NETWORKING

Guide to TCP/IP

Fourth Edition

Jeffrey L. Carrell
Laura A. Chappell
Ed Tittel
with James Pyles

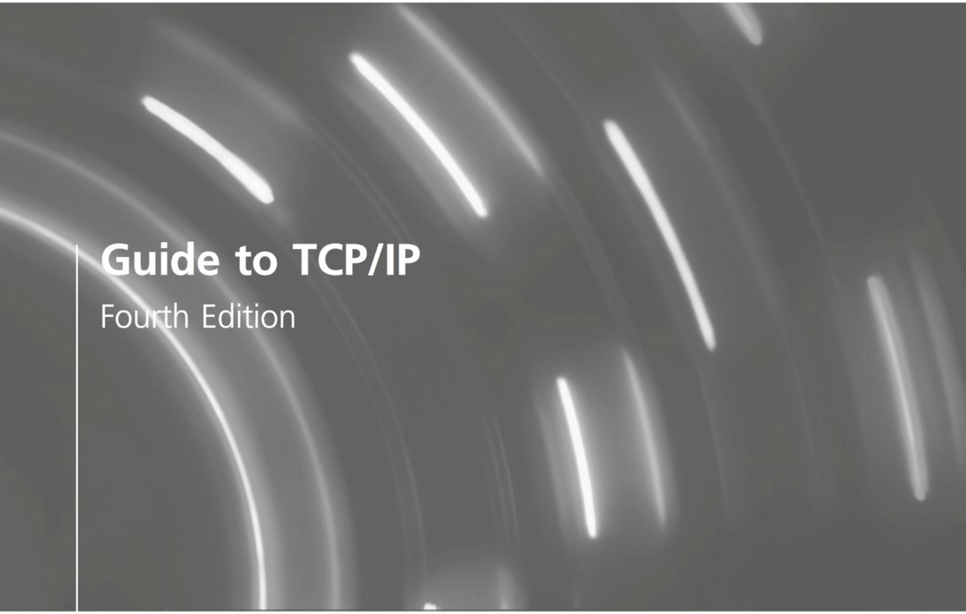

Guide to TCP/IP

Fourth Edition

Jeffrey L. Carrell

Laura A. Chappell

Ed Tittel

with James Pyles

COURSE TECHNOLOGY
CENGAGE Learning®

Australia • Brazil • Japan • Korea • Mexico • Singapore • Spain • United Kingdom • United States

Brief Table of Contents

Table of Contents

CHAPTER 5

Preface

Welcome to *Guide to TCP/IP*, *Fourth Edition*! TCP/IP stands for Transmission Control Protocol/Internet Protocol and defines the broad family of protocols and services that make the Internet able to function as we know it today. In covering TCP/IP, this book offers you real-world examples, interactive examples, and many Hands-On Projects that reinforce key concepts and teach the use of important monitoring and management tools. This book also includes voluminous protocol traces, or decodes, that will help you understand what TCP/IP looks like, and how it behaves, on your networks.

This book offers in-depth coverage of all the salient models, protocols, services, and standards that govern TCP/IP and that guide its behavior on modern networks. Throughout the book, we provide pointed questions to reinforce the concepts introduced in each chapter and to help prepare you to interact with TCP/IP in its native habitat—that is, on the vast majority of networks in use in the world today. In addition to the review questions, we provide detailed Hands-On Projects that provide you with firsthand experience in installing, configuring, using, and managing TCP/IP on a working network. Finally, to put a real-world slant on the concepts introduced in each chapter, we also include Case Projects that pose problems and require creative solutions that should prepare you for the kinds of situations and needs you'll face on a real, live network.

Intended Audience

This book is intended to serve the needs of individuals and information systems professionals who are interested in learning more about working with and on TCP/IP-based networks. These materials have been specifically designed to prepare individuals to take an active role in administering a network infrastructure that uses TCP/IP, either as its only protocol suite or in concert with other protocol suites. Those students who work their way through this entire book should be well equipped to recognize, analyze, and troubleshoot a broad range of TCP/IP-related networking problems or phenomena.

Chapter Summaries

Chapter 1, "Introducing TCP/IP," presents the broad outlines of TCP/IP's capabilities and identifies its most important constituent elements—namely, the protocols and services that TCP/IP provides. In addition, it explores the Open Systems Interconnection (OSI) reference model for networking, as standardized by the International Organization for Standardization (ISO), and compares and contrasts this standard model to the model around which TCP/IP is built. This chapter then covers the structure and origins of the standards documents known as Requests for Comments (RFCs), which describe and govern TCP/IP protocols, services, and practices. The chapter concludes with an overview of the key tool that will play a significant role throughout the remainder of the book—a special software utility called a protocol analyzer that captures, unpacks, and displays the contents of traffic on a network, including TCP/IP. In this book, we use a protocol analyzer named Wireshark.

Chapter 2, "IP Addressing and Related Topics," covers the intricacies involved in managing unique IP addresses for both 32-bit IPv4 addresses and 128-bit IPv6 addresses. Beginning with the anatomy of a numeric IPv4 address, the chapter explores IPv4 address classes, special cases such as broadcast and multicast addresses, subnets and supernets, and reviews the reasons for classless IPv4 addressing, public versus private IPv4 addresses, and IPv4 addressing schemes. The rest of the chapter repeats this coverage for IPv6, including a review of address formats and notation, address layouts and types, and address allocations. You'll also find addressing schemes and subnetting considerations covered, along with some discussion about how to manage the transition from IPv4 to IPv6 addresses.

Chapter 3, "Basic IP Packet Structures: Headers and Payloads," covers the key components of any IP packet (both for IPv4 and IPv6): the header that describes the packet for routing, forwarding, and filtering, and the payload that contains the data that the packet is meant to convey. IPv4 and IPv6 headers are laid out and dissected in detail, including IPv6 Extension Headers, and the use of transport and packet handling controls are described and explored. The chapter concludes with a comparison of header structures in IPv4 versus IPv6, with a rationale to explain redesign and changes involved.

Chapter 4, "Data Link and Network Layer Protocols in TCP/IP," explores and explains the TCP/IP protocols that operate at the Data Link and Network layers in the OSI reference model. In that context, it discusses data link protocols in general, examines IP frame types, and talks about hardware addresses in the IP environment and the various protocols—particularly ARP and RARP, for IPv4, and the Neighbor Discovery Protocol, or NDP, for IPv6—that support their use. The chapter also covers TCP/IP's most important protocol at the Network layer, the Internet Protocol, along with routing protocols, mechanisms, and characteristics for IPv4 and IPv6, including RIPv1 and v2, OSPF, EIGRP, and BGP, with considerations for both IPv4 and IPv6 protocols and behaviors.

Chapter 5, "Internet Control Message Protocol," covers a key Network layer protocol for TCP/IP whose job is to ferry status and error messages about IP traffic back to its senders and to other "interested devices," such as routers or switches. Starting with a review of ICMPv4 and ICMPv6 structures and functions, this chapter examines ICMP testing and troubleshooting methods, security issues, and ICMP message types and codes, and concludes with a thorough review of testing and troubleshooting sequences for ICMP and decoding ICMP packets.

Chapter 6, "Neighbor Discovery in IPv6," digs into NDP to explain how neighbor discovery works on IPv6 networks. Topics covered include comparing NDP to related IPv4 protocols, various NDP message formats and options, and the overall neighbor discovery process on IPv6 networks.

Chapter 7, "IP Address Autoconfiguration," describes various auto-addressing schemes and mechanisms used on IPv4 and IPv6 networks, including the Dynamic Host Configuration Protocol, or DHCP, as well as autoconfiguration mechanisms used for IPv4 (APIPA and DHCP) and IPv6 (host/interface address determination, stateless and stateful address autoconfiguration, and DHCPv6).

Chapter 8, "Name Resolution on IP Networks," deals with key services used to resolve symbolic, human-readable network names and addresses into machine-intelligible network addresses. Topics covered include name resolution fundamentals and various network name resolution protocols. IPv4 and IPv6 name resolution via the Domain Name Service, or DNS, is described in detail, as is name resolution support for Windows operating systems, including issues related to setup, configuration, troubleshooting, and relevant utilities.

Chapter 9, "TCP/IP Transport Layer Protocols," covers two key protocols that operate at the Transport layer of the OSI reference model: the heavy-duty, robust, reliable Transmission Control Protocol (TCP) and the lighter-weight but faster User Datagram Protocol (UDP). TCP is examined in great detail, with particular attention on its packet structures and functions (including IPv6 extension headers for TCP), whereas UDP gets the brief coverage it deserves. The chapter concludes with a discussion of common and appropriate uses for these two protocols.

Chapter 10, "Transitioning from IPv4 to IPv6: Interoperation," deals with issues and techniques that apply when IPv4 and IPv6 must coexist on the same networks, as will surely be the case for many networks for the foreseeable future. It explains the means whereby IPv4 and IPv6 can interact, explains hybrid IPv4/IPv6 networks and node types, and explores transition addresses and switchover mechanisms to make the change from IPv4 to IPv6 as straightforward as possible. Tunneling mechanisms and protocols, including ISATAP, 6to4, and Teredo, are described in detail.

Chapter 11, "Deploying IPv6," jumps into an area of great interest to Internet professionals—namely, what's involved in understanding, planning, deploying, and using IPv6 on modern TCP/IP networks. Topics covered include evaluating potential software and hardware changes, addressing schemes and autoaddressing, and priority schemes for various classes or types of network services.

Chapter 12, "Securing TCP/IP Environments," covers general network security basics, with a particular emphasis on IP security topics. It also addresses key topics that include perimeter security, infrastructure security, and host device security.

The book also includes **Appendix A,** which explains the required software and trace files available on the book's online resources Web site. In addition, this site includes the following, plus much more:

- A list of the important RFCs mentioned throughout the text and the available IPv6-specific RFCs

NOTE

RFCs are a dynamic collection of documents, so anything collected in static form represents a snapshot of what was current at the time the snapshot was taken. Always consult online RFCs for information about the most current documents and standards.

- A reference to TCP/IP-related command-line utilities for Windows desktop and Windows Server
- A list of the Windows desktop and Windows Server Registry settings found in numerous tables in this book

Features

To ensure a successful learning experience, this book includes the following pedagogical features:

- **Chapter Objectives**: Each chapter in this book begins with a detailed list of the concepts to be mastered within that chapter. This list provides you with a quick reference to the contents of that chapter as well as a useful study aid.
- **Illustrations and Tables**: Numerous illustrations of server screens and components aid you in the visualization of common setup steps, theories, and concepts. In addition, many tables provide details and comparisons of both practical and theoretical information and can be used for a quick review of topics. This book also includes a great number of protocol traces from both IPv4 and IPv6 protocols. Because of formatting differences between the two protocol families, these traces differ slightly, but they present more or less the same information, subject only to minor differences.
- **End-of-Chapter Material:** The end of each chapter includes the following features to reinforce the material covered in the chapter:
  - **Summary:** A bulleted list providing a brief but complete summary of the chapter
  - **Key Terms List:** A list of all new terms and their definitions
  - **Review Questions:** A list of review questions that test your knowledge of the most important concepts covered in the chapter
  - **Hands-On Projects:** Projects that help you to apply the knowledge gained in the chapter
  - **Case Study Projects:** Projects that take you through real-world scenarios
- **Student and Instructor Online Resources:** The book's online resources Web site provides self-extracting files that contain the trace (data) files and the software required to work through the Hands-On Projects in this book—Wireshark for Windows and the Bitcricket IP Subnet Calculator. In addition, you will find documents containing descriptions of other handy networking tools and utilities. Student and instructor resources for this book are available at *www.cengage.com*. To locate the resources, search for **Guide to TCP/IP** in the Higher Education catalog.

Text and Graphic Conventions

Wherever appropriate, additional information and exercises have been added to this book to help you better understand what is being discussed in the chapter. Icons throughout the text alert you to additional materials. The icons used in this textbook are as follows:

CAUTION

The Caution icon warns you about potential mistakes or problems and explains how to avoid them.

The Note icon is used to present additional helpful material related to the subject being described.

Tips based on the authors' experience provide extra information about how to attack a problem or what to do in real-world situations.

Each Hands-On Project in this book is preceded by the Hands-On icon and a description of the exercise that follows.

Case Project icons mark the case projects. These are more involved, scenario-based assignments. In this extensive case example, you are asked to implement independently what you have learned.

Instructor Resources

The following supplemental materials are available when this book is used in a classroom setting. All the supplements available with this book are provided to the instructor on a single CD-ROM (ISBN: 978-1-1330-1987-9) and online at www.cengage.com.

Electronic Instructor's Manual. The Instructor's Manual that accompanies this textbook includes additional instructional material to assist in class preparation, including suggestions for classroom activities, discussion topics, and additional projects.

Solutions. The answers to all end-of-chapter material, including the Review Questions and, where applicable, Hands-On Projects and Case Projects, are provided.

ExamView®. This textbook is accompanied by ExamView, a powerful testing software package that allows instructors to create and administer printed, computer (LAN-based), and Internet exams. ExamView includes hundreds of questions that correspond to the topics covered in this text, enabling students to generate detailed study guides that include page references for further review. The computer-based and Internet testing components allow students to take exams at their computers and also save the instructor time by grading each exam automatically.

PowerPoint presentations. This book comes with Microsoft® PowerPoint® slides for each chapter. These are included as a teaching aid for classroom presentation, to make available to students on the network for chapter review, or to be printed for classroom distribution. Instructors, please feel to add your own slides for additional topics you introduce to the class.

Figure Files. All the figures in the book are reproduced on the Instructor Resources CD, in bitmap format. Similar to the PowerPoint presentations, these are included as a teaching aid for classroom presentation, to make available to students for review, or to be printed for classroom distribution.

Acknowledgments

The authors would like to thank Course Technology for this opportunity to revise *Guide to TCP/IP* to include detailed coverage of IPv6. We deeply appreciate their patience and indulgence, especially that of Nick Lombardi, our acquisitions editor; Natalie Pashoukos, our product manager; Brooke Baker, our content project manager and Susan Pedicini, our technical editor in charge of manuscript quality assurance. Thanks also to Kent Williams, our wonderful developmental editor, whose in-depth and detailed work turned these materials into the finely polished form they now take.

The authors would also like to thank the behind-the-scenes author team who helped to bring this book to fruition, such as Tom Lancaster, who provided Chapter 11, and especially James Pyles, who revised several chapters and added plenty of new information about IPv6. His diligence and hard work earned him a spot on the cover. Thanks also to Kim Lindros and Mary Kyle, who jumped in to help manage the project on behalf of the authors.

Jeff Carrell: With God's help, all things are possible. Thank you to my wife and best friend Cynthia for all your love, encouragement, and patience; I am truly blessed to have you in my life. Thank you to my friends and colleagues who provided input and encouragement along the way. Thanks to Ed Tittel for the opportunity, inspiration, and mentoring. The project was huge, exciting, and awesome to be a part of. Thanks to Kim Lindros and Mary Kyle who not only kept us moving but made sure we stayed on task. I could not imagine working on this project without y'all. Finally, thanks to James Pyles and Tom Lancaster, who provided updates and much new content; we could not have completed this project without you guys.

Ed Tittel: My profound thanks to Jeff Carrell for taking over as lead author and architect for the fourth edition of this book, and to James Pyles and Tom Lancaster for helping us provide new content and exercises. Also, thanks again to Kim Lindros and Mary Kyle for making this book so much easier to finish than it was to start without their able and competent presence. Finally, thank you to my lovely wife Dina and son Gregory, who have brought me much joy and happiness.

Laura Chappell: Special thanks to Ed Tittel and Jeff Carrell for their enthusiasm and wonderful writing efforts on this book. Thanks also to James Pyles and Tom Lancaster for their tremendous technical contributions to this title—this book could not have been completed in such a timely manner without your assistance. Also, very special and sincere thanks to Kent Williams and Kim Lindros for keeping all the loose ends tied on this project. Finally, my deepest thanks to my children, Scott and Ginny, who make me laugh and enjoy life way beyond the packet-level.

James Pyles: I appreciate the opportunity to contribute to updating this fine book for its fourth edition. I especially want to thank Ed Tittel and Kim Lindros for inviting me along for the ride. I also am extremely grateful to Jeff Carrell, at whose feet I would gladly sit at any time in order to learn the arcane mysteries of IPv6. I also want to thank Mary Kyle for her excellent management skills and her infinite patience, as well as Tom Lancaster for his invaluable contributions.

No acknowledgement would be complete without my recognition of my lovely wife Lin and the invaluable support she has provided me in all of my endeavors. Given the rapidly evolving nature of the Internet as we proceed into the twenty-first century, I can't help but think of my three-year old grandson Landon, who will inherit the future from us. May we leave him and his generation a worthy legacy.

Readers are encouraged to e-mail comments, questions, and suggestions regarding *Guide to TCP/IP, Fourth Edition* and the accompanying student and instructor resources Web site to *tcpip4e@networkconversions.com.*

Read This Before You Begin

To the User

This book is intended to be read in sequence, from beginning to end. Each chapter builds upon those that precede it to provide a solid understanding of TCP/IP concepts, protocols, services, and deployment practices. Readers are also encouraged to investigate the many pointers to online and printed sources of additional information that are cited throughout this book.

Some of the chapters in this book require additional materials to complete the end-of-chapter projects. The student and instructor resources Web site for this book contains the necessary supplemental files. To download the resources, go to *www.cengage.com* and search for **Guide to TCP/IP** in the Higher Education catalog.

This Web site includes:

- Software required to complete the Hands-On Projects, which includes Bitcricket IP Calculator and the Wireshark for Windows protocol analyzer
- A link to the Student Data Files (referred to as “trace” or “packet” files in this book) required to complete the Hands-On Projects
- Additional resources for topics in select chapters

NOTE

This book was written using the popular Wireshark protocol analyzer. The Wireshark version used in the Hands-On Projects is available for download from the companion Web site for this book. You may also download the latest version of the software from the Wireshark Web site at *www.wireshark.org*.

To the Instructor

When setting up a classroom lab, make sure each workstation has Windows Vista or Windows 7 Professional, Internet Explorer 9 or later, and a network interface controller (NIC) capable of working in promiscuous mode. Students will install Wireshark and the Bitcricket IP Subnet Calculator on these computers in the course of working through the book. In addition, students will need administrative rights on their workstations to perform many of the operations covered in the Hands-On Projects throughout the book. Students will also need access to Windows Server 2008 R2 for a small number of projects.

Coping with Change on the Web

Sooner or later, at least a few of the Web links in the book will go stale or be replaced by newer information. In that case, there's always a way to find what you want on the Web, if you're willing to invest some time and energy. To begin with, most large or complex Web sites—and Microsoft's qualifies on both counts—offer a search engine. As long as you can get to the site itself, you can use this tool to help you find what you need.

Finally, feel free to use general search tools such as *www.google.com*, *www.bing.com*, or *www.yahoo.com* to find information related to topics in this book. For example, although certain standards bodies may offer the most precise and specific information about their standards online, there are plenty of third-party sources of information, training, and assistance in this area that do not have to follow the party line like a standards group typically does. The bottom line: If you can't find something where this book says it's supposed to be, start looking around. It's got to be around there somewhere! If you find it on your own, send an e-mail to *tcpip4e@networkconversions.com* and we will do our best to update this book's companion Web site in short order. Plus, you will have the satisfaction of knowing you helped all the instructors and students who are using this book.

Lab Requirements

Following are the recommended hardware and software requirements to perform the end-of-chapter projects:

- 1 GHz CPU or higher, 2 GB of RAM, 80 GB hard disk with at least 2 GB of storage available
- CD-ROM drive
- NIC in promiscuous mode connected to a LAN
- Windows Vista or Windows 7 Professional (Service Pack 1 or later) and Internet Explorer 9 or later
- Access to a Windows Server 2008 R2 system with TCP/IP installed and configured; an IP address must be defined either statically or via DHCP
- Internet access

Ideally, you should have two computers with the same hardware specifications listed above—one running Windows 7 and one running Windows Server 2008 R2, as well as a Layer 3 switch or router that supports both IPv4 and IPv6.

chapter 1

Introducing TCP/IP

After reading this chapter and completing the exercises, you will be able to:

- Describe TCP/IP's origins and history
- Explain the process by which TCP/IP standards and other documents, called Requests for Comments (RFCs), are created, debated, and formalized (where appropriate)
- Describe the "huge difference" between IPv4 and IPv6 and explain why a switch to IPv6 is both necessary and inevitable
- Describe the Open Systems Interconnection network reference model, often used to characterize network protocols and services, and how it relates to TCP/IP's own internal networking model
- Define the terms involved and explain how TCP/IP protocols, sockets, and ports are identified
- Describe data encapsulation and how it relates to the four layers of the TCP/IP protocol stack
- Describe and apply the basic practices and principles that underlie network protocol analysis

This chapter introduces the background and history of the collection of networking **protocols** known as **TCP/IP**, which is an abbreviation for **Transmission Control Protocol/Internet Protocol**. Two of the most important protocols in the overall collection known as TCP/IP give their names to this protocol collection—namely, the **Transmission Control Protocol (TCP)**, which handles reliable delivery for messages of arbitrary size, and the **Internet Protocol (IP)**, which manages the **routing** of network transmissions from sender to receiver, among other capabilities.

In addition, this chapter covers TCP/IP's networking model, its various ways of identifying specific protocols and services, how TCP/IP standards are defined and managed, and which elements of the TCP/IP collection are most noteworthy. It also includes coverage of the original version of TCP/IP, sometimes called IPv4, and the newer versions, known as IPv6. The chapter concludes with an overview of the art and science of protocol analysis, which uses special tools to gather data directly from a network itself, characterize a network's traffic and behavior, and examine the details inside the data that's moving across a network at any given point in time.

What Is TCP/IP?

The large collection of networking protocols and services called TCP/IP comprises far more than the combination of those two key protocols that gives this collection its name. Nevertheless, these two protocols deserve an introduction. Transmission Control Protocol, or TCP, offers reliable delivery for messages of arbitrary size, and it defines a robust delivery mechanism for all kinds of data across a network. Internet Protocol, or IP, manages the routing of network transmissions from sender to receiver, along with issues related to network and computer addresses, and much more. Together, these two protocols ferry the vast bulk of data that moves across the Internet, even though they represent only a tiny fraction of the total TCP/IP protocol collection.

To gain a better appreciation for the importance of TCP/IP, keep in mind that anyone who uses the Internet must also use TCP/IP because the Internet runs on TCP/IP. Its roots run deep and long, as computing technologies go—all the way back to 1969. Knowing where TCP/IP comes from and what motivated its original design can enhance one's understanding of this essential collection of protocols (often called a **protocol suite**). For that reason, we will explore this protocol suite's roots and design goals in the following section.

The Origins and History of TCP/IP

In 1969, an obscure arm of the United States Department of Defense (DoD), known as the **Advanced Research Projects Agency (ARPA)**, funded an academic research project involving a special type of long-haul (long-distance) network, called a **packet-switched network**. In a packet-switched network environment, individual chunks of data (called **packets**) can take any usable path between the sender and receiver. The sender and receiver are identified by unique network addresses, but the packets are not required to follow the same path in transit (although they often do). The network built as a result of this project was known as the **ARPANET**.

TCP/IP's Design Goals

The design of the ARPANET and protocols that evolved to support it were based on the following government needs:

- *A desire for a communications network with the ability to withstand a potential nuclear strike*—This explains the need for packet switching, in which the routes from sender to receiver can vary as needed, as long as a valid route exists. It also explains why, in a world that could blow up at any time, robust and reliable delivery was a concern.
- *A desire to permit different kinds of computer systems to communicate easily with one another*—Because proprietary networking was the order of the day, and because the government owned many different kinds of incompatible networks and systems, it was necessary that this technology permit dissimilar systems to exchange data.
- *A need to interconnect systems across long distances*—The late-1960s was an era of "big iron," in which large, expensive individual systems with terminals dominated the computer landscape. At that time, interconnecting multiple systems meant interconnecting far-flung locations. Thus, the original ARPANET linked systems at the Stanford Research Institute (SRI), the University of Utah in Salt Lake City, and campuses in the University of California system at Los Angeles and Santa Barbara.

These design goals may not seem terribly important in the early twenty-first century. That's because the concern about a global nuclear holocaust has largely subsided and networking is now taken for granted. Likewise, high-bandwidth, long-distance data communication is a big business. However, some would argue that the Internet is responsible for the prevalence of high-bandwidth, long-distance data communication in today's modern world!

TCP/IP Chronology

TCP/IP appeared on the scene in the 1970s. By that time, early networking researchers realized that data had to be moved across different kinds of networks as well as among multiple locations. This was especially necessary to permit **local area networks (LANs)**, such as those using **Ethernet**, to use long-haul networks, such as the ARPANET, to move data from one local network to another. Although work on TCP/IP began in 1973, it wasn't until 1978 that **Internet Protocol version 4** (also known as **IPv4**—the very same version used today on most TCP/IP networks) was developed.

The original Internet (notice the initial capital letter) helped establish a model for a network composed of other networks. Thus, the term **internetwork** (notice the lack of an initial capital letter) refers to a single logical network composed of multiple physical networks, which may all be at a single physical location or may be spread around multiple physical locations. We distinguish the "Internet," a proper name for the worldwide collection of publicly accessible networks that use TCP/IP, from an "internetwork," which can appear anywhere in the world and may or may not be part of the Internet (and may not use TCP/IP, even though the majority do).

In 1983, the Defense Communications Agency (DCA, now known as the **Defense Information Systems Agency**, or **DISA**) took over operation of the ARPANET from DARPA (Defense Advanced Research Projects Agency). This allowed more widespread use of the Internet, as more colleges and universities, government agencies, defense facilities, and government contractors began to rely on it to exchange data, e-mail, and other kinds of information. In the same year, the DoD instituted its requirement that all computers on the Internet switch to TCP/IP from a hodgepodge of earlier, mostly experimental protocols that had been used on the ARPANET since its inception. In fact, some people argue that the Internet and TCP/IP were born at more or less the same time.

By no coincidence whatsoever, 1983 also was the year that the Berkeley Software Distribution version of UNIX known as 4.2BSD incorporated support for TCP/IP in the operating system. There are those who argue that this step, which exposed computer professionals at colleges and universities around the world to TCP/IP, helps explain the birth and proliferation of Internet protocols and how they became the behemoths they are today.

At roughly the same time—we're still in 1983—the all-military MILNET was split off from the ARPANET. This divided the infrastructure of the Internet into a military-only side and a more public, freewheeling side that included all nonmilitary participants. Also in 1983, the development at the University of Wisconsin of name server technology, which allowed users to locate and identify network addresses anywhere on the Internet (this remains a hallmark of its operation to this day), capped off a banner year in the Internet's history.

After that, the Internet and TCP/IP experienced a series of landmarks that led to the global Internet we know today. Here are some additional highlights:

- 1986—The **National Science Foundation (NSF)** launches a long-haul, high-speed network known as **NSFNET**, with a network backbone running at 56 Kbps. NSF also imposes a set of policies known as **Acceptable Use Policies (AUPs)**, which governs Internet use and sets the tone for how users interact on the Internet.
- 1987—The number of hosts on the Internet breaks 10,000.
- 1989—The number of hosts on the Internet breaks 100,000.
  - The NSFNET backbone is upgraded to T1 speeds, at 1.544 megabits per second (Mbps).
- 1990—ARPANET ceases doing business under that name, and commercial enterprises, academic institutions, and government organizations begin supporting the Internet.
  - Work begins on the **Hypertext Transfer Protocol (HTTP)**; the notion of the World Wide Web is born at the **Centre Europeen de Researche Nucleaire (CERN)** in Switzerland.
- 1991—The **Commercial Internet Exchange (CIX)**, a consortium of Internet operators, system providers, and other commercial operations with Internet interests, is formed.
- 1992—The Internet Society (ISOC) is chartered.
  - The number of hosts on the Internet breaks one million.
  - The NSFNET backbone is upgraded to T3 speeds, at 44.736 Mbps.
  - CERN releases HTTP and Web server technology to the public ("birth of the Web").

- 1993—The **Internet Network Information Center (InterNIC)** is chartered to manage **domain** names.
  - The first-ever, high-powered graphical browser, Mosaic, emerges from the **National Center for Supercomputing Applications (NCSA)**. This starts the Web revolution.
  - The U.S. White House goes online at *www.whitehouse.gov*.
- 1994—The U.S. Senate and House of Representatives establish Internet Web servers.
  - Online junk mail and shopping malls begin to proliferate.
- 1995—Netscape launches Netscape Navigator and begins the commercialization of the Web.
  - The number of hosts on the Internet breaks five million.
- 1996—Microsoft launches its Internet Explorer Web browser, even though Netscape dominates the Web browser marketplace.
- 1997—The number of registered domain names breaks 31 million.
- 2000—The Love Letter worm infects over one million personal computers.
- 2001—The number of hosts on the Internet breaks 150 million.
  - Sircam virus and Code Red worm infect thousands of Web servers and e-mail accounts.
- 2002—The number of hosts on the Internet breaks 204 million.
  - The Internet2 backbone utilizes native IP version 6.
- 2003—Public Interest Registry (PIR) assumes responsibility as the .org registry operator.
- 2005—The number of hosts on the Internet breaks 250 million.
- 2008—The number of hosts on the Internet breaks 600 million.
- 2009—The number of hosts on the Internet breaks one billion, and the number of Chinese users surpasses the number of U.S. users.

Today, there are few aspects of commerce, communications, and information access that do not involve the Internet in one way or another. Living without e-mail, the Web, and online e-commerce has become unthinkable. As we progress through the twenty-first century, new Internet services and protocols continue to appear, but TCP/IP keeps going strong.

NOTE

For more information about the fascinating history of the Internet, visit the ISOC's "A Brief History of the Internet" Web page at *www.isoc.org/internet/history/brief.shtml*.

Who "Owns" TCP/IP?

Given that its roots are everywhere and its reach is unlimited, who owns and controls TCP/IP can seem puzzling. Even though TCP/IP and related protocols are under the purview of specific standards-making bodies, which we'll discuss later, TCP/IP also falls into the public domain because it's been publically funded since its inception. In other words, both everybody and nobody owns TCP/IP.

Standards Groups That Oversee TCP/IP

The standards groups involved with TCP/IP are:

- **Internet Society (ISOC)**—This is the parent organization for all Internet boards and task forces. It is a nonprofit, nongovernmental, international, professional membership organization funded through membership dues, corporate contributions, and occasional support from governments. For more information, visit *www.isoc.org*.
- **Internet Architecture Board (IAB)**—Also known as the Internet Activities Board, this arm of the ISOC is the parent organization for standards-making and research groups that handle current and future Internet technologies, protocols, and research. The IAB's most important tasks are to provide oversight for the architecture of all Internet protocols and procedures and to supply editorial oversight regarding the documents known as Requests for Comments (RFCs), in which Internet Standards are stated. For more information, visit *www.iab.org*.
- **Internet Engineering Task Force (IETF)**—This is the group responsible for drafting, testing, proposing, and maintaining official Internet Standards (in the form of RFCs) through the participation of multiple working groups under its purview. The IETF and the IAB use a process accurately described as "rough consensus" to create Internet Standards. This means that all participants in the standards-making process, a type of peer review process, must more or less agree before a standard can be proposed, drafted, or approved. Sometimes, that consensus can be pretty rough indeed! For more information, visit *www.ietf.org*.
- **Internet Research Task Force (IRTF)**—This group handles the more forward-looking activities of the ISOC, including research and development work for topics too far out or impractical for immediate implementation but that may play a role on the Internet some day. For more information, visit *http://irtf.org/*.
- **Internet Corporation for Assigned Names and Numbers (ICANN)**—This group has the ultimate responsibility for managing Internet domain names, network addresses, and protocol parameters and behaviors. However, it delegates the management of customer interaction, money collection, database maintenance, and so forth to commercial authorities. For more information, visit *www.icann.org*. Also, there's a list of accredited and accreditation-qualified name registrars on the ICANN site at *www.icann.org/registrars/accredited-list.html*.

Of all these organizations, the most important one for TCP/IP is the IETF because it is responsible for creating and managing RFCs, in which the rules and formats for all related protocols and services are described.

IPv4 and IPv6

By the time TCP/IP had established itself, in the mid- to late-1980s, IPv4 was the only Internet protocol around. It uses 32-bit addresses, which means it supports around four billion distinct network addresses, of which over three billion are usable on the public Internet. At the time, this address space seemed inexhaustible. However, by the early 1990s, when the public Internet became a global phenomenon, it became obvious that IPv4 addresses would run out someday.

By February 2011, ICANN had dispensed its last few unallocated Class C address blocks. (These are explained in Chapter 2.) By June 2011, all those addresses had been assigned to specific organizations. Thus, the entire IPv4 address space is now occupied, and the only way to obtain a public IPv4 address these days is to buy or otherwise acquire an address from some other organization or user.

IPv6, on the other hand, supports 128-bit addresses, which means that its address space is roughly $3.4 * 10^{38}$, whereas IPv4's is roughly $4.3 * 10^{9}$; that means the IPv6 address space is roughly $8 * 10^{28}$ larger than the IPv4 space. This difference is almost too large to comprehend, but here's a way to put it in perspective: these days, Internet service providers who supply IPv6 addresses to customers routinely supply them with what's called a /64 (pronounced "slash 64") public IPv6 address. That means that each individual user gets a 64-bit address. Each individual address space is 4.3 billion times larger than the entire IPv4 address space.

There's no question that the future of TCP/IP networking involves a switchover to IPv6. That's why this fourth edition of our book covers IPv6 topics in detail. Because IPv4 won't disappear in the foreseeable future, however, this edition provides information you will need to work with IPv4-only networks, IPv6-only networks, or the increasingly likely IPv4/IPv6 hybrids.

TCP/IP Standards and RFCs

Although "**Request for Comments**" sounds like a pretty tentative name for a document, the impact of RFCs on TCP/IP is nothing less than overwhelming. After going through a multistep process of proposals, drafts, test implementations, and so forth, they become official standards, providing the documentation necessary to implement and use TCP/IP protocols and services on the Internet.

Older versions of RFCs are often replaced by newer, more up-to-date versions. Each RFC is identified by a number, the most recent of which are in the 6300 range. (Visit *www.rfc-editor.org/new_rfcs.html* to see the latest ones.) When two or more RFCs cover the same topic, they usually share the same title; the RFC with the highest number is considered the current version, and all the older, lower-numbered versions are considered obsolete.

One special RFC is titled Internet Official Protocol Standards. It provides a snapshot of current prevailing standards and best practices documents. If you visit your favorite Internet search engine (e.g., Yahoo! or Google), you can find many online locations for Internet RFCs. We recommend using, say, the search string "RFC 5000" to find RFC 5000, or "RFC 2026" to find RFC 2026. (Depending on which search engine you use, you may need to put the entire string in quotation marks.) We recommend the Internet RFC/STD/FYI/BCP Archives site, where an index for all RFCs is available at *www.faqs.org/rfcs/*.

RFC 2026 is another important document. It describes how an RFC is created and what processes it must go through to become an official standard adopted by the IETF. It also describes how to participate in that process.

A potential Standard RFC begins its life when a process or protocol is developed, defined, and reviewed, and it is then tested and reviewed further by the Internet community. After it is revised, tested further, proven to work, and shown to be compatible with other Internet Standards, it may be adopted as an official Standard RFC by the IETF. It is then published as a Standard RFC and assigned a number.

Actually, an RFC passes through several specific phases while becoming a standard and acquires specific status designations during the process. These are fully defined in RFC 2026. For example, a potential Standard RFC goes through three phases on its way to becoming a standard. It starts as a **Proposed Standard**, moves up to a **Draft Standard**, and, if formally adopted, becomes an **Internet Standard**, or a Standard RFC. Eventually, when replaced by a newer RFC, it can be designated an **Historic Standard**.

Best Current Practice (BCP) is another important category of RFCs. A BCP does not define a protocol or technical specifications; rather, it defines a philosophy, or a particular approach, to a network design or implementation that is recommended as tried and true, or that enjoys certain desirable characteristics worthy of consideration when building or maintaining a TCP/IP network. BCPs are not standards per se, but because they present highly recommended design, implementation, and maintenance practices, they are well worth reading and applying where appropriate.

OSI Network Reference Model Overview

Before discussing TCP/IP protocols and services in further detail, let's explore a model of how networks operate in general. This will help you better understand what protocols are for and what roles they play on contemporary networks. This kind of model is often called a **network reference model**, but it is formally known as the **International Organization for Standardization Open Systems Interconnection** network reference model and may sometimes be called the **ISO/OSI network reference model**. No matter what you call it, its job is to describe how networks operate from the hardware or device level (signals and bits) all the way to the software or program level (application interfaces).

Governed by ISO Standard 7498, the ISO/OSI network reference model, also known as the reference model or the seven-layer model, was developed as part of an international standards initiative in the 1980s that was intended to usher in a new and improved suite of protocols to replace TCP/IP. Although the OSI protocols were never widely adopted outside Europe, the network reference model provides a standard way to talk about networking and explain how networks operate. Despite the 10-year, multibillion-dollar effort to complete the OSI protocols and services, TCP/IP remained the open standard protocol suite of choice and remains so to this day.

Breaking Networking into Layers

The network reference model's value lies in its ability to break a big technical problem into a series of interrelated subproblems and then solve each subproblem individually. Computer scientists call this approach **divide and conquer**.

The network reference model handles networking all the way from hardware to the high-level software involved in making networks work. The divide-and-conquer approach keeps concerns related to networking hardware completely separated from those related to networking software. It even further divides networking software into multiple **layers**, each of which represents a separate kind or class of networking activity. (You will read more about this in the section titled "The ISO/OSI Network Reference Model.") Thus, the whole series of independent but interconnected layers of hardware and software work together to enable one computer to communicate with another across a network.

In fact, a layered approach to networking is a good thing. That's because the kind of expertise that makes it possible for an electrical engineer to specify how a network medium must behave, as well as specify what kinds of physical interfaces are necessary to attach to such a networking medium, is quite different from the kinds of expertise that software engineers need. In fact, software engineers must not only write drivers for network interfaces, they must implement the networking protocols that operate at various layers in the network reference model (or in another layered model for whatever networking protocols may be in use).

Before we dive into the details of the network reference model and describe its layers, you should understand and appreciate these key points about networking:

- The challenges of networking are easier to overcome when big tasks are broken into a series of smaller tasks.
- Layers operate more or less independently of one another, enabling modular design and implementation of specific hardware and software components that perform individual network functions.
- Because individual layers encapsulate specific, largely independent functions, changes to one layer need not affect other layers.
- Individual layers work together on pairs of computers. The sending computer performs operations on one layer that are in some sense "reversed" or "undone" by the operations performed at the same layer on the receiving computer. Because such layers work in concert across the network, they are called **peer layers**.
- Different expertise is needed to implement the solutions necessary for the networking functions or tasks handled at each layer.
- The layers in a network implementation work together to create a general solution to the general problem known as networking.
- Network protocols usually map to one or more layers of the network reference model.
- TCP/IP itself is designed around a layered model for networking.

In fact, breaking down networking into an interconnected series of layers defines a general abstract reference model that explains what networks do and how they work. The same kind of model is expressed in somewhat different terms as part of TCP/IP's very definition. The key insight that makes divide and conquer such a powerful tool for implementing networks has been part of TCP/IP's design from its earliest days. This abstraction into layers explains why TCP/IP is so good at allowing different computers, operating systems, and even types of networking hardware to communicate with one another.