprets the HTML codes. When you place the mouse cursor on this underlined text and click the mouse button, you initiate a request by the browser for a new Web page or—if the text references an internal link to information in the same document—direct your browser to scroll to another, specific point within the same document.

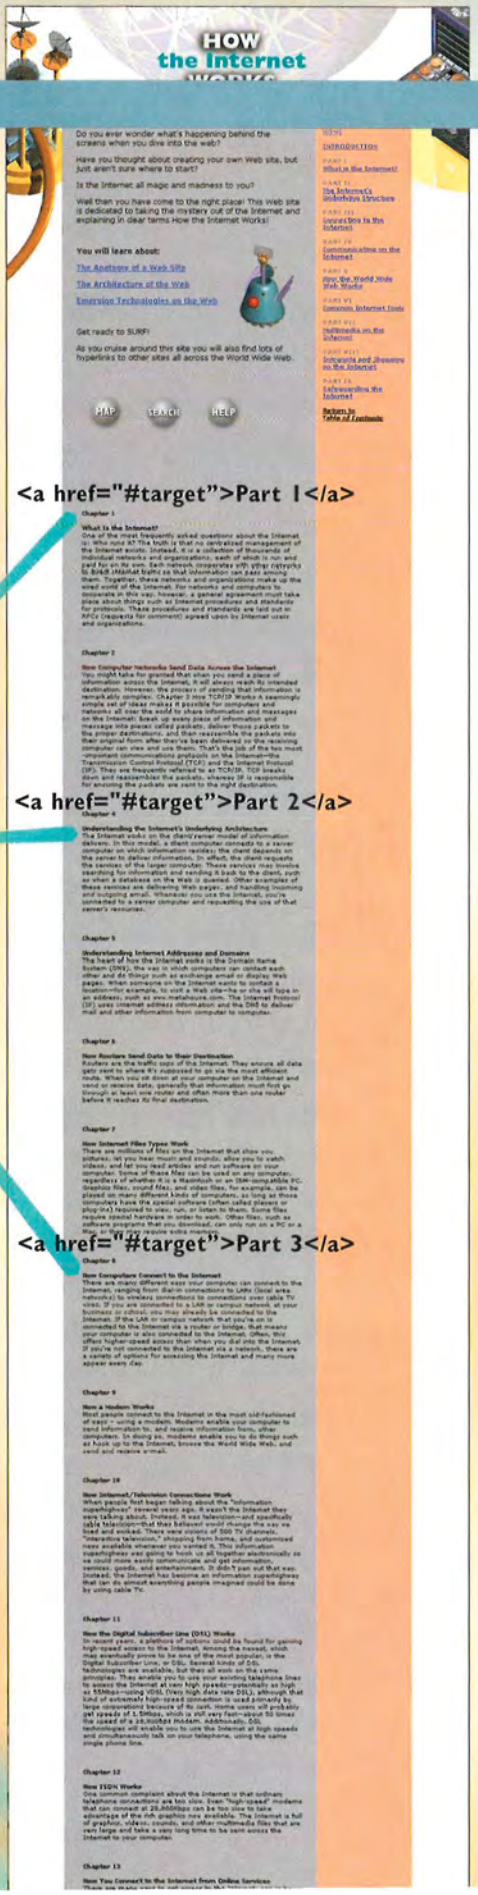
Images or icons can also act as hyperlinks. When you move the mouse cursor over the icon or graphic and click the mouse button, you launch the request to retrieve the linked information.

How Hyperlinks Work

1. The "hyperlinking" begins when you first retrieve a Web page from a remote Web server. Target links within the page move you quickly from one part of the page to the next.







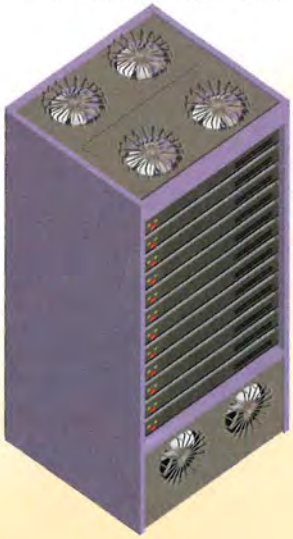
Part 1

Part 2

Part 3

2 A *relative* link is used to initiate a request by your Web browser to retrieve a page located on the same Web server as the page from which you are linking. Web developers use relative links when they're pointing to a page under their control because this enables the to more easily maintain their HTML coding if their server locations change.

Chapter 4







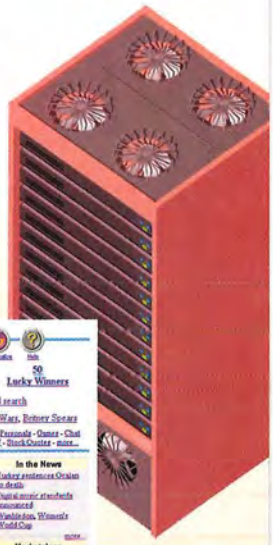


Chapter 7

3 A hyperlink that leads to a completely different Web server uses an *absolute* link. Because these links might not be under the control of the Web developer who created the page, there is a chance that this will be a "dead link"—the page will have been moved without the Web developer knowing about it.









<img src="searchyahoo.gif">

CHAPTER 25

How URLs Work



THE Web pages and the hosts that make up the World Wide Web must have unique locations so that your computer can locate and retrieve the pages. The unique identifier for a host is called the Internet Protocol (IP) address, and the unique identifier for a page is called the uniform resource locator (URL). A URL functions much like a postal or e-mail address. Just as postal and e-mail addresses list a name and specific location, a URL, or Web address, indicates where the host computer is located, the location of the Web site on the host, and the name of the Web page and the file type of each document, among other information.

A typical URL looks like this:

```
http://www.zdpress.com/internetworks/index.html/
```

If you were to interpret the instructions in this URL from left to right, it would translate to: "Go to the host computer called `zdpress` (a commercial business), in a directory called `internetworks`, and retrieve a hypertext document with the filename `index.html`." The URL, or address, tells the browser which document to fetch and exactly where to find it on a specific remote host computer somewhere on the Internet.

The first part of the URL indicates what type of transfer protocol will be used to retrieve the specified document. The most common request is for a hypertext document that uses Hypertext Transfer Protocol (HTTP).

The second portion of the URL refers to the specific host computer on which the document resides, which is to be contacted by the browser software. This part of the address is also called the *domain name*. See Chapter 5, "How Internet Addresses and Domains Work," for more information about domains.

The third part of the URL is the directory on the host computer that contains a specific Web site or multiple Web sites. This is always located after the first single slash in the URL and is essentially the subdirectory on the hard disk that houses the Web site. Subdirectories might also be indicated in this part of the address. For example, if the previous URL were changed to `http://www.zdpress.com/internetworks/partone/chapters/chapte.html` there would be two subdirectories—`part one` and `chapters`.

In the preceding example, the filename is `chapte.html`. This is always the last portion of the URL. If you see an address without a filename, it is assumed that the filename `index.html` contains the requested Web page. Therefore, the default document a Web server will deliver to the client when no other filename is listed is `index.html`. (Note that sometimes the last portion of the URL might not be a filename—it could be other types of information required by a Web server, such as codes required to log on to the Web server.)

The illustration in this chapter shows the process necessary to request and retrieve a Web document. When a request for a document occurs for the first time in a Web-browsing session, the host computer must first be located to find the file. After that, the specific subdirectory and document are retrieved.

How URLs Are Structured

1 The first part of the URL indicates which type of transfer protocol will be used to retrieve the specified document. The most common request is for a hypertext document that uses the HTTP protocol.

3 The third part of the URL is the directory on the host computer that contains a specific Web site. A host computer can house multiple Web sites. This third segment of the address is essentially the root directory that houses the HTML document, CGI, image, and so on that is being requested. Subdirectories might also be indicated in this part of the address.

http://www.sample.com/samples/sample.html

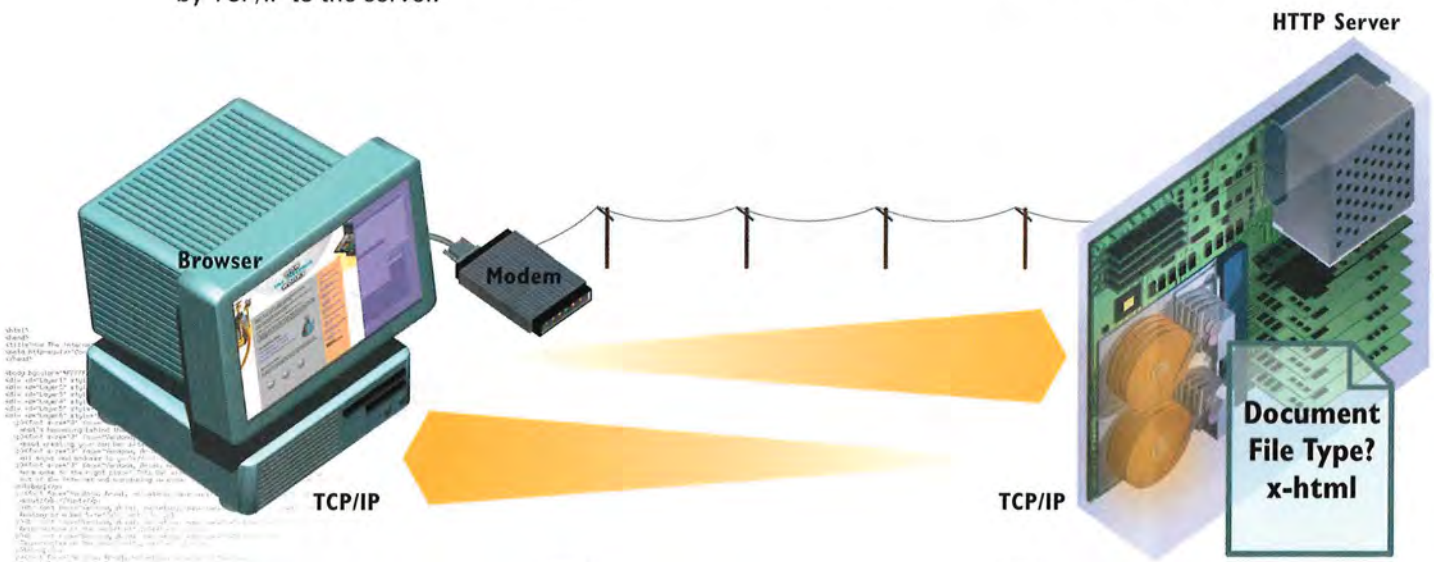
2 The second portion of the URL is the specific host computer on which the document resides, which is to be contacted by the browser software. This part of the address is also called the *domain*. Domain names end in a suffix that indicates which type of organization the domain is. For example, `.com` indicates a commercial business, `.edu` indicates a college or university, `.gov` indicates a government office, `.mil` indicates a military facility, and `.org` indicates a not-for-profit organization. The suffix also can indicate the country in which the host computer is located. For example, `.ca` is in Canada, and `.au` is in Australia.

4 The last segment of the URL is the filename of the specific Web page you are requesting. If no filename is indicated, the browser assumes a default page, usually called `index.html`.

How URLs Help Retrieve Web Documents

1 The Web browser installed on your local computer sends your TCP/IP software a signal that it is ready to request a document. TCP/IP makes a connection with the host TCP/IP software. After the connection is established, your browser makes a request for a document by sending its URL through the two-way connection maintained by TCP/IP to the server.

2 The HTTP server is the portion of the host computer that runs HTTP server software. TCP/IP makes and maintains the connection this way. The browser can use HTTP to send requests and receive pages through the host's Web server software. This software enables the host to communicate with the client browser, in HTTP, over TCP/IP.



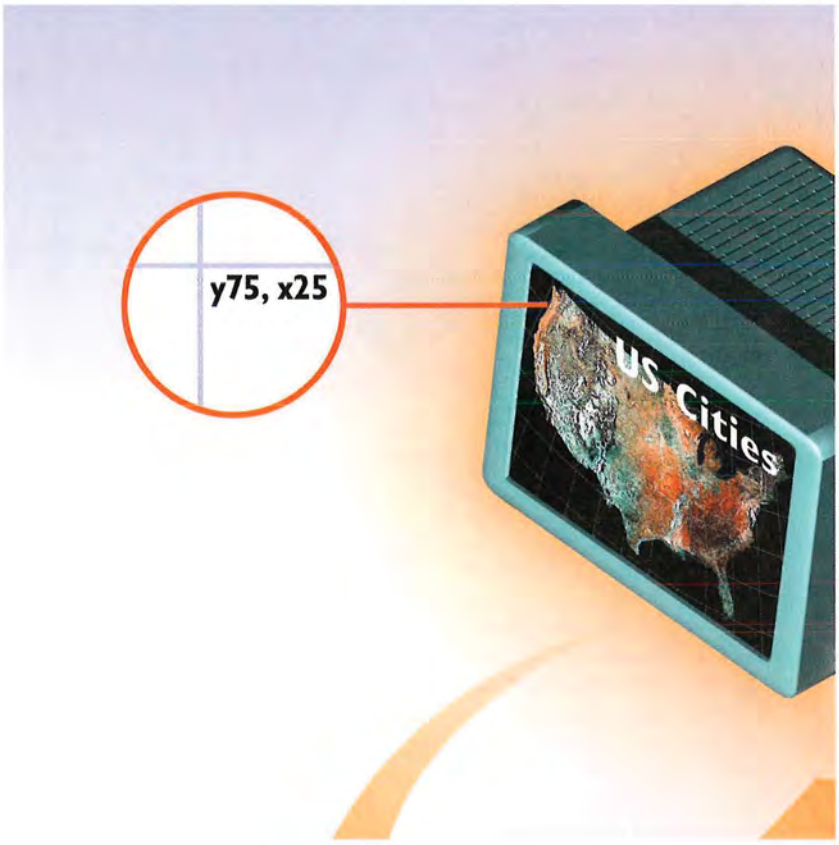
5 The browser on your local computer reads the file type. If it is an HTML document, the browser examines the content, breaking it down into meaningful parts. Two general parts include text, which is displayed by the browser word for word. The other part consists of HTML markup information called *tags*, which are not displayed but display formatting information, such as normal text, bold headers, or colored hypertext. The results are displayed on your monitor.

4 If the document is found, the host checks its file type (usually either x-html or x-text) and sends this information to the client with the requested page. When the client receives the page, it first checks the file type. If the type is one it can display, it does so; otherwise, it prompts the user to see whether she wants to save it to disk or open it using a helper application. The x-html file type is by far the most common one used when transmitting Web pages.

3 The server then receives the transmitted URL and responds in one of three ways. It follows the directory path given in the URL; the server finds the file on its local hard disk and opens it; the server runs a CGI script or detects an error (such as `file not found`) and generates an error document to be sent back to the client.

CHAPTER

26 How Image Maps and Interactive Forms Work



GRAPHICS called image maps and functions called interactive forms demonstrate two of the more common and helpful uses of HTML. *Image maps* are static images that have been turned into clickable images with various clickable parts. *Interactive forms* are HTML-based pages you fill out by providing information such as your name, e-mail address, and similar information. Both image maps and interactive forms are created using Common Gateway Interface (CGI), a communications protocol by which a Web server can communicate with other applications. (For more information about how CGI works, turn to Chapter 35, "How CGI Scripting Works.")

Image maps can be thought of as fancy hyperlinks. However, instead of a word, an entire icon, or an image linking to another page, an image is divided into different segments, or coordinates, that link to different HTML pages. That is, image maps link to another document through a predefined *hot* area within an image. As soon as you click your mouse on a hot spot, a CGI script and special image map coordinates file with the suffix .map go to work. A CGI application reads the map file to match coordinates of a mouse click with a corresponding URL. For instance, imagine an electronic map of the United States in which you click Washington, D.C. In the HTML code for that page, the electronic map is surrounded by a tag and an attribute called `ISMAP`. The code looks something like this:

```
<A HREF="some.server/maps/clickable.map>
<IMG SRC+"clickable.map" ISMAP>
</A>
```

The x,y coordinate of your mouse click is sent to the server. The coordinates are received by the server and then redirected to a CGI application. The CGI application scans the file for matching coordinates and then forwards the corresponding URL to the server. Lastly, if the Web page resides on the same server, it delivers that Web page to the client browser. If not, the server returns the URL to the client browser, which in turn sends a request to the correct server for the page. You then see the page about Washington, D.C. begin to load on your browser. Behind the scenes, the server passed your mouse click coordinates to a CGI application via the CGI. Then, the CGI application matched those coordinates to its URL in a .map file. Finally, the URL sent the URL back to the server, which redirected the client browser to the new Web page.

Forms work differently, although they also use CGI. In a form, when you fill in information on a Web page, that information goes to the server for processing. Next, the server redirects the information to a CGI application that is called by the form *submit*. (CGI scripts are activated by the server in response to an HTTP request from the client.) Lastly, a CGI application might send form data to another computer program, such as a database; save it to a file; or even generate a unique HTML document in response to the user's request.

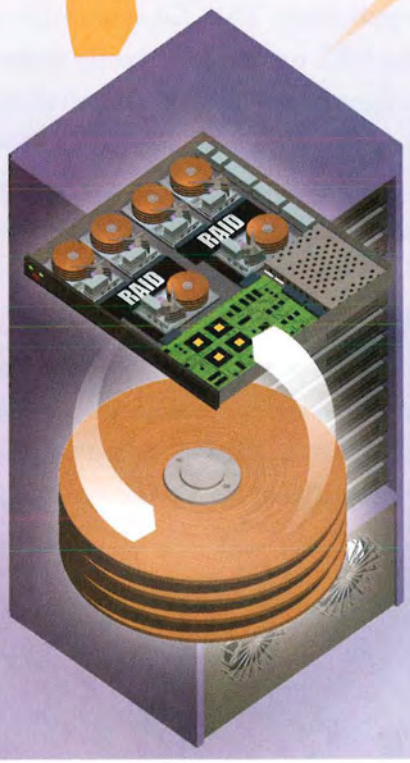
How Image Maps Work

1 In this map example, the user clicks Seattle. The x,y grid coordinate is 75, 25. In the HTML code, the browser recognizes the ISMAP image tag attribute. The mouse click activates the browser to send the x,y coordinate of the click to the server. The location of the `National.map` file is also sent to the server.



2 The server hands off the coordinate and map file data to a CGI application. The CGI application matches the coordinates to the URL that has been requested by the user by clicking that portion of the map. This URL is handed back to the server, and the server sends the page to the client.

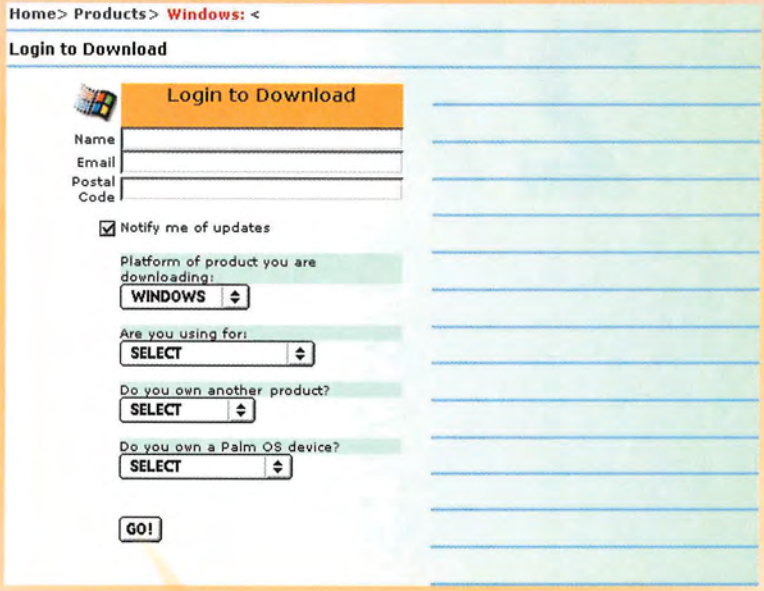
3 The Web document is either served up (if it resides on the same server) or the client browser if forwarded to the new URL.



4 The client browser either displays the returned page or (based on the returned URL) sends a request to the correct server for the page.

How Interactive Forms Work

1 In working with an interactive form, the user clicks a data-entry submit button. The data in the data fields is sent to the server with the request.



Submit Form Data



2 When a form submission is received by the server, it activates a CGI application, a program, or a script that interacts with a Web server; it then passes the resulting information to a Web form. (The application could add the form data to a database or compare it to a password list of eligible users, among other tasks.) The program's output goes either to another program, such as a database, or into a unique HTML document, or both.

CHAPTER

27 How Web Host Servers Work



TO serve up pages, Web sites need a *host*—a computer—and server software that runs on the host. The host manages the communications protocols and houses the pages and the related software required to create a Web site on the Internet. The host machine often uses the Unix, Windows NT, Linux, or Macintosh operating systems, which have the TCP/IP protocols built in.

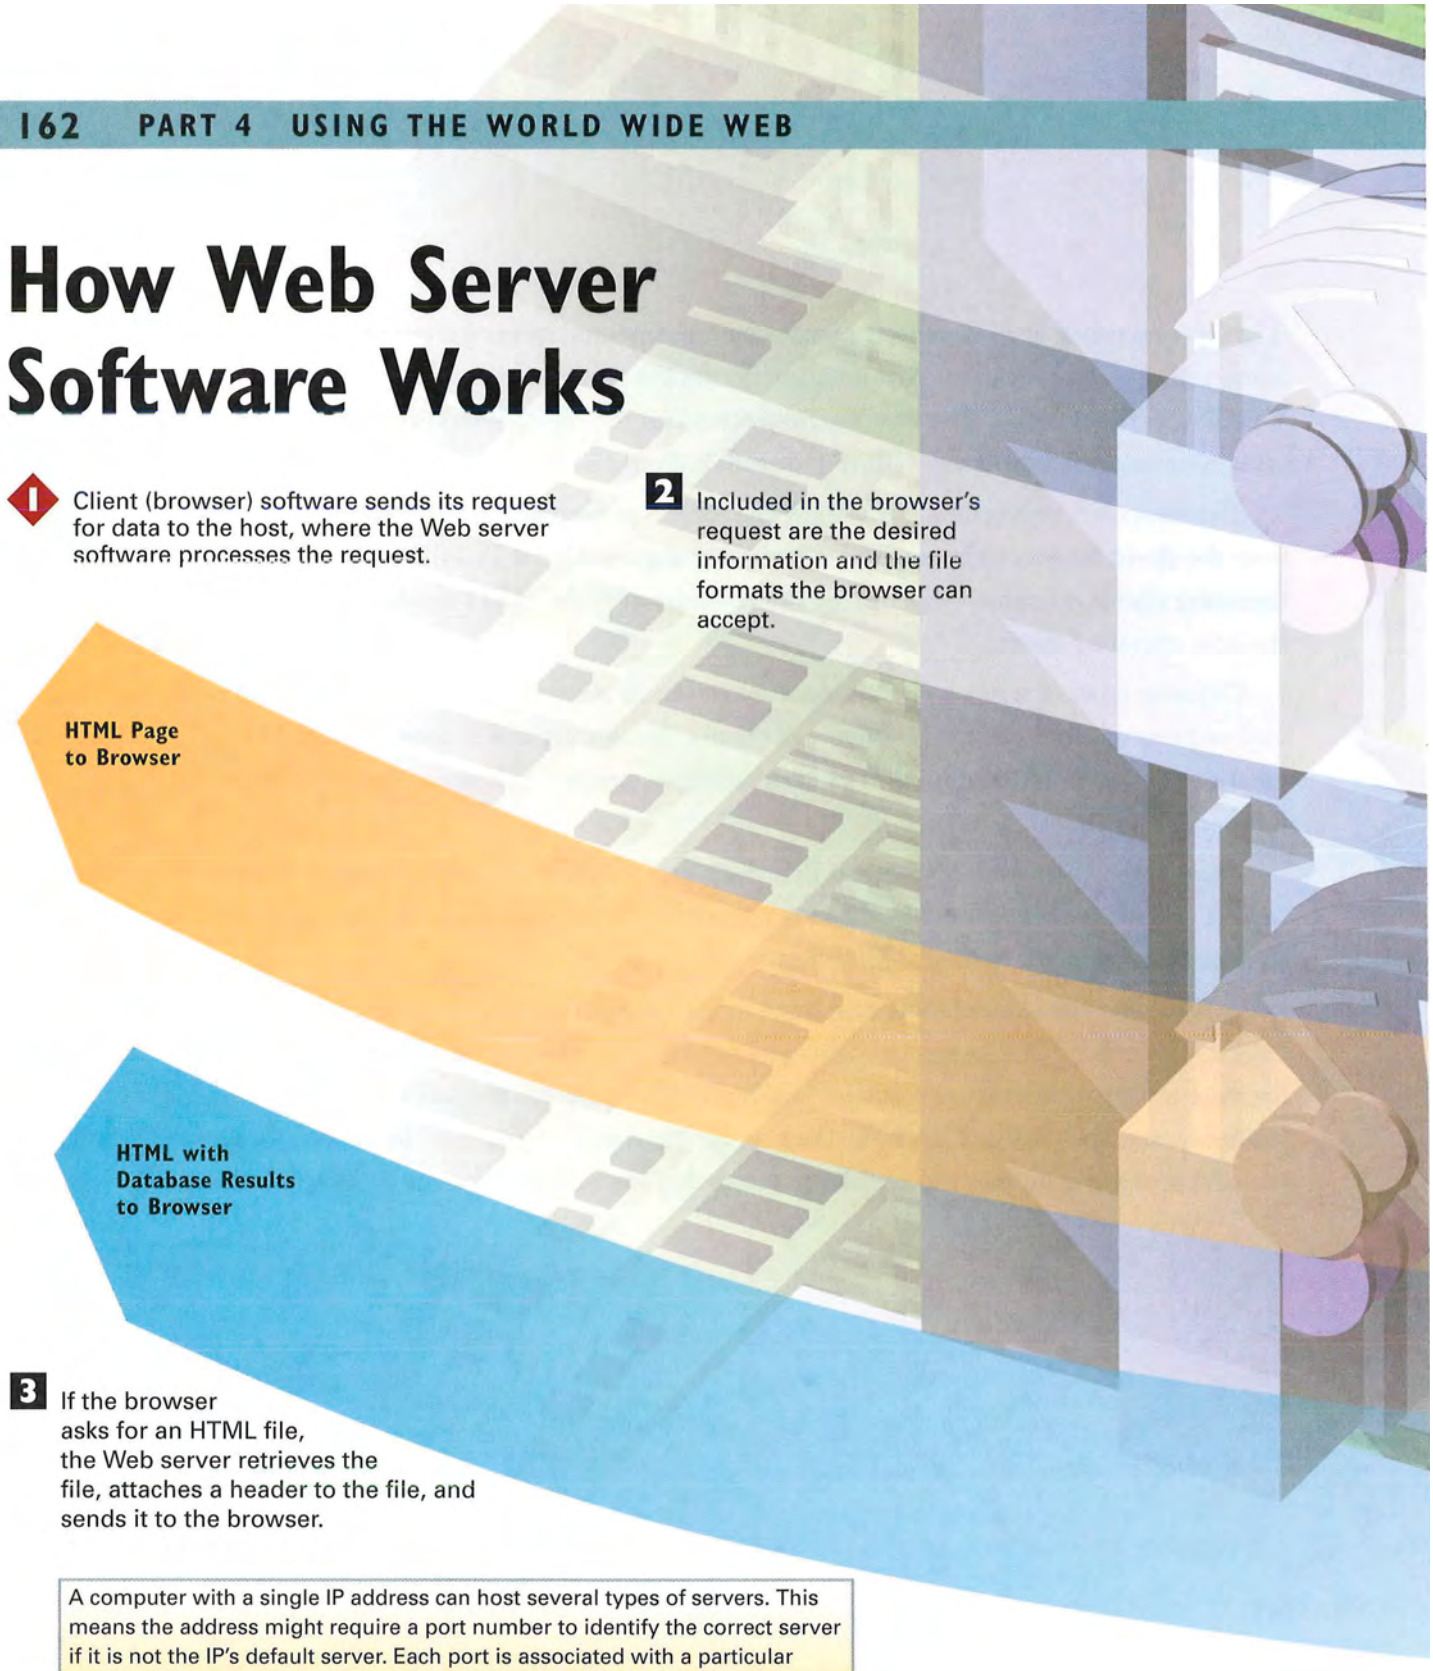
The server software resides on the host and serves up the pages and otherwise acts on the requests sent from the client browser software. The server is not responsible for TCP/IP communications—the host operating system does that—but instead the server handles the HTTP requests and communications with the host operating system.

Different types of server software (database servers or network servers, for example) exist that perform various types of services for various types of clients. Specifically, a *Web server* is an HTTP server, and its function is to send information to the client software (typically a browser) using the Hypertext Transfer Protocol (HTTP).

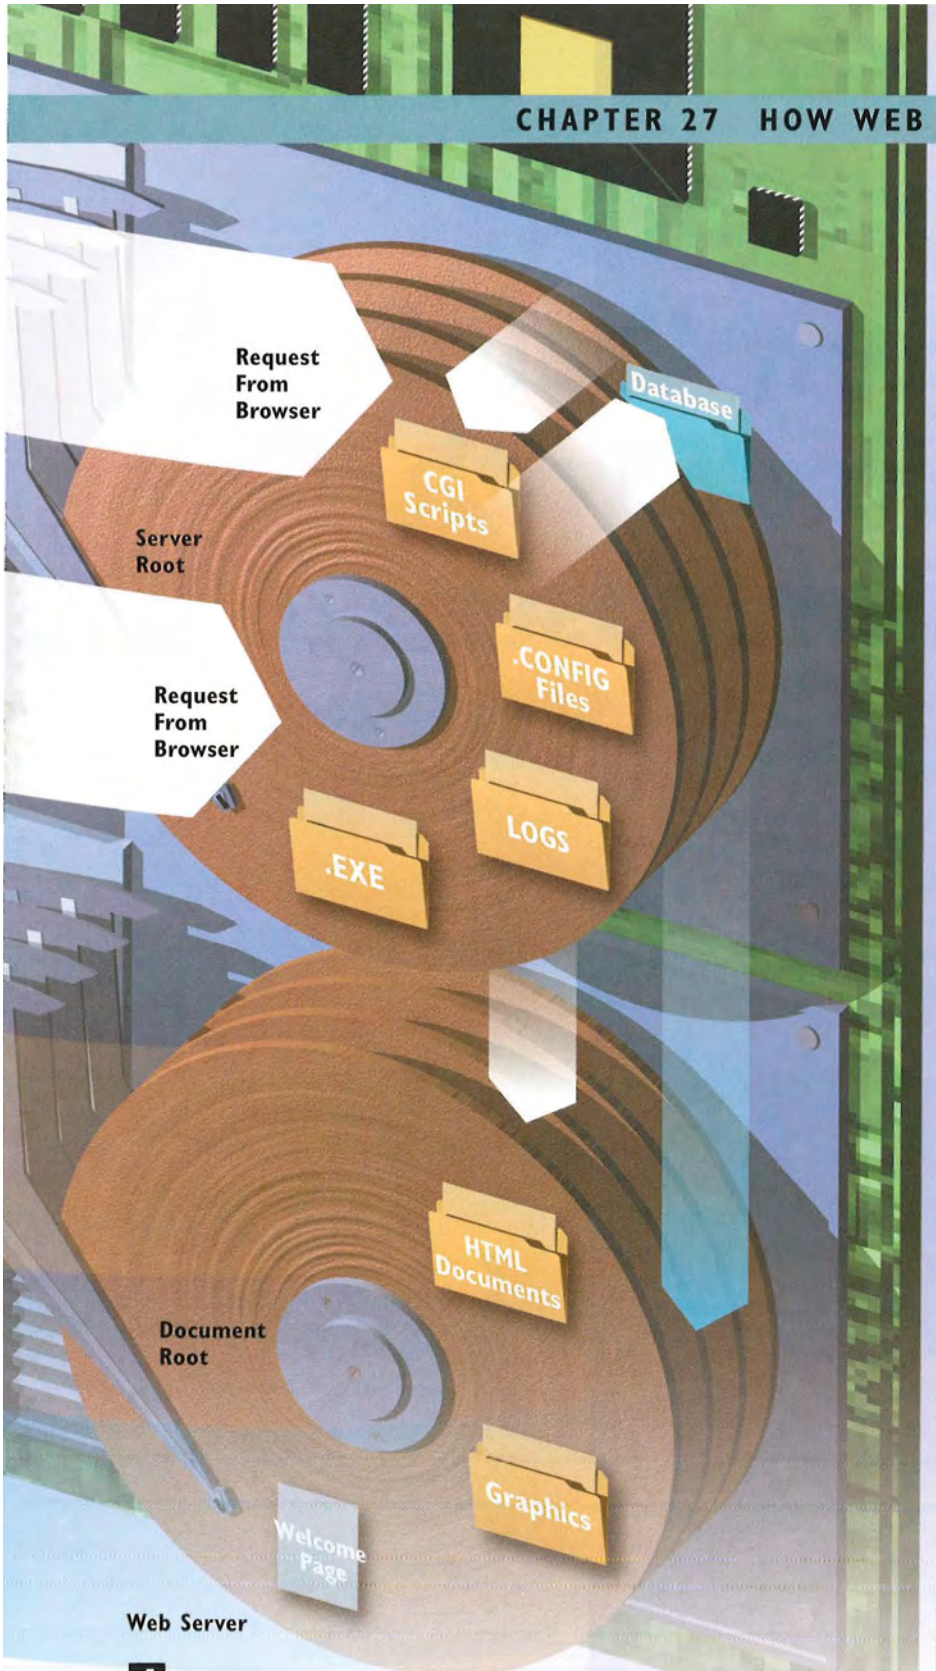
Usually, the client browser requests that the server return an HTML document. The server receives this request and sends back a response. The top portion of the response includes transmission information, and the rest of the response is the HTML file.

A Web server does more than send pages to the browser, however. It passes requests to run Common Gateway Interface (CGI) scripts to CGI applications. These scripts run external mini-programs, such as a database lookup or interactive forms processing. The server sends the script to the application via CGI and communicates the results of the script back to the browser, if appropriate. Moreover, the server software includes configuration files and utilities to secure and manage the Web site in a variety of ways.

How Web Server Software Works



A computer with a single IP address can host several types of servers. This means the address might require a port number to identify the correct server if it is not the IP's default server. Each port is associated with a particular server. Ports are identified by a number from 0 to 65,535, but common server types, such as FTP servers, are given the same number by convention.



4 If the browser has asked for specific database information, the Web server passes a request through CGI to the application, which performs a database lookup, for example. The CGI script returns the results to the Web server, which in turn attaches a header to the data and sends it to the browser.

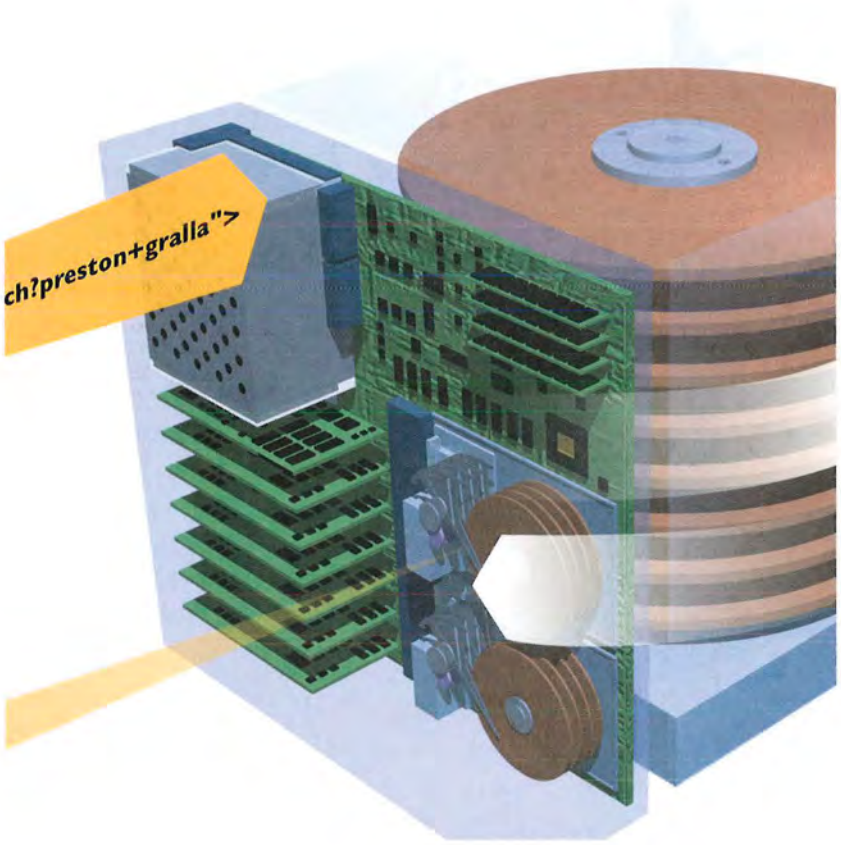


The host computer must have a unique IP address to send and receive information across the Internet. Because raw IP addresses are very intimidating, they are assigned a unique domain name, which is less daunting. The domain name is part of a hierarchical lookup system called the Domain Name System (DNS).

IP Address	=	Domain Name
83.382.1.838		books.que.com

CHAPTER

28 How Web Sites Work with Databases



ONE of the most useful applications of the Web is its capability to link a Web site with a database so Web surfers can search for information. In essence, the Web page becomes the front-end for database applications, enabling you to select search criteria and execute even complex searches of a database that resides on the host computer.

A well-known and widely used example of this type of linking between Web sites and databases is the popular Yahoo! Web site (`www.yahoo.com`). The Yahoo! site serves as a front-end to an extensive database of Web site descriptions, which can be searched according to keywords. The Welcome page includes a search dialog box in which you enter a keyword that represents the subject matter you are looking for. Clicking the Search button sends a request from the browser to the Web server to bring back a list of all Web sites that contain your keyword.

Furthermore, not only can the Web serve up data, it can also collect it. For example, many Web sites ask users to "register" their names, addresses, and other demographic information that is captured and stored in a database.

But how does this all work? You don't have to be a corporate giant—or for that matter even an able programmer—to link your Web site to a database. In fact, linking a Web site to a database can be relatively simple. The database can take just about any form and can be as simple as a FileMaker Pro database or as complex as an Oracle SQL database. The bridge that brings together Web sites and databases is the Common Gateway Interface (CGI).

On the client side of the database, you see a Web page that includes a form in which you enter your search terms. When you execute the search, the Web server passes your search information to a CGI script, which then searches the database. So a search on the Yahoo! site for public relations firms looks like this:

```
http://search.yahoo.com/bin/search?p=public+relations
```

When the Web server receives this URL, it identifies the URL as a trigger for a CGI script (called *script* in this example) and passes it along with the search criteria ("public relations," in this example) to the miniprogram using CGI. The CGI script then sends the search to the database, receives the results of the query, and passes it on to the Web server to be sent back to the client. That's a lot of handing off of requests and data, but typically even a search of a large database is very fast because the majority of Unix and Windows NT databases—the types most often used—can perform these tasks simultaneously. All of this happens behind the scenes, of course—you won't need to do any kind of database work or scripting yourself. Instead, the Web sites you visit have easy-to-use interfaces that take care of interacting with databases; you'll have to type only what you're looking for.