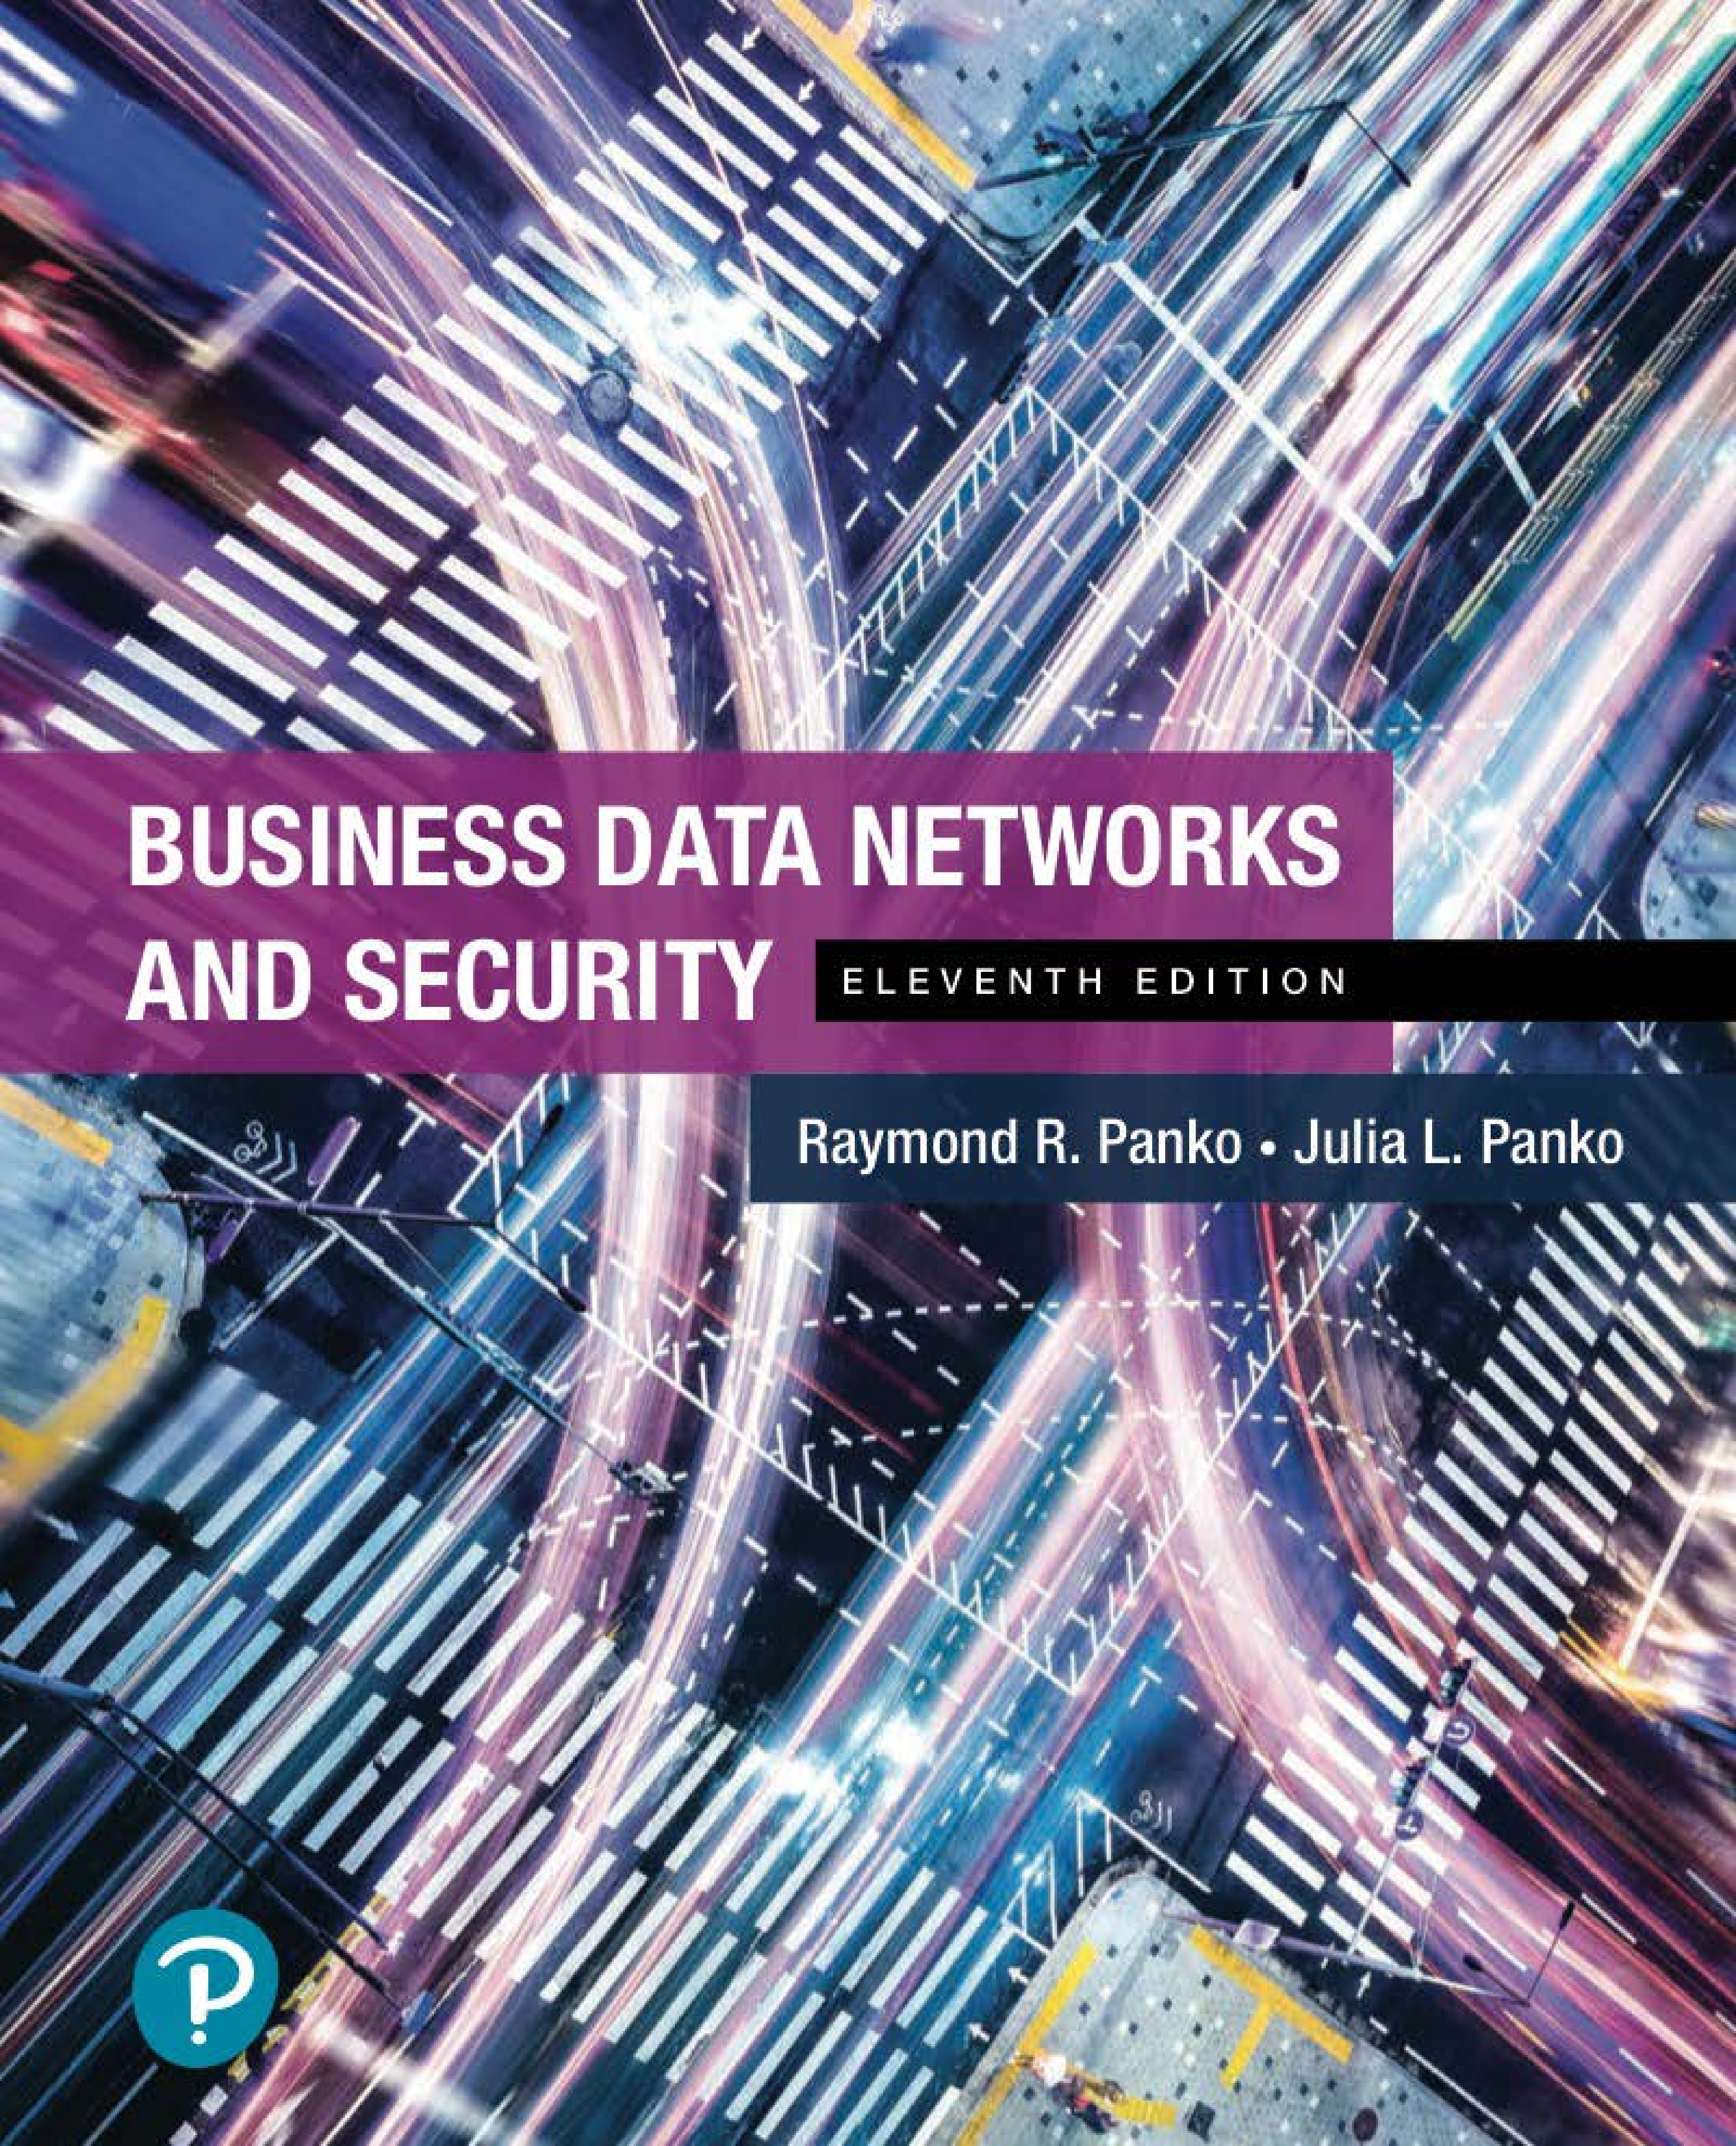

BUSINESS DATA NETWORKS AND SECURITY

ELEVENTH EDITION

Raymond R. Panko • Julia L. Panko

ELEVENTH EDITION

BUSINESS DATA NETWORKS AND SECURITY

Raymond R. Panko
University of Hawai`i at Mānoa

Julia L. Panko
Weber State University

Pearson

330 Hudson Street, NY NY 10013

Vice President, IT & Careers: Andrew Gilfillan
Senior Portfolio Manager: Samantha Lewis
Managing Producer: Laura Burgess
Associate Content Producer: Stephany Harrington
Portfolio Management Assistant: Madeline Houpt
Director of Product Marketing: Brad Parkins
Product Marketing Manager: Heather Taylor
Product Marketing Assistant: Jesika Bethea
Field Marketing Manager: Molly Schmidt
Field Marketing Assistant: Kelli Fisher
Cover Image: uschools/E+/Getty Images
Vice President, Product Model Management: Jason Fournier
Senior Product Model Manager: Eric Hakanson
Lead, Production and Digital Studio: Heather Darby
Digital Studio Course Producer: Jaimie Noy
Program Monitor: Christopher Rualizo, SPi Global
Project Manager: Neha Bhargava, Cenveo® Publisher Services
Composition: Cenveo Publisher Services
Printer/Binder: LSC Communications
Cover Printer: Phoenix Color
Text Font: Palatino LT Pro

Credits and acknowledgments borrowed from other sources and reproduced, with permission, in this textbook appear on the appropriate page within text or at the end of book.

Microsoft and/or its respective suppliers make no representations about the suitability of the information contained in the documents and related graphics published as part of the services for any purpose. All such documents and related graphics are provided "as is" without warranty of any kind. Microsoft and/or its respective suppliers hereby disclaim all warranties and conditions with regard to this information, including all warranties and conditions of merchantability, whether express, implied or statutory, fitness for a particular purpose, title and non-infringement. In no event shall Microsoft and/or its respective suppliers be liable for any special, indirect or consequential damages or any damages whatsoever resulting from loss of use, data or profits, whether in an action of contract, negligence or other tortious action, arising out of or in connection with the use or performance of information available from the services.

The documents and related graphics contained herein could include technical inaccuracies or typographical errors. Changes are periodically added to the information herein. Microsoft and/or its respective suppliers may make improvements and/or changes in the product(s) and/or the program(s) described herein at any time. Partial screen shots may be viewed in full within the software version specified.

Trademarks

Microsoft® Windows®, and Microsoft Office® are registered trademarks of the Microsoft Corporation in the U.S.A. and other countries. This book is not sponsored or endorsed by or affiliated with the Microsoft Corporation.

Library of Congress Cataloging-in-Publication Data

Names: Panko, Raymond R., author. | Panko, Julia L., author.
Title: Business data networks and security / Raymond R. Panko, University of
Hawai'i at Mānoa, Julia L. Panko, Weber State University.
Description: Tenth edition. | Boston : Pearson, [2018] | Includes
bibliographical references and index.
Identifiers: LCCN 2017048586 | ISBN 9780134817125 (alk. paper) | ISBN
0134817125 (alk. paper)
Subjects: LCSH: Business enterprises–Computer networks–Security measures. |
Computer networks–Management. | Computer networks–Security measures. |
Computer security.
Classification: LCC HD30.37 .P36 2018 | DDC 658.4/78–dc23 LC record available at
https://lccn.loc.gov/2017048586

ISBN 10: 0134817125
ISBN 13: 9780134817125

Pearson

To Sal Aurigemma. A great partner in crime in research and teaching.

This page intentionally left blank

BRIEF CONTENTS

Online Modules

This page intentionally left blank

CONTENTS

Online Modules

This page intentionally left blank

PREFACE FOR ADOPTERS

SIX QUESTIONS

This preface begins with six questions that adopters have when considering a textbook.

- What courses is this book used in?
- Why all the security?
- Does this book have the content your students need on the job market?
- Why does it have four principles chapters followed by chapters on specific technologies?
- Does this book have the support you need?
- Does this book have the support your students need?

What Courses use this Book?

- Introductory networking courses in information systems that prepare graduates to work in corporate IT departments use this book. It has the kind of knowledge they need to manage networking in corporations.
- It is used at both the undergraduate and graduate levels.
- Due to its extensive security content, some schools use it in a combined networking and security course. This requires covering the Appendix. Compared to the last edition, the Appendix considerably expands security content. Ideally, schools will have separate introductory network and security courses. Unfortunately, not all schools have that luxury.
- It *does not* focus on the very different needs of computer science students, who will build routers and switches in companies such as Cisco Systems. Instead, it focuses on how to manage and secure them, which is what networking professionals actually do in corporate IT departments. This still requires a lot of technical knowledge but not at the expense of job-required content.

Why all the Security?

In the last two decades, the need for network security knowledge has grown enormously in networking departments. It must be covered pervasively in networking courses. General security courses do not cover network-specific security, such as protecting access points with 802.11i security and knowing ways in which 802.11i security is bypassed in the real world.

Too many IS programs have had to choose between offering an introductory security course and an introductory networking course. This book lets the networking course serve as a decent introduction to security.

Does this Book have the Content your Students need?

This book is based on discussions with networking professionals and focuses on their current and emerging needs. We are especially concerned with potentially disruptive trends such as software-defined networking and high-density Wi-Fi networks. Here is a sampling of this type of job-ready content.

- The Internet of Things. The IoT will keep networking professionals very busy. Obviously, connecting lots and lots of small devices that talk to each other is going to require a lot of work. More broadly, IoT transmission standards and security are pretty raw, requiring even more effort to manage them. Chapter 7 deals with the standards and technologies competing for dominance (or at least survival) in the new market for the Internet of Things.
- Network management. Networking, like security, is more about management than it is about technology. Chapter 3 focuses on network management principles that must be applied in all networking projects. It also focuses on the pervasive importance of SNMP and the potentially disruptive impacts of SDN.
- Security threats and protections. Sun Tzu, in *The Art of War*, exhorted military leaders to know their enemies and to also know themselves. Chapter 4 covers the threat environment facing firms today and the countermeasures that companies can put into place to protect themselves. However, security begins with the first paragraph of the first chapter and continues throughout the book.
- Ethernet is covered in Chapter 5 with a holistic approach. The chapter covers the explosion in Ethernet standards, including those driven by Wi-Fi trends.
- Chapter 6 and much of Chapter 7 deal with Wi-Fi. They again cover technology, which is multifaceted and complex, and they cover wireless management and security. They deal with the current explosion in emerging standards, such as the potentially disruptive 802.11ax standard. Importantly, they show how 802.11i security can be broken.
- Chapter 7 also covers Internet of Things transmission protocols. IoT transmission turns many networking ideas on their heads, such as the desirability of high speed and long transmission distance.
- Chapters 8 and 9 deal with the Internet in context. A special focus is IPv6, which has now gone well beyond its infancy in both technology and use. This material is considerably updated from the previous edition. The material on IPsec is considerably stronger.
- Chapter 10 deals with networking beyond the customer premises. It focuses first on access technologies, then on WAN technologies that must be used beyond the Internet with its limited QoS abilities. The WAN technologies section focuses on leased lines, carrier Ethernet, and MPLS.
- Chapter 11 deals with networked applications—applications that need networks to operate. It focuses on management and security. In the past, some schools skipped this chapter because the material was covered in introductory courses. Actually, intro courses did not focus on the needs of networking professionals, and that is even more true today. This chapter brings the student into the worlds of cloud computing, HTTP/HTML, email, VoIP, and peer-to-peer applications, and it does so in terms of the knowledge that IT professionals need.

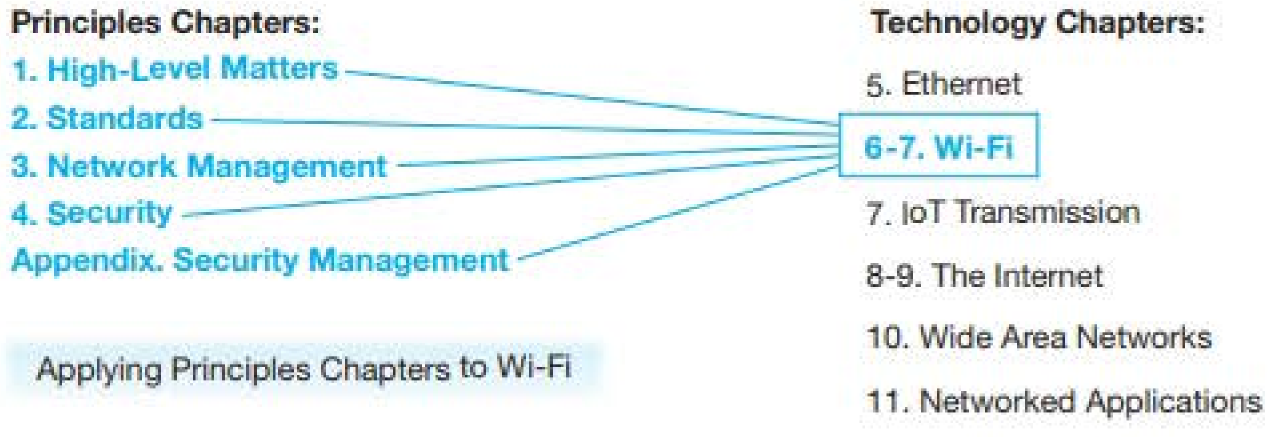

FIGURE P-1 Principles and Applications

Why have four Principles Chapters followed by Chapters on Specific Technologies?

Networking professionals want students to be able to *apply principles* to *real networking situations*. The book begins with four chapters that cover core network principles. It then applies these principles in a series of chapters that deal with Ethernet, Wi-Fi, Internet of Things transmission, the Internet, wide area networks, and networked applications. Figure P-1 illustrates this logical flow for Wi-Fi in Chapters 6 and 7. These chapters deal with how 802.11 Wi-Fi is used in business, how Wi-Fi operates at the physical and data link layers, Wi-Fi security threats and countermeasures, and key points in network management. This approach not only has students deal with technologies holistically. It also reinforces difficult core concepts such as layering.

Traditionally, networking books go "up through the layers." At the end of the course, students have all the knowledge of concepts and principles they need. However, they have limited experience in applying them, which is the whole point of the networking job.

Does this Book have the Support you need?

Teaching is hard. Teaching networking is harder. This book tries to make it a little easier.

PowerPoint Presentations and the Centrality of Figures The PowerPoint presentations are full lectures, not "a few significant figures." A core design principle of this book is that all key concepts are expressed in figures. Most of these figures are Illustrations. Some are "Study Figures," which essentially take notes for the student in areas that do not lend themselves to illustrations.

A core design principle of this book is that all key concepts are expressed in figures.

In line with this focus, the PowerPoint presentations are created directly from the figures. Figures are designed for this. Font size is larger in the PowerPoint slides, and several slide builds are often used to cover a figure well, but making them consistent with the figures has proven to be a great help for both teachers and students.

Adopters get an annotated version of each PowerPoint presentation. This can help you present the material in the slide. Sometimes we even add a little extra information for you to present.

The Instructor's Manual: The Usual Suspects with a Twist Of course, there is an Instructor's Manual with chapter teaching hints and answer keys for chapter questions. There is also a multiple-choice test item file and a test generator for exams.

Test Your Understanding Questions Now for the twist. Each chapter is broken into fairly small and highly targeted sections that end in a handful of Test Your Understanding questions. The Test Item File questions are linked to specific Test Your Understanding Questions. This means that you can assign certain questions for study and exclude others from exams. This lets you tailor exams to exactly the content points you wish your students to be responsible for.

Chapter-Opening Caselets Most chapters begin with brief caselets that students find interesting. In Chapter 1, for example, the caselet deals with how KrebsOnSecurity.com was hit with a denial-of-service attack that used small Internet of Things devices. Try assigning them for reading before the class and go over them as an interaction starter.

Does this Book have the Support your Students Need?

Let's face it. Networking and security are tough. They are highly conceptual. It is not primarily a matter of building cumulative skills as in programming courses. There are a lot of concepts, and they are often abstract or require the student to understand multiple steps. Networking professionals know that their careers are governed by the few things they need to know but don't in particular situations. Students must understand a lot just to be minimally competent.

Guided Reading One way the book helps students is by guided reading. There usually is a chapter-opening caselet to get the juices flowing. The flow that follows is broken up into fairly small pieces, with many headings. This helps the student focus on specific points. Figures show them how they fit together in a broader framework. Important concepts are displayed as key words. The index and glossary are linked to these key words. In addition, critically important concepts are often shown as callouts:

Students quickly learn to pay special attention to these callouts.

Fun Footnotes? Then there are fun footnotes. No, that is not an oxymoron. We limit chapter content to what all students should be able to master in an introduction to networking course. Sometimes, it is useful for some students if a bit more information is available to satisfy their curiosity. We put them in footnotes. They are not required reading, so they are not deadly detailed. Sometimes, footnotes are used for illustrative (semisnarky) comments.

Test Your Understanding Test Your Understanding questions help students stop after a section and see if they understood it. The best students learn that this is the best way to learn because networking is so cumulative, and moving on too fast is a capital mistake. At the end of the chapter are integrative questions that provide exercises for putting the things the student has learned together.

Exam Study These and other design elements help students prepare for exams as well.

- It is good for students to begin their exam prep by skimming for callouts and key words and being sure that they know them.
- Importantly, they should look at all of the figures and see if they can explain them. Again, figures include nearly all major content in the book.
- With this grounding, they should go over the test our understanding questions to see if they understand the detail. If they aren't sure, the text is right there to reread.

CONTENT FLOW

This section describes the flow of content in the book. It discusses each chapter briefly, giving its role in the book. It also describes changes from the previous edition. Overall, this edition is a 70% rewrite.

Chapter 1

This is the first of four "principles chapters" that give the student broad grounding in core concepts and principles needed to understand and deal with specific networking technologies such as Ethernet, Wi-Fi, and the Internet of Things.

Chapter 1 covers basic Internet terminology, concepts, and architectural principles. It begins with a broad introduction to the Internet. It then looks at the Internet from the outside by focusing on what hosts do to send and receive packets. It then looks inside the Internet to show how packets are delivered. On the Internet, routers are connected by data links, which may be single networks. The chapter ends with the distinction between Internet routers, personal access routers, and wireless access points. Students tend to confuse these terms. The Internet of Things is a major theme of the chapter.

Caselet The chapter begins with a caselet to show how KrebsOnSecurity.com was the victim of a distributed denial-of-service attack that used IoT devices.

Objectives After mastering the chapter, the student should be able to . . .

- Discuss how the Internet is changing and the security challenges these changes are creating.
- Explain basic concepts and terminology for the hosts (devices) that connect to the Internet.
- Explain basic concepts and terminology for the Internet itself.
- Explain basic concepts and terminology for single networks and their role on the Internet.
- Explain the distinctions between Internet routers and personal access routers; explain the differences between personal access routers and wireless access points.

Changes In the previous edition, Chapter 1 began with single networks and then showed how they are connected to the Internet. Students said that they wanted to know about the Internet first, so that is how we wrote it. Speed details were moved to Chapter 3 to give a cleaner flow. Students already know speed basics enough to put off details. Cloud computing was moved to Chapter 11 because it primarily deals with application architecture, which deals with the locus of processing. Application architecture is a major theme of that chapter. Standards architectures were moved to Chapter 2, although the first chapter introduces terminology that readies students for standards architectures.

Chapter 2

Chapter 2 presents standards principles and patterns that the student will see throughout networking. This chapter also introduces the main syntax elements of IP, TCP, UDP, and Ethernet.

The chapter, like the rest of the book, is based on the hybrid standards architecture that companies use in real life. They use OSI standards at the physical and data link layer. They primarily use TCP standards at the internet and transport layers. They use standards from a variety of sources for applications. TCP/IP have no problem working with OSI standards at lower layers, and nearly all applications can interface with TCP or UDP. Focusing only on OSI standards makes no sense in terms of corporate realities.

Caselet The chapter opens with a caselet on how Internet standards came to be and why they are sometimes weird.

Objectives After mastering the chapter, the student should be able to . . .

- Explain how Internet standards are made and why this approach is valuable.
- Provide the definitions of network standards and protocols and articulate their importance.
- Explain the OSI, TCP/IP, and Hybrid TCP/IP-OSI architectures and their standards agencies.
- Explain the purpose of each standards layer in the hybrid TCP/IP-OSI architecture, what is standardized at each layer, and which standards agency dominates standards at each layer.
- Explain message ordering in general and in HTTP and TCP.
- Explain message syntax in general and in IP packets, TCP segments, and UDP datagrams, and Ethernet frames.
- Demonstrate how application programs encode alphanumeric, decimal, and alternative data into bits (1s and 0s) before passing their messages to the transport layer.

Changes Compared to the previous edition, standards architectures have been moved entirely from Chapter 1 to this chapter. The syntax of HTTP has been moved entirely to Chapter 11, the Networked Applications chapter.

A specific significant change is that the chapter discusses the Ethernet II frame, not the 802.3 MAC Layer frame. The Internet Protocol standards call for IP packets to be carried inside Ethernet II frames, and this practice appears to be general. Now that IP dominates at the Internet layer, it is Ethernet II frames that students have to understand. Conveniently, the Ethernet II frame is simpler.

Chapter 3

Chapter 3 covers core concepts and principles in network management. It introduces students to the importance of centralized management and to software-defined networking (SDN), which is potentially a fundamentally disruptive technology for changing how we manage networks.

Objectives After mastering the chapter, the student should be able to . . .

- Discuss network quality of service (QoS) and specify service level agreement (SLA) guarantees.
- Design a network layout based on required traffic volumes between pairs of sites.
- Describe options for dealing with momentary traffic peaks.
- Describe the benefits and importance of centralized network management; discuss and compare three tools for centralizing network management: Ping, traceroute, and the Simple Network Management Protocol (SNMP).
- Describe Software-Defined Networking (SDN), including why it is potentially revolutionary.

Changes In the previous edition, Chapter 4 covered both network and security management. That was too much to cover well. Chapter 3 in this edition has the network management information. It also centralizes SDN information, which was spread across multiple chapters in the previous edition. The section on network design has additional examples and exercises and introduces a new tabular approach. Redundancy is shown, but no computations are made because that is for an advanced course.

Chapter 4

Chapter 4 is primarily Chapter 3 in the previous edition. It introduces security threats and countermeasures. It may seem odd to put off security to the end of the principles chapters, but the material in the chapter requires full knowledge of core networking principles and concepts.

Caselet This chapter's caselet is the Target Breach, which was a complex hack. It takes several years before the details of such hacks are understood.

Objectives After mastering the chapter, the student should be able to . . .

- Describe the threat environment, including categories of attacks and attackers.
- Explain how to protect dialogues by cryptography, including encryption for confidentiality, electronic signatures, and host-to-host virtual private networks (VPNs).
- Evaluate alternative authentication mechanisms, including passwords, smart cards, biometrics, digital certificate authentication, and two-factor authentication.
- Describe firewall protection, including stateful packet inspection, next-generation firewalls, and related intrusion prevention systems.
- Describe the role of antivirus protection.

Changes Everything has been updated. The stateful inspection and next-generation firewall sections have been considerably redone.

Appendix

Most teachers who cover the Appendix cover it after Chapter 4, although some will wait until the end because it is a fun read. It includes much of the material from Chapter 4 on security management. It goes into more depth on planning principles and adds a discussion of the response phase. Covering the Appendix after Chapter 4 allows teachers to talk about defense in depth, weakest link thinking, and other principles throughout the discussion of security for specific technologies.

Caselet This caselet builds on the Target Breach discussed at the beginning of Chapter 4. It describes how critical security policies were violated, making the breach possible.

Objectives After mastering the chapter, the student should be able to . . .

- Describe the threat environment, including types of attacks and types of attackers.
- Explain how to protect dialogues by cryptography, including encryption for confidentiality, electronic signatures, and host-to-host virtual private networks (VPNs).
- Evaluate alternative authentication mechanisms, including passwords, smart cards, biometrics, digital certificate authentication, and two-factor authentication.
- Describe firewall protection, including stateful packet inspection, next-generation firewalls, and related intrusion prevention systems.
- Describe the role of antivirus protection.

Chapter 5

Now that the student has mastered basic principles and concepts regarding the Internet, standards, network management, and security, they are ready to apply this knowledge to key network technologies. In Chapter 5, this is Ethernet. Ethernet is covered before Wi-Fi because it is impossible to talk about Wi-Fi management without understanding Ethernet.

Objectives After mastering the chapter, the student should be able to . . .

- Explain basic Ethernet terminology and how Ethernet is standardized.
- Describe basic physical propagation concepts: digital and binary signaling, full-duplex transmission, and parallel transmission.
- Explain the technologies of 4-pair UTP and optical fiber. Compare their relative strengths and weaknesses, including cost and transmission distances.
- Design an Ethernet network based on knowledge of transmission requirements and Ethernet physical link standards, including link aggregation.
- Describe the Ethernet II frame. Explain basic Ethernet data link layer switch operation.
- Describe security threats to Ethernet and ways to deal with them.

Changes Compared to the last edition, this chapter relegates some nice to know but advanced features to footnotes. Power Over Ethernet is one of them. There is just too much stuff to learn about Ethernet to cover everything in an introductory course. The discussion of UTP and fiber media has also been streamlined, and single-mode fiber is moved to a box for additional information. As in Chapter 2, the focus is on Ethernet II frames.

Chapter 6

This chapter and most of the next deal with 802.11 Wi-Fi. This chapter focuses on what students need to know about the core technologies of Wi-Fi. The box at the end deals with the ongoing explosion of new physical layer standards and their relative strengths and issues.

Objectives After mastering the chapter, the student should be able to . . .

- Explain basic Wi-Fi 802.11 terminology and the role of access points.
- Explain basic radio signal propagation concepts, including frequencies, antennas, and wireless propagation problems. These are physical layer concepts.
- Explain the frequency spectrum, service bands, channels, bandwidth, licensed versus unlicensed service bands, and spread spectrum transmission used in 802.11 Wi-Fi LANs. These are also physical layer concepts.
- Describe 802.11 Wi-Fi WLAN operation with access points and a switched Ethernet distribution system to link the access points. Distinguish between BSSs, ESSs, and SSIDs. Discuss communication between access points. These are data link layer concepts.
- If you read the box "Media Access Control (MAC)," compare CSMA/CA+ACK and RTS/CTS for media access control. These are data link layer concepts.
- Compare and contrast the 802.11n and 802.11ac transmission standards. Discuss emerging trends in 802.11 operation, including channels with much wider bandwidth, MIMO, beamforming, and multiuser MIMO. These are physical layer concepts.
- If you read the box "802.11/Wi-Fi Notes," be able to know what happens when devices follow different Wi-Fi standards, explain how devices that follow new Wi-Fi standards get released in profile waves, and describe emerging 802.11 standards and what they will bring.

Changes Compared to the previous edition, a number of topics have been streamlined. The material in the closing box is new. It deals more specifically with the current standards explosion and how products implement standards in profile waves.

Chapter 7

This chapter deals heavily with Wi-Fi security. A key point is that 802.11i security is mandatory but can be defeated by evil twin and rogue access point attacks. Centralized wireless LAN (WLAN) management is critical because access points are so widely dispersed. There is a boxed section on decibel calculations. You can decide how much, if anything, to cover. The chapter ends with a section on the wireless technologies that underpin Internet of Things transmission, including Bluetooth Low Energy, ZigBee, Wi-Fi Direct, and near-field communication (including radio frequency IDs).

Caselet How easy is it to crack an unprotected Wi-Fi hot spot? This caselet shows how seven-year-old Betsy Davies did it in just under 11 minutes. Including reading a tutorial on how to do it.While drinking a milkshake.

Objectives After mastering the chapter, the student should be able to . . .

- Explain 802.11i Wi-Fi security.
- Explain why 802.11i security is not enough for WLANs.
- Discuss 802.11 WLAN management.
- Work with decibel representations of power ratios (if he or she reads the box on decibels).
- Compare peer-to-peer local wireless technologies that will be important for the Internet of Things.

Changes The challenging evil twin section has been broken into more pieces and simplified to the extent possible. The decibel section has been heavily rewritten. The section on IoT transmission technologies is expanded considerably to reflect today's explosion in IoT transmission standards and technology.

Chapter 8

Chapters 5 through 7 dealt with single-network technologies that use standards at the physical and data link layers. With this as a basis, we can now move into TCP/IP at the Internet and transport layers. This chapter looks at how routers make their routing decisions and looks at the syntax of IPv6 main headers, extension headers, and higher-layer content. IPv6 is an important topic in networking today because IPv6 is no longer just a percent or less of all IP traffic. Students need to know how to write IPv6 addresses for human reading.

Some have asked why the book waits so long to move into TCP/IP. The answer is that TCP/IP is substantially more complex than Ethernet and Wi-Fi technology. Learning simpler technologies first makes it easier to learn TCP/IP and its many standards.

Objectives After mastering the chapter, the student should be able to . . .

- Define hierarchical IPv4 addresses, networks and subnets, border and internal routers, and masks.
- Given an arriving packet's destination IPv4 address, explain what the router will do with the packet based on its routing table.
- Explain the IPv4 packet header fields we did not see in earlier chapters.
- Explain the IPv6 packet's main header fields and IPv6's use of extension headers.
- Convert a 128-bit IPv6 address into canonical text notation consistent with RFC 5952.

- Explain TCP segment fields, UDP datagram fields, and TCP session closings.
- Explain why application message fragmentation is not possible with UDP.

Changes Relatively little is new to this edition, although almost all topics have been rewritten to help student comprehension.

Chapter 9

This chapter takes the TCP/IP discussion into management and security. TCP/IP uses many supervisory protocols beyond TCP, UDP, and IP. This chapter discusses a few of them.

Objectives After mastering the chapter, the student should be able to . . .

- Explain IPv4 subnet planning and do the calculations needed for working with subnet and host parts and deciding on part lengths.
- Do the same for IPv6.
- Explain the purposes of Network Address Translation (NAT) and how NAT operates.
- Explain in more detail than you learned in Chapter 1 about how the Domain Name System (DNS) and the Dynamic Host Configuration Protocol (DHCP) operate.
- Describe the object model in the Simple Network Management Protocol (SNMP) and describe the enabling value of good security in the use of Set commands.
- Describe how the DNS was modified to deal with IPv6 addresses for host names.
- Describe how dynamic routing protocols work and how to select among alternative dynamic routing protocols.
- Describe the Internet Control Message Protocol (ICMP).
- Explain central concepts in IPsec (IP security), including its strategic importance, transport versus tunnel mode operation, ESP versus AH protection, security associations, important cryptographic methods and options, session initiation with IKE, and how IPsec compares to SSL.

Changes Again, relatively few things were changed, but there was a good deal of rewriting and streamlining. One specific change is that subnetting for IPv6 now follows immediately after subnetting for IPv4. Another is that the section on IPsec has been expanded to include such things as how session initiation is done. More IPv6 material is an obvious need.

Chapter 10

This chapter deals specifically with wide area networking. In WANs, companies must deal with carriers instead of doing things themselves. They also face much higher costs per bit transmitted, so efficiency is critical. Isn't wide area networking just the Internet? No, it isn't. Companies must have quality of service guarantees for some of their site-to-site traffic, and the Internet does not provide that. Carrier WAN services for corporations today are dominated by lease lines, carrier Ethernet, and MPLS.

Most carriers have moved all of their Frame Relay and other customers to carrier Ethernet or MPLS. The chapter looks at cellular data communication, ADSL, and cable modem services as well as the carriers' local loop, which serves the premises of home and business users.

Changes After mastering the chapter, the student should be able to . . .

- Contrast LANs and WANs in terms of technology, diversity, economics, speed, and need for optimization.
- Describe the three carrier WAN components and the two typical business uses for carrier WANs.
- Describe how the telephone system is organized, including its hierarchy of switches. (Most carrier WAN networks use the public switched telephone network for some or all of their communication.)
- Explain and compare the ADSL and cable modem residential Internet access services and how fiber to the home is changing the residential access market.
- Discuss trends in cellular data transmission speeds.
- Distinguish between access lines and leased lines. Select a leased line for a given application speed requirement. Explain how companies use leased lines in Internet access.
- Explain how networks of leased lines, carrier Ethernet, and MPLS can be used for site-to-site communication within a firm. Discuss the relative advantages and disadvantages of each.
- Explain the capabilities of WAN optimization devices.

Changes This chapter mostly covers the same topics that Chapter 10 did in the previous edition. However, rewriting and streamlining is very heavy. There is more clarity on why the Internet does not meet the quality of service levels needed in most firms, requiring them to used technologies beyond the Internet for their much of their long-distance communication

Chapter 11

This chapter is about application architectures—where application processing is done and why it is done there. Falling prices for both computers and transmission have taken us from stand-alone mainframes to mainframes with dumb terminals, to client/server processing, cloud computing, and peer-to-peer computing. The chapter begins by noting that most computer hacks today involve taking over an application and receiving its permissions. The chapter looks at cloud computing and P2P computing. In between, it looks at the behavior of today's most central networked applications.

Objectives After mastering the chapter, the student should be able to . . .

- Explain core concepts in networked applications and application architectures.
- Describe how taking over an application can give an attacker the ability to control the computer.

- Describe how Netflix uses cloud computing and how this illustrates the importance of host technology (and cloud computing specifically) as a driving force for networking.
- Describe the World Wide Web in terms of standards and explain how a webpage with text, graphics, and other elements is downloaded.
- Describe electronic mail standards and security.
- Describe voice over IP (VoIP) operation and standards.
- Explain why peer-to-peer (P2P) computing is both desirable and dangerous.

Changes This chapter brings cloud computing from Chapter 1. The treatment of the "Big Three" business applications—the WWW, e-mail, and VoIP—is somewhat expanded. Peer-to-peer computing is reduced. It focuses on traditional VoIP versus P2P VoIP to show what peer-to-peer computing changes. It also discusses Tor, which is a P2P tool for anonymizing IP transmission. Tor is used both by people seeking anonymity and by cybercriminals.

The "a" Chapters

Several chapters are followed by an "a" chapter (1a, 3a, etc.) that provides some hands-on experience for students.

Chapter 1a. Hands-On: A Few Internet Tools This "a chapter" gives the student a bit of basic hands-on experience to help them make the concepts in Chapter 1 more concrete while learning a few useful tools. After mastering the chapter, the student should be able to . . .

- Test his or her Internet connection speed.
- Look up a host's IP address by querying a DNS server.
- Use Ping and traceroute to diagnose an Internet connection.

Chapter 3a. Hands-On: Microsoft Office Visio As the name suggests, this is a quick tutorial on Visio basics. Visio is widely used in network representation. Some schools have free versions for students. For those that do, Visio is useful in doing some homework questions.

Chapter 5a. Hands-On: Cutting and Connectorizing UTP If students are still cutting and connectorizing wire on a regular basis three or four years into their careers, they have probably made a wrong turn somewhere. However, learning how to do it is a good skill, and it makes 4-pair UTP less abstract. It is also fun, and it gives students something to take home to show their parents. After mastering the chapter, the student should be able to . . .

- Cut, connectorize, and test 4-pair UTP cabling.
- Explain the difference between solid wire and stranded-wire UTP.
- Know when to use patch cables.

Chapter 8a. Hands-On: Wireshark Packet Capture This chapter has the student capture a stream of IP packets and then analyze their headers in some detail. This exercise makes the syntax of IP, TCP, and UDP far more real to the student.

Chapter 9a. Hands-On: Cisco's IOS Command Line Interface (CLI) This chapter addition introduces the student to the flavor of Cisco's command line interface used in switches, routers, and other devices. It walks the student through a few sample interactions. After the class, some students may wish to master IOS in detail to help them pass valued Cisco certifications. This chapter was not in the previous edition.

Online Modules

Teachers who want to cover material not in the text may find it useful to look at online modules that cover additional matters. These are available for both teacher and student download. The purpose is to allow you to cover certain additional topics without having to do more preparation. A word of caution. There is a lot of material. Only small amounts of the material in the online modules are likely to fit into courses.

Module A: More on TCP This module is for teachers who wish to cover TCP sequence and acknowledge numbering and flow control using the Windows Size field. It comes most naturally after Chapter 8.

Module B: More on Modulation The main text does not deal with modulation. Covering this short module will help your students understand how the most advanced 802.11 physical layer standards can transmit data more efficiently by sending more bits per clock cycle.

Module C: More on Telecommunications Some courses have titles that include Telecommunications. This normally means telephony. This chapter has material for these courses.

Module D: Directory Servers Directory servers are a big thing in the corporate world. This module looks at directory servers in more detail, including Microsoft's Active Directory and authentication using directory servers. The latter is covered briefly in the Appendix. This module adds metadirectory servers.

PREFACE FOR STUDENTS

THIS BOOK

Most textbooks start by trying to convince you that the subject matter is important. This one doesn't need to do so. Everybody knows that the Internet is important. Ditto on security.

Networking and Security Why *both* networking and security? The reason is that security pervades professional networking today. There is no way to separate them. Every network project has a sizeable security content. The traditional view that networking is the moving of bits and packets is no longer sufficient. Nor is it enough to slap a security chapter at the end of the book. Security must be deeply integrated into your knowledge of Ethernet, Wi-Fi, TCP/IP, applications, and everything else. Some teachers cover the Appendix to give you an even deeper view of security planning and response when security failures happen.

Principles and Their Application Figure 1 shows how this book will help you learn networking and security. First, you will learn concepts and principles. You will learn core ideas such as how the Internet operates, the nature and idiosyncrasies of network standards, keys to managing networking projects, and core security concepts. Why does security come last among these core chapters? The answer is simply that you can't learn network security without understanding core networking ideas first.

The rest of the book takes you through a series of technologies. For each, you will apply the concepts and principles you mastered in the first four chapters. For instance, when you learn about Wi-Fi in Chapters 6 and 7, you will understand its basic operation, physical transmission, switch operation, standards, management, and, of course, security. You will do the same for the other technologies and applications shown in Figure 1.

Job-Relevant Knowledge We have done everything we could to fill this book with job-relevant knowledge. You will not have to learn about technologies that haven't been seen in this century. There simply isn't time to cover history when companies need

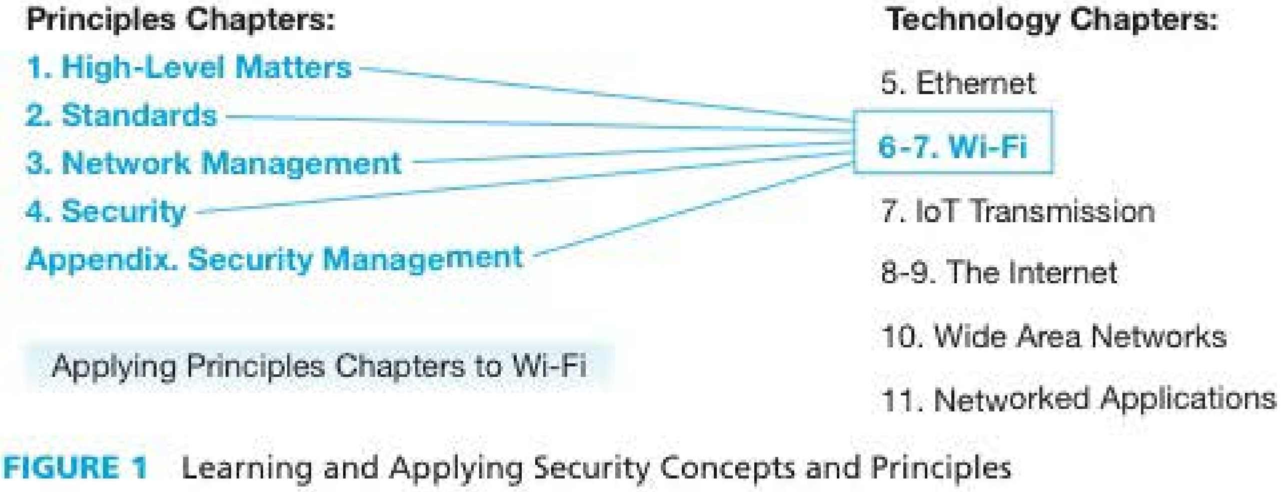

FIGURE 1 Learning and Applying Security Concepts and Principles

students who understand IPv6, IPsec, the current explosion in Ethernet standards, the current explosion in Wi-Fi standards, Internet of Things transmission protocols, and many other recent developments. You will learn all the general principles that all networking books cover, but you will learn about them in the context of today's important technologies. If you can, work through the hands-on "a" chapters that follow several main chapters. These things are kind of fun, and they will make concepts a lot more concrete.

Information Systems versus Computer Science How does an information systems book differ from computer science books? Our friends in computer science teach students how to design routers in networking and how to create ciphers in security. Our students will work in IT departments. They will never build a router, but they will buy them and need to understand how to manage and secure them. Would it help to teach you how to build a router? Perhaps. But that would mean not teaching you how to use them in real organizations because there wouldn't be time. Design your own cipher? We teach our students that doing that is stupid. You do not have to know how to design a cipher to know how to select a cipher to use in a project, and 99.9% of all developed ciphers are broken quickly.

STUDYING NETWORKING

Although networking and security are exciting, many find them hard to learn. It is not that they are terribly difficult inherently. The main problem is that you do not have a mental framework when you start, so it is hard to absorb individual pieces of knowledge. You need to learn frameworks and individual pieces at the same time.

Frameworks and Individual Pieces Unfortunately, this means that you need to jump back and forth between frameworks and individual pieces until both settles into place. Once you master that discipline, you will be able to grasp major constellations of concepts. If you do not, this course is going to be very hard.

Intelligent Choices This class requires upper-level college thinking. In the first years of college, you are learning individual facts. In your final years, you need to master comparisons between concepts so that you know which to apply. This is exactly what networking professionals need to do. To design a network, you need to make complex decisions requiring you to evaluate alternatives. You also need a complex mental model to troubleshoot problems, which takes up a surprising amount of professional work time. It has been said that artists are known for their best moments but engineers are known for their worst. Any piece you do not master comes back to haunt you.

TLAs and FLAs Then there is the problem of TLAs and FLAs (three-letter acronyms and four-letter acronyms). You will see a lot of them. Why not just avoid acronyms? The problem again is the environment in which network professionals work. If you pick up any trade magazine, you will see that few acronyms are ever spelled out. You will have to learn a lot of them. Think of them as abbreviations when you text people on your phone.

There is a comprehensive Glossary at the back of the book. If you aren't sure what a term means, go to it for a quick definition. If that isn't enough, the index will tell you what pages to read. If a page number in the Index is boldfaced, look at that page first.

No Escape By this point, you may have decided that networking and security are rather challenging and that programming and database are beginning to seem attractive. Unfortunately, they won't get you away from networking and security. Today, most programs in industry are written to work with other programs on other machines; and all of their interactions take place over networks. Database management systems and systems analysis also require solid networking knowledge. So learn networking as much as you can. We have cute kittens to watch and alien ships to destroy. For security, we have fascinating stories, and you are not just going up against hardware reliability and software bugs. You will find yourself matched against determined attackers who will respond to whatever you do.

STRUCTURE OF THE MATERIAL

If you page through the book, you will see that it is set up a little differently than other textbooks you have seen.

Fun Footnotes Fun Footnotes? Footnotes are dry and academic. Ours are little bits of knowledge that take you beyond the book. Some students are really turned on by them. No, honestly. In any case, they are never required reading. If you find them interesting, enjoy them.[1] If not, ignore them. Some are different; they take a swipe or two at what standards agencies do.

Small Sections Long blocks of text are daunting to read. This book breaks things into a lot of small digestible sections with a lot of headings.

Short Sections with Level Three Headings (Like This One) If you just read a title, you often can get the gist of what follows. This will make it easier to know what the section does. Learning small chunks of information also increases comprehension.

Key Terms Key concepts and their acronyms are shown in **boldface**. That alerts you to their importance. If you forget this key term, you can always go to the Glossary to refresh your memory. The index also lets you see where a key term appears. If a page number is shown in boldface, that is where the concept is defined or characterized.

Callouts As you read a section, pay attention to callouts like the one below. They emphasize an important fact or idea and often things that are points of frequent confusion. Before exams, first go over the callouts until you have them cold.

[1] This is our way to put in some material that is good to know but that is more than an introductory course should include and that generally has proven difficult for even well-prepared undergraduates to master.

As you read a section, pay attention to callouts like this one. They emphasize an important fact or idea.

Comprehensive Figures Nearly every important concept in the book is covered in a figure. The figures are very carefully designed to show the flow of actions or ideas. As you read a section, look at the figures carefully. See if you can teach each to an imaginary friend. First set the stage. What are the pieces? Then step through the various parts of the figure.

Some figures end with (Study Figure). These are essentially notes on what the section covers. It gives you a view of a block of material from 10,000 feet and helps link frameworks with individual facts.

Test Your Understanding Questions The material in networking is highly cumulative, so you want to master the material in a section before going on. Each section ends with Test Your Understanding questions designed to help you see if you have understood what you just read. When you reach them, you want to go on instead of testing yourself. If can get yourself to go over the questions immediately, it will help you learn whether you understand the material you just read. If you aren't comfortable, go back and learn the material again.[2]

STUDYING FOR EXAMS

If you think you won't have to study for exams, it will probably end in tears. Given this reality, some advice about how to study for exams is in order.

- Again, a good place to begin is the callouts. Go through them and make sure you understand them all. They include a lot of the chapter's important content in little chunks.
- A good place to go next is the figures. Go through them one at a time, teaching them to your imaginary friend. This again packs a lot of material in small packages. Let the study figures help you understand the structure of the relevant section and its key points. To tell a story, first set the stage. What is the problem being solved or presented in the figure? What are the devices and programs involved? Then walk through the rest of the figure. Often, steps to do so are numbered. If you understand all the figures, you should do well.
- After you have done these things, go over your Test Your Understanding answers. If you did them from homework, don't just study your original answers. When you wrote them, your knowledge was less mature than it will be just before exams, and many of your early answers will be science fiction. One helpful trick is to ask yourself why each question is important. Why do you have to know it?

[2] A key idea in answering Test Your Understanding questions is to maximize what you learn, ask yourself, "Why is this question important?" Each question has a reason for being there. See if you can understand what it is and why it is important.

- Yes, you are going to have to reread much of the text. This is especially important for parts of the chapter that deals with complex frameworks with multiple parts. As discussed previously, you will learn them, forget them, learn them again, and so forth.

CERTIFICATIONS

In high school, you may have taken advanced placement exams. Passing AP exams impress college admissions committees. Analogously, IT certification exams let you demonstrate some in-depth knowledge and also tell companies that you are serious and proactive. The problem is that there are many certifications, and they offer different levels of knowledge about different topics. Many require hands-on expertise in working with networking technology. Most require two to five years of work experience for full certification, although some of these allow you to receive associate status if you pass but have not yet acquired the work experience. All of them cost money, in some cases thousands of dollars.

Network+ and Security+ The least ambitious certifications are CompTIA's Network+ and Security+ certifications. Both are quite doable with some extra study. Neither impresses IT departments highly. However, they are achievable with reasonable effort. A major practical problem with these certifications is that they spend far too much time on technologies and concepts that have been irrelevant for thirty years or more.

Vendor Certifications Vendors offer certification exams that are prized by IT departments. The introductory certifications show that the bearer has the knowledge to do entry-level tasks in the exam's area.

The problem with vendor certification is that they see things only from that particular vendor's point of view. For example, Cisco will cover a great deal about Cisco routers, switches, and other network devices. In contrast, Microsoft will focus on networking from the client and server point of view, including various types of network servers such as DNS servers.

Passing a vendor certification will require you to learn more than an introductory network course will cover. You will need to buy a book to study. Many of the concepts will be the ones you learned in this course. You will also see quite a few topics in depth. Sadly, in our opinion, you will also have to master quite a few legacy technologies that have not been seen in this century. We understand that businesses must support some obsolete network technologies, so learning about them in a vendor certification course makes sense. Given that only some students go on to networking makes it silly to cover these topics in introductory networking courses, however. It takes too much time away from job-relevant material.

For new graduates, Cisco now offers the Cisco Certified Entry Network Technician certification. A CCENT certification validates skills for entry-level work. Those who pass have the skills to install and manage a small branch office network in an enterprise. This includes relevant network security. To be attractive to corporations, students should achieve the next-level Cisco certification, Cisco Certified Network Associate (CCNA).

Professional Association Security Certifications Security has professional associations for people working in security. They generally offer certification programs.

- For broad security professionals, $(ISC)^2$ offers certifications in a number of security domains. Passing most or all of them will validate a good level of mastery of security. For new graduates, there is the Associate of $(ISC)^2$ certification, which allows a student with no work experience to demonstrate a good level of knowledge before obtaining the experience requirements for more advanced certifications. In turn, the Systems Security Certified Practitioner (SSCP) certification requires one year of experience in one of eight content domains. The most important initial certification is the Certified Information Systems Security Professional (CISSP). This requires five years in two or more of the eight domains.
- For information systems auditors, there are more focused certifications. These are offered by ISACA, the Information Systems Auditing and Control Association. ISACA offers the Certified Information Systems Auditor (CISA) and Certified Information Systems Manager (CISM) certifications.

Advanced Certification Programs and Master's Degrees At a higher level of knowledge and skills, there are advanced certification programs and master's degrees. The predominant advanced certification program in security is offered by SANS, which offers advanced courses in specific areas leading to a broad level of knowledge. These courses are quite expensive. Most SANS participants are sponsored by their employers. The first author has found them to be great courses.

ABOUT THE AUTHORS

Ray Panko is a professor of IT management and a Shidler Fellow at the University of Hawai'i's Shidler College of Business. His main courses are networking and security. Before coming to the university, he was a project manager at Stanford Research Institute (now SRI International), where he worked for Doug Englebart, the inventor of the mouse and creator of the first operational hypertext system. He received his B.S. in physics and his M.B.A. from Seattle University. He received his doctorate from Stanford University, where his dissertation was conducted under contract to the Office of the President of the United States. He has been awarded the Shidler College of Business's Dennis Ching award as the outstanding teacher among senior faculty. His e-mail is Ray@Panko.com.

Julia Panko is an assistant professor on the faculty at Weber State University. She received her doctorate from the University of California, Santa Barbara. Her research interests include the twentieth- and twenty-first-century novel, the history and theory of information technology, and the digital humanities. Her dissertation focused on the relationship between information culture and modern and contemporary novels.

This page intentionally left blank

Chapter **1**

Core Network Concepts and Terminology

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Discuss how the Internet is changing and the security challenges these changes are creating.
- Explain basic concepts and terminology for the hosts (devices) that connect to the Internet.
- Explain basic concepts and terminology for the Internet itself.
- Explain basic concepts and terminology for single networks and their role on the Internet.
- Explain the distinctions between Internet routers and personal access routers; explain the differences between personal access routers and wireless access points.

A STATE OF SIEGE[1]

On September 15, 2016, criminals launched a massive cyberattack on KrebsOnSecurity.com. This is the blogsite of Brian Krebs, whose posts are often the first analyses of major cybercrime incidents (such as the Target breach we will see in Chapter 4).

[1] Kyle York, "Dyn Statement on 10/22/2016 DDoS Attack," Dyn, April 19, 2017, https://dyn.com/blog/dyn-statement-on-10212016-ddos-attack/; Brian Krebs, "KrebsOnSecurity Hit With Record DDoS," KrebsOnSecurity.com, September 16, 2016, https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/; Brian Krebs, "Source Code for IoT Botnet 'Mirai' Released," KrebsOnSecurity.com, October 16, 2016, https://krebsonsecurity.com/2016/10/source-code-for-iot-botnet-mirai-released/; Brian Krebs, "Who Makes the IoT Things Under Attack?" KrebsOnSecurity.com, October 16, 2016, https://krebsonsecurity.com/2016/10/who-makes-the-iot-things-under-attack/; Brian Krebs, "Hacked Cameras, DVRs Powered Today's Massive Internet Outage," KrebsOnSecurity.com, October 16, 2016, https://krebsonsecurity.com/2016/10/hacked-cameras-dvrs-powered-todays-massive-internet-outage/; Brian Krebs, "Akamai on the Record KrebsOnSecurity Attack," KrebsOnSecurity.com, November 16, 2016, https://krebsonsecurity.com/2016/11/akamai-on-the-record-krebsonsecurity-attack/.

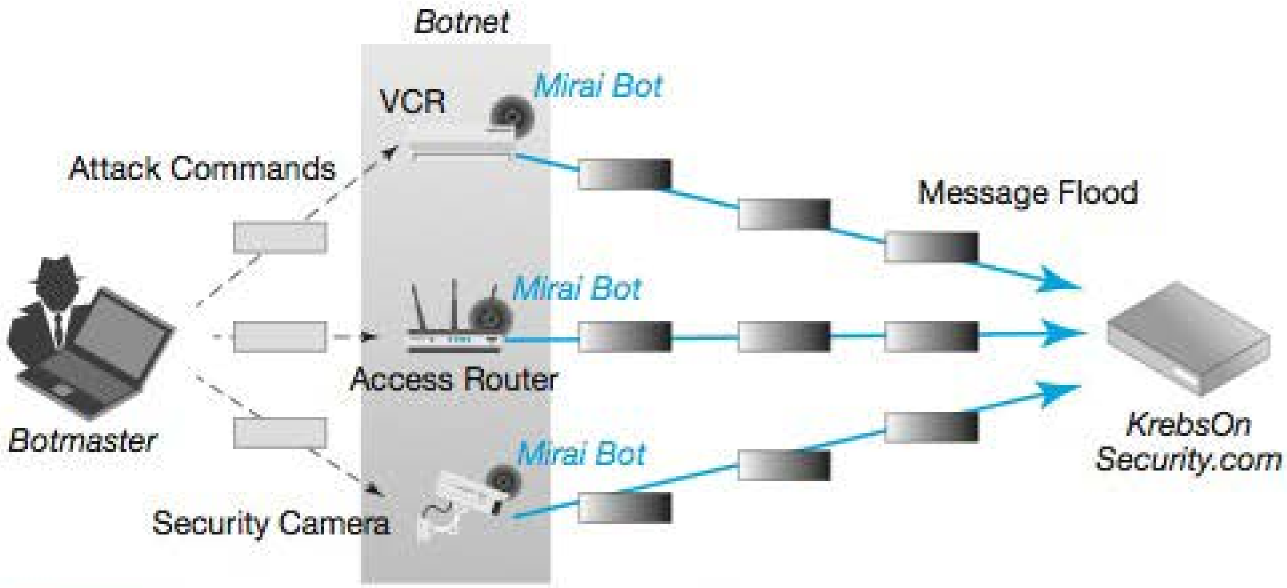

FIGURE 1-1 Simplified Depiction of Mirai Distributed Denial-of-Service Attack

Cybercriminals hate him, and they had attacked his site 269 times in the previous four years.[2] (This was about one attack every five days.) The attacks that began on September 15, 2016, however, were unprecedented.

DDoS Attack These attacks were distributed denial-of-service (DDoS) attacks. Figure 1-1 shows a simplified view of a DDoS attack.[3] In advance, a cybercriminal called a botmaster installs malware on hundreds or thousands of computers without the owners' knowledge. This malware is called a bot. Like a physical robot, a malware bot can be given goals, which it will then execute in detail. In Figure 1-1, the botmaster commands the bots to attack a certain target site. Each bot then sends a flood of packets at the target host. The traffic overwhelms transmission lines to the target. The particular botnet malware that attacked Krebs' site was called Mirai.

Enormous Traffic The September 2016 attack was remarkable for two reasons. The first was the deluge of traffic it threw at Krebs' site. The Mirai bots were able to flood the site with traffic at an astounding 620 Gbps[4] (billions of bits per second).[5] According to Akamai, which was protecting KrebsonSecurity.com at the time, this was almost twice the volume of any DDoS attack it had ever encountered.[6] Mitigating such an attack was daunting, and it took considerable time.

Internet of Things (IoT) Devices The second reason the attack was remarkable was the nature of the devices used in the attack. Normally, DDoS attacks use

[2] Krebs, "Akamai on the Record KrebsOnSecurity Attack."

[3] We will look at these attacks in more depth in Chapter 4.

[4] Speeds are measured in bits per second, kilobits per second (kbps), megabits per second (Mbps), and gigabits per second (Gbps).

[5] Krebs, "Source Code for IoT Botnet 'Mirai' Released."

[6] Krebs, "KrebsOnSecurity Hit With Record DDoS."

compromised desktop computers, laptops, and other traditional IT devices. In the attack on Krebs' site, however, the attacking computers were small nontraditional devices, including home access routers, home security cameras, and home VCRs. In a trend called the Internet of Things (IoT), we are seeing explosive growth in Internet connections by devices previously too lacking in power to use the internet. The size of the IoT is difficult to discuss because it is growing so explosively. However, Gartner, Inc. estimated the number of active IoT devices at 5 billion and forecasts that 2020 will see almost 21 billion.[7] Even if that forecast is highly optimistic, IoT devices are already about as widespread as humanly used computers and will soon be far more numerous.

Weak IoT Security The cybercriminals realized that IoT devices often have weak security. Many come with a login account paired with a well-known default password. If the default password is not changed, anyone can take over the device over the Internet. Users often fail to change them. In fact, some default passwords are hardcoded into IoT devices and cannot be changed by the user.[8] The Mirai malware jumped from one device to another by trying a mere 68 device-password combinations.[9] In many ways, this attack was a coming of age for the Internet of Things. IoT may still be in its infancy overall, but it is now mature as a destructive force.

Dyn There have been many other Mirai victims. On October 21, 2016, Dyn, Inc. was the target of a similar attack. In a postmortem on the attack, Dyn reported that it had been attacked by tens of millions of discrete IP addresses known to be part of the Mirai botnet.[10] Dyn is a Domain Name System (DNS) hosting service. We will see DNS later in this chapter. If you know the name of a site, such as panko.com, you cannot send it messages until you learn its official Internet Protocol (IP) address. (To give an analogy, if you know someone's name, you cannot call that person until you learn his or her telephone number.) A DNS server gives your computer a named site's IP address. If a DNS server that serves hundreds or thousands of popular sites is disabled, the result can be chaos. Among the sites at least temporarily disrupted in the Dyn attack were Amazon, Netflix, Twitter, Spotify, Reddit, and Tumblr.[11] This incident did not merely attack a site. It attacked a critical piece of the Internet infrastructure.

Perspective The great promise of the Internet has been to give access to "anything, anytime, anywhere." Unfortunately, criminals are quick to exploit new technologies. The Internet has evolved with breathtaking speed, bringing both new applications and new types of attacks. Networking people are involved in a protracted arms race with cybercriminals, and the cybercriminals have been winning too often.

[7] Ibid.

[8] Krebs, "Hacked Cameras, DVRs Powered Today's Massive Internet Outage."

[9] Krebs, "Who Makes the IoT Things Under Attack?"

[10] York, "Dyn Statement on 10/22/2016 DDoS Attack."

[11] Krebs, "Hacked Cameras, DVRs Powered Today's Massive Internet Outage."

All this does not mean that the Internet or other networks are bad. The very reason denial-of-service attacks are so damaging is that the Internet's benefits have become indispensable for people and organizations. However, every garden has snakes. Networking cannot be managed without understanding security, and security cannot be managed without understanding networks.

Test Your Understanding

1. a) What is a DDoS attack? b) In what two ways was the KrebsOnSecurity.com DDoS attack unusual? c) What do we mean by the "Internet of Things?" d) What happens when a host cannot reach a Domain Name System server? e) What specific security weakness did the Mirai malware use to propagate from machine to machine?

ANYTHING, ANYTIME, ANYWHERE

The Internet used to be the "*New Thing*." It caught fire in the public's imagination in 1995 when the Internet first became commercial. Before then, the Internet's Acceptable Use Policy explicitly prohibited most commercial activity. This was done because the Internet's transmission backbone was supplied by the National Science Foundation (NSF). Using the NSF to subsidize commercial activity was simply not in the cards. In 1995, however, the NSF pulled out. The rationale for the Acceptable Use Policy vanished. The Internet could be used commercially. It was, immediately.

Test Your Understanding

2. When was commercial activity on the Internet first allowed?

The Internet Reorganizes to Get Commercial

Internet Service Providers In 1995, commercial **Internet service providers (ISPs)** took over the backbone of the Internet. They also became the onramps to the Internet. Anyone wanting to use the Internet must go through an ISP. The Internet today is simply a collection of ISPs that collectively deliver traffic from source to destination computers. Figure 1-2 illustrates this situation.

Internet transmission is handled by commercial Internet service providers (ISPs).

Hosts Figure 1-2 notes that all devices connected to the Internet are called **hosts**. You will encounter this term throughout this book. A laptop is a host when it connects to the Internet. So is a mobile phone. So are the webservers and other servers that provide the services you use when you use the Internet.

Devices that connect to the Internet are called hosts.

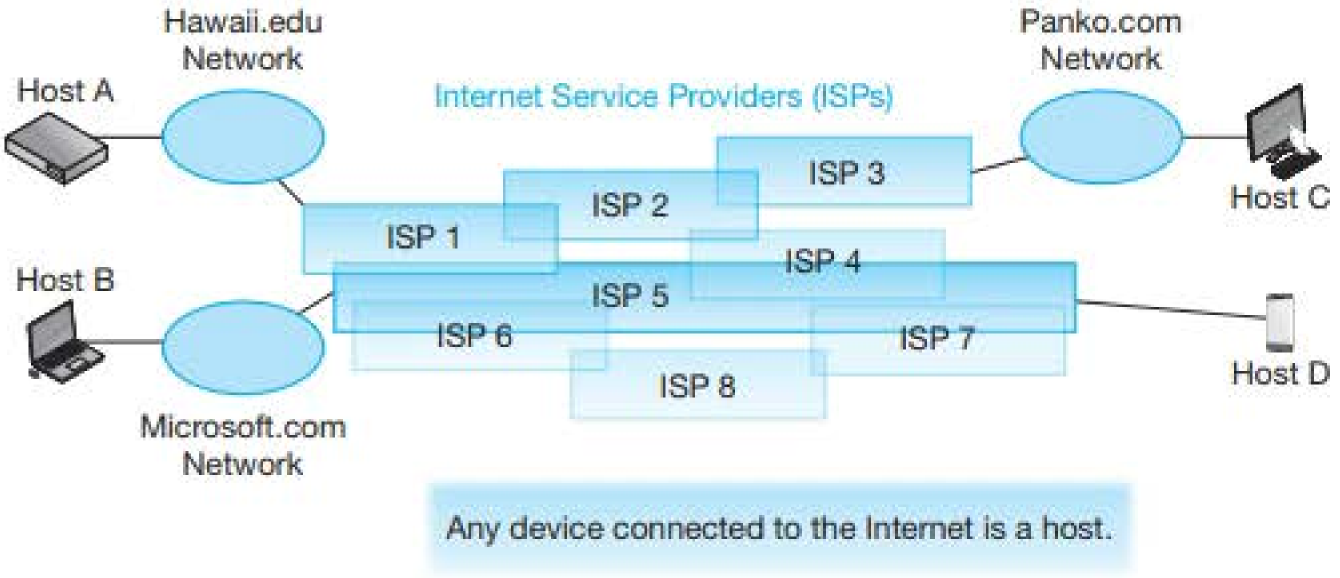

FIGURE 1-2 The Internet: Internet Service Providers, Organizational Networks, and Hosts

E-Commerce The year 1995 saw an immediate rush of commercial companies to ply their businesses over the Internet. Companies with such familiar names as Amazon and eBay were ready and waiting. Amazon's entry was especially interesting. Jeff Bezos wanted to create a company that would sell everything over the Internet, not just books. He chose the name of the company to indicate that it would be a very wide torrent for delivering goods and services. When you look at the Amazon logo, note the arrow at the bottom. It points from A to Z.

Why start with books? Bezos realized that the book industry had almost everything needed for online sales. Publishers and distributors had huge warehouses of books and the ability to do single-item packaging. More importantly, everything was on their computers. Amazon could reach into those databases and provide an online sales front end, complete with the company's innovative one-click ordering. Many organizations and individuals developed simpler non-interactive informational websites to provide information. Soon the Internet became the first place to go for information, some of it correct.

Test Your Understanding

3. a) What services do Internet service providers provide? b) In Figure 1-2, through which ISP(s) will traffic pass if a packet from Hawaii.edu goes to Panko.com? (Answer: ISP 1, ISP 2, and ISP 3) c) Through which ISP(s) will traffic pass if a packet from Microsoft.com goes to the mobile phone in the lower right of Figure 1-2? d) Through which ISP(s) may traffic pass if a packet from Microsoft .com goes to Panko.com? (Hint: There are multiple possible answers.)
4. a) What do we call any device connected to the Internet? b) When you use a laptop to connect to the Internet, is it a host? Explain in terms of the definition of *host*. c) When you use the Internet, are *you* a host? Explain in terms of the definition.

Old Yet Always New

No Longer New? The Internet today, more than a human generation after its creation, is no longer new. Many of the young pioneers who created it are no longer with us. Both e-commerce and informational websites that appeared only about twenty years ago are also old hat.

Commercial for More than Twenty Years

- In 1995, the U.S. government pulled out transmission funding
- Now, e-Commerce was possible

Yet Still New Applications, Even Entire Classes of Applications

- Social Media, etc.

Growing Speed

- High-definition and 4K video, large data transfers, full-computer backup, etc. are now possible
- Companies can locate servers far from expensive city locations, even rent servers "in the cloud"
- Back-end artificial intelligence processing for speech recognition, more

Growing Ubiquity and Reliability

- Almost never out of touch with the Internet and your resources there

The Emerging Internet of Everything

- Traditionally, there was a human user involved
- Growing technology allows devices to talk to one another, without human involvement
- These devices can now be very small, such as thermostats
- These devices now communicate by low-cost radio directly with one another

FIGURE 1-3 The Ever-Changing Internet (Study Figure)

However, what the Internet offers to people and organizations who use it is constantly new. Social media are relatively recent developments, as are high-quality video streaming and teleconferencing. Now we are beginning to see augmented reality, and not just to find and fight Pokémons. What will come next? Based on what we know about the past, it will surprise us. Since the emergence of the Internet, we have always been shocked when new "killer app" categories emerge to create a whole new set of billionaires and addicted users.

Growing Speed How can the Internet change so frequently and so radically in the applications it supports? The answer is that the Internet itself is changing technically at an enormous rate and will continue to do so. Simple speed is the most obvious change. Today, wired Internet connections can bring multiple high-definition videos to homes. Increasing velocity also allows you to use programs like Box and Dropbox to back up your files in real time and use them immediately or later despite their sizes. At the corporate level, companies even back up their massive transaction databases in real time. If one corporate site fails, another site can pick up the computing load almost instantly. At a very broad level, many companies have decided to stop buying and running their own servers. They have turned to cloud computing, in which you rent just the number of servers you need. If your load varies, you can even rent servers by the hour. Netflix, which generates about a third of the Internet traffic going into American homes in the evening, varies the number of servers it uses throughout the day. It even has a self-service portal to add and drop servers instantly. These cloud servers are clustered in massive server farms that each has thousands of servers under one roof. Where are these server farms? It really doesn't matter. The Internet can connect computers anywhere with minimal delay.

Growing Ubiquity and Reliability Another continuing disruptive change is the ubiquity of Internet access. Initially, you needed a desktop or at least a laptop computer. It was probably stationary in your home or office. When you were away from it, your access to the Internet depended on the presence of an Internet cafe or the kindness of friends. Mobile phones changed that, but only gradually. Early mobiles either could not access the Internet at all or were limited to a stupefying slow 10 kbps (10,000 bits per second). No mega, and certainly no giga. Today, speeds are far greater.

Of even greater importance, we can use the Internet everywhere. Our connection to the Internet is "always on." Our mobiles provide that to us nearly all the time, and we are increasingly able to plug into lower-cost (and higher-speed) Wi-Fi as we travel.

This always-on connectivity of mobile devices is used in ways we do not even realize. For example, the speed recognition processing needed for voice commands is usually done on a distant server with massive processing power, not on our puny little phones and tablets. This allows far richer and smarter interactions.

Along with this near ubiquity of Internet access is nearly perfect reliability. When the first author initially used the ARPANET (the forerunner of the Internet), he was astounded to see that he had new mail. It was from a colleague at MIT, welcoming him to the 'Net. A week later, he was still amazed, but like everybody else who was using networks, he also thought, "Too bad it doesn't work more often." Today, it does.

The Emerging Internet of Everything When the Internet was created in the late 1970s, computers were the size of rooms. Users worked at dumb terminals on their desks. These terminals were basically keyboards and low-quality displays. Microprocessors had just been invented, and they were too expensive for individuals to use. When the Internet was designed, it was widely assumed that only these large computers would connect to the Internet. However, Moore's Law forecast that microprocessor prices would soon fall dramatically and would continue doing so for many years. Personal computers (PCs) began to communicate. Then came smartphones.

Today, we increasingly have residential thermostats, air conditioners, and even coffee makers with enough processing power to run applications and communicate over the Internet. Furthermore, these devices increasingly talk to one another, with no human involvement. As noted at the beginning of this chapter, this trend is being called the **Internet of Things (IoT)**. In fact, the "devices" connected to the Internet may not even be physical. Today, a computer can run two or more virtual machines, which are programs and related data that act like full computers when they talk to real devices and humans on the Internet.

Test Your Understanding

5. a) What continuing changes in the Internet are contributing to its ability to support new applications constantly? b) What are the characteristics of the Internet of Things?

Owning and Managing the Internet

When the U.S. government pulled out of the Internet, the Internet needed a way to fund itself. This task was left to the ISPs. To use the Internet, you must connect through an ISP. Doing this is not free. As an individual or part of a family, you probably pay about

Commercial ISPs Handle Transmission

You must have an ISP to use the Internet
You pay the ISP money
Corporations pay a lot more
ISPs deliver packets across one another
Settlements for sharing revenue from users
Nobody owns the Internet. The ISPs do collectively

Nobody Controls the Internet, Either

The Internet Engineering Task Force (IETF) sets standards, but compliance is voluntary
A few things are centralized, including controlling Internet addresses to prevent duplication

FIGURE 1-4 Owning and Managing the Internet (Study Figure)

$50 per month to your ISP for Internet access. Organizations pay far more—often tens of thousands or even millions of dollars each year. Traffic must flow across ISPs, so the ISPs have financial settlement agreements among themselves to compensate for cross traffic.

Under these conditions, "Who owns the Internet?" The answer is, "nobody." Each ISP owns its own resources, and the Internet is the sum of these resources. This may seem like an odd situation, but this is exactly how the worldwide telephone network works. There are thousands of telephone companies around the world. Like ISPs, they exchange traffic and use financial settlements to balance costs and revenues.

An obvious related question is, "Who controls the Internet?" The answer, again, is, "nobody." A few things about the Internet are controlled. For example, the **Internet Assigned Numbers Authority (IANA)** controls internet addresses to avoid address duplication. However, remarkably little else is controlled.

What about standards? There is certainly a need for standards to govern how devices talk to one another. However, things are a little complicated. The organization that creates standards is the **Internet Engineering Task Force (IETF)**. This is a volunteer and sometimes rowdy organization that creates great standards. However, it has no power to impose these standards on ISPs and user organizations. In fact, quite a few of its standards have been ignored by ISPs. Keep this in mind when we talk about Internet standards created by the IETF.

The Internet Engineering Task Force (IETF) creates Internet standards.

Test Your Understanding

6. a) Who owns the Internet? b) Who is in charge of the Internet? c) What is the role of the IETF?

The Snake in the Garden

The Internet promises to give users access to almost everything, anytime, anywhere. Unfortunately, it does the same for criminals, national governments, and just plain jerks. As the Internet has grown in size and complexity, so have the adversaries and

Anything, Anytime, Anywhere

Works for Attackers As Well As Legitimate Users

Security Underlies Everything That Network Professionals Do

FIGURE 1-5 Security: The Snake in the Internet Garden (Study Figure)

the attacks they use. Networking practitioners are not the only professionals who are responsible for stopping security threats, but security underlies almost everything that networking professionals do. We will hold off looking more deeply into security until Chapter 4. This is not because security is unimportant but because you need a solid grasp of networking concepts, standards, and management before you can understand security threats and countermeasures.

Test Your Understanding

7. a) Why is the Internet's ability to give broad access a good thing? b) What danger does it bring?

Next Steps

So far, we have been looking at the Internet at a very high level. For the rest of this chapter (and this book in general), we look at the Internet and other networks in the detail that professionals in IT, networking, and security need to understand to enter the profession. In this chapter, we focus on the core Internet terms and concepts we will see throughout this term.

This introduction ends with a fundamental point. So far, we have been talking about the Internet. However, the Internet is not the only network. In fact, "Inter" means "between." The Internet was specifically created to link many individual networks together. We begin with the Internet, but later in the chapter we also look at two types of networks that can be standalone networks or parts of the Internet.

The Internet is not the only network.

Test Your Understanding

8. a) What does "Inter" in Internet mean. b) Why is this important?

OUTSIDE THE INTERNET

We will spend most of this term looking *inside the Internet* and other networks. However, we begin by looking at the Internet *from the outside*—focusing on the user devices attached to it. Figure 1-6 shows some devices attached to the Internet. However, it depicts the Internet itself as an opaque cloud. The **cloud** indicates that the average

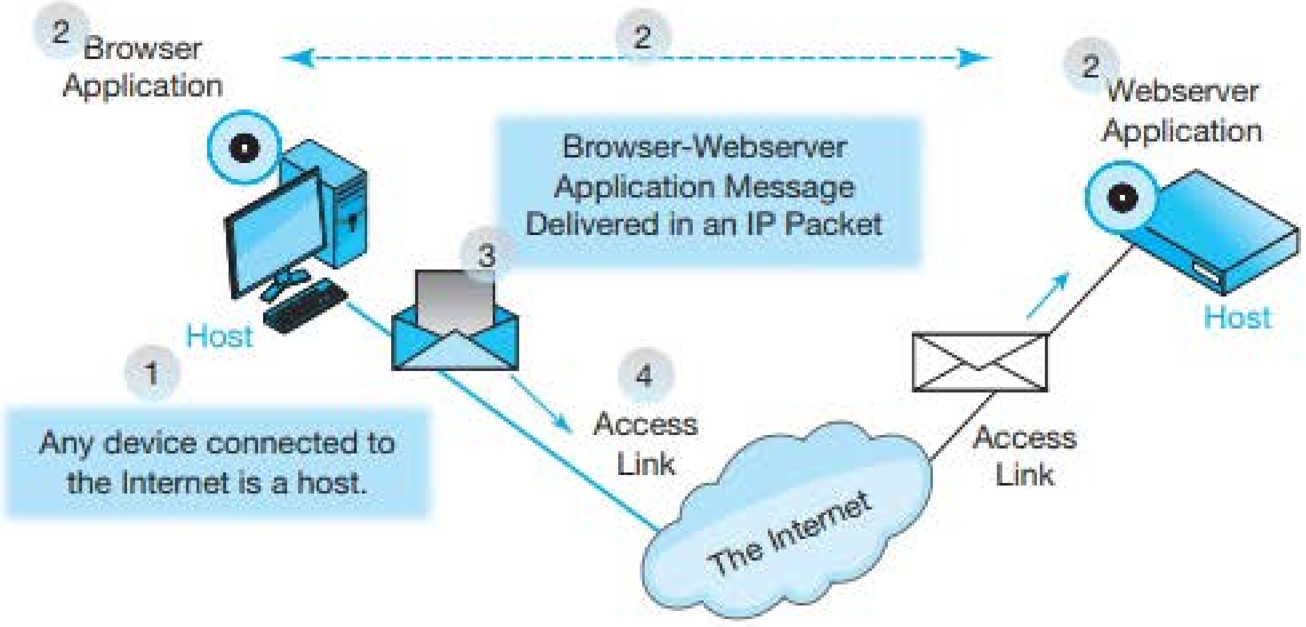
2 Browser Application
2
2 Webserver Application
Browser-Webserver Application Message Delivered in an IP Packet
3
Host
Host
1
4
Access Link
Access Link
Any device connected to the Internet is a host.
The Internet

FIGURE 1-6 Outside the Internet

user does not have to know what is happening inside the cloud. Things (should) just work. The electrical system works in a way similar to this cloud. When you turn on a light switch, you do not have to know how the electricity is delivered to you. It just is. Depicting the Internet and other networks as a cloud is very common.

Figure 1-6 shows that the job of the Internet is to deliver application messages from application programs running on hosts. As we saw earlier, a host is any device connected to the Internet.[12] Your mobile phone, tablet, and PC are all hosts.

Note, however, that the Internet is not just about connecting hosts. It is about connecting *applications* running on these hosts, and it is about connecting them by delivering **application messages** between them.

The Internet connects hosts by delivering application messages between them.

Test Your Understanding

9. a) Why is the Internet often depicted as a cloud? b) Why is the Internet not about sending messages between hosts?

Client and Server Hosts

Most hosts are clients or servers. **Server hosts** provide services to **client hosts**. For example, when you browse the Web on your mobile phone, your mobile phone is a client host. The server is the webserver to which you send requests. It processes your requests and sends you the files you specified.

[12] When the Internet was first designed, it was assumed that only large computers would have enough processing power to connect to it. Typically, these computers served users at dumb terminals without processing power. These large computers were usually called hosts in standards documents. As processing power grew cheaper, PCs began to connect to the Internet directly, and today we even have Internet coffee pots. Faced with these changes, the Internet Engineering Task Force had to rewrite its early standards with a more inclusive name or broaden the name "host" to embrace smaller devices. Intellectual energy conservation won.

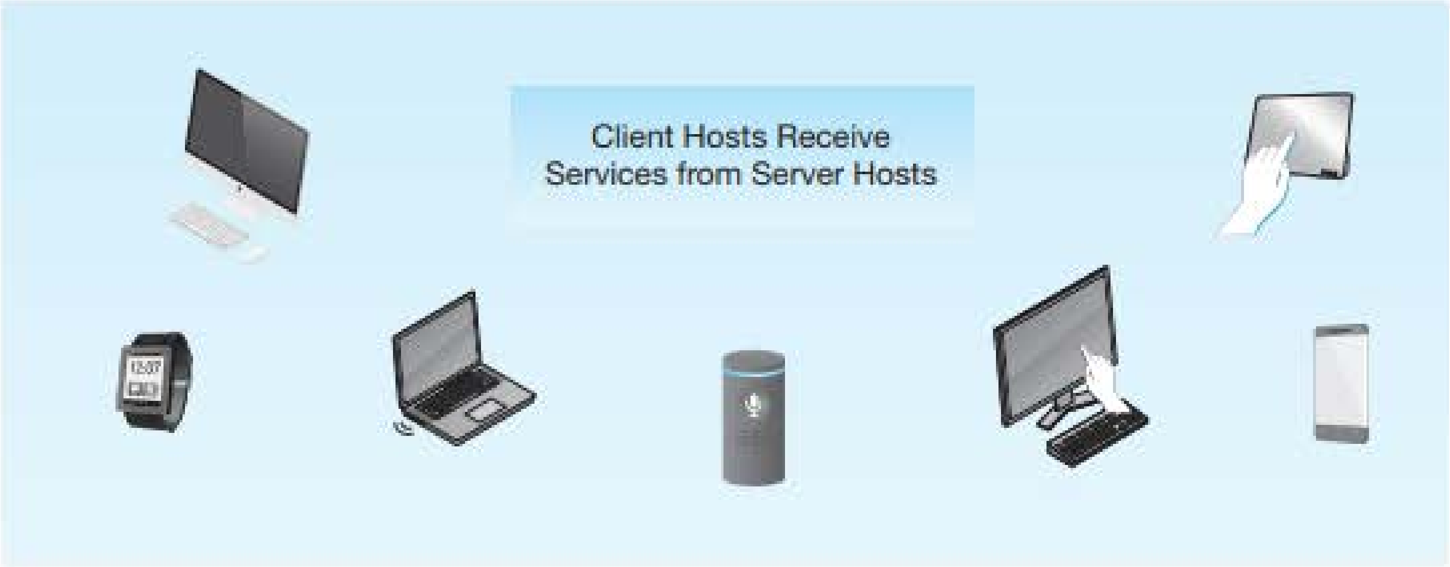

FIGURE 1-7 Client Hosts

Server hosts provide services to client hosts.

Client Hosts As a user, you are personally most familiar with client hosts. Client hosts include your mobile phone, tablet, and PC. They also include your fit bit, smart watch, or anything else you use to access services on the Internet. Figure 1-7 shows some common client hosts you may use.

Server Hosts You personally see client hosts every day. However, you may never have seen a server. If you suspect that servers are interesting to see, you will be disappointed. Most servers are stored in equipment racks that are 48 cm (19 inches) wide. These **rack servers** are installed one above another—often a dozen or more servers in a rack. Rack servers are usually plain-looking boxes, with a few connections on the outside, usually at the rear. Rack server heights are measured in multiples or fractions of U, which is 1.75 inches. The smallest are ½U tall. Larger servers are 2 to 5Us in height. Figure 1-8 shows three server racks and one rack server being installed. Computer centers and server farms have hundreds or thousands of racks.

Most servers today are rack servers that fit into 48 cm (19 inches) wide equipment racks.

Small size does not mean that rack servers have little power. For example, Netflix delivers streaming content to users via open connect appliance servers. (An appliance is something you just plug in and use, like a toaster.) Each of these appliances is about 4U tall. The fastest can stream 90 Gbps of content to users.[13] This is enough to give 2,000 customers simultaneous streaming high-definition video.[14]

[13] Michelle Clancy, "Netflix Moves All Global Traffic to Open Connect CDN," http://www.rapidtvnews.com/2016031942170/netflix-moves-all-global-traffic-to-open-connect-cdn.html#axzz4YY3OCBs8.

[14] In fact, even small rack servers are too powerful for many uses. It is common for a single server's power to be divided into several virtual servers. Virtual servers are programs that act like physical servers.

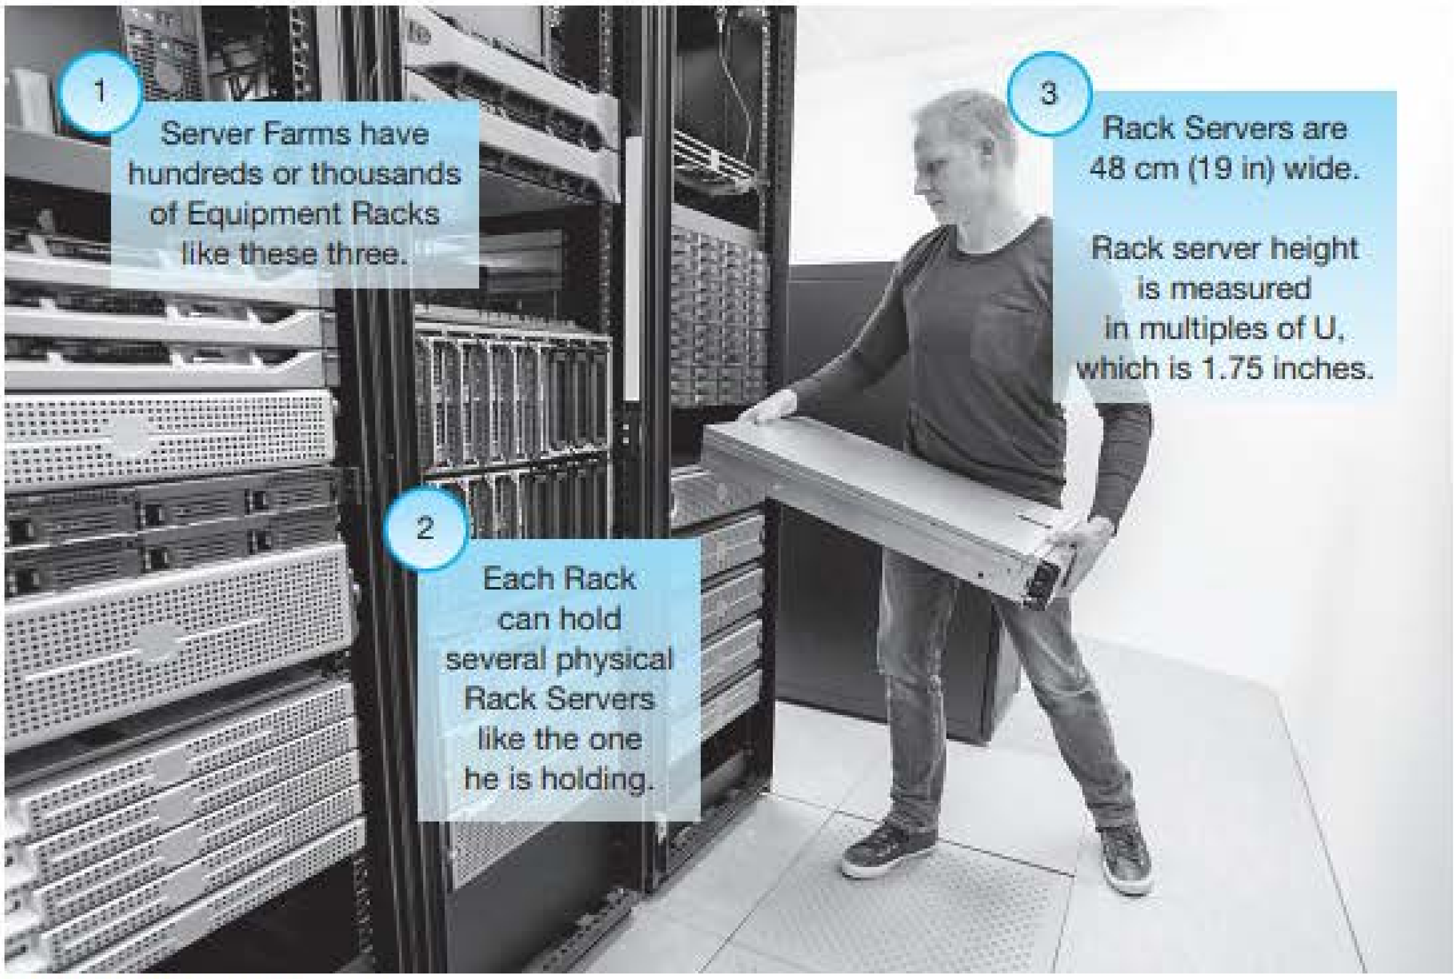
1
Server Farms have hundreds or thousands of Equipment Racks like these three.
2
Each Rack can hold several physical Rack Servers like the one he is holding.
3
Rack Servers are 48 cm (19 in) wide.
Rack server height is measured in multiples of U, which is 1.75 inches.

FIGURE 1-8 Rack Server Host

Test Your Understanding

10. a) Distinguish between client and server hosts. b) What type of devices are most servers?

Networked Applications

Networked applications are simply those that require a network to communicate with one another. For example, when you use the Internet, you have a browser on your device. Your browser communicates with a webserver application program on a webserver. Figure 1-9 illustrates this situation.

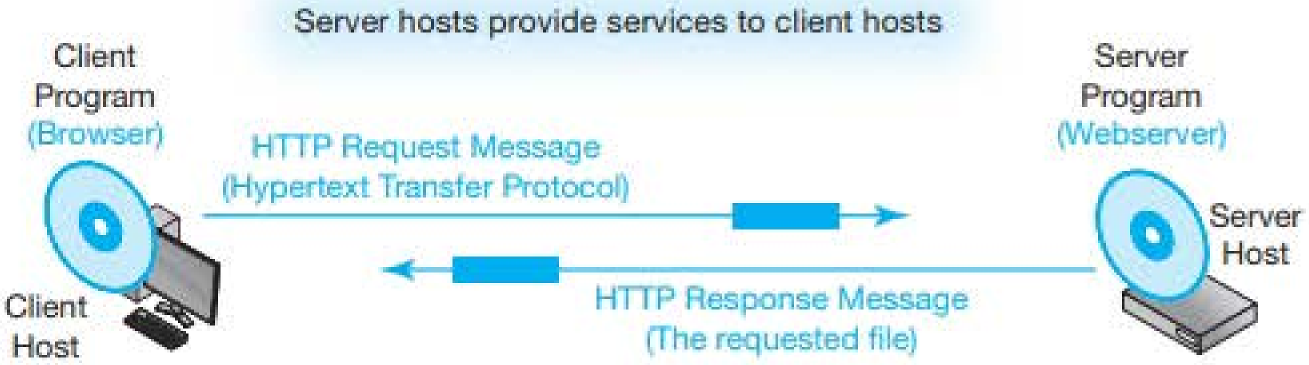
Server hosts provide services to client hosts
Client Program (Browser)
HTTP Request Message (Hypertext Transfer Protocol)
Server Program (Webserver)
Server Host
Client Host
HTTP Response Message (The requested file)

FIGURE 1-9 Client/Server Application: Webservice

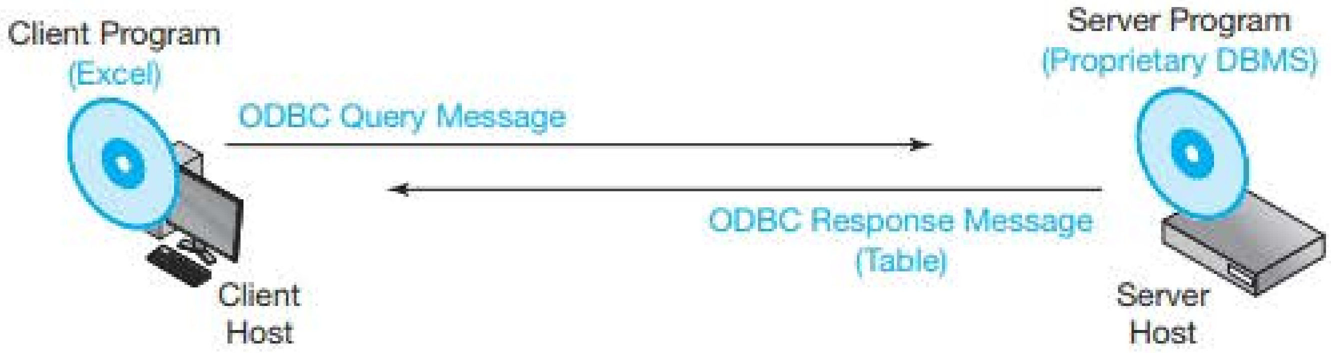

FIGURE 1-10 Not Always Browsers and Webservers: Excel Querying a Propriety Database Management System (DBMS) Using the Open Database Connectivity Protocol (ODBC)

Networked applications are simply those that require a network to communicate with one another.

Your browser, which is your **client program**, sends a **request message** to the webserver. This is an **HTTP request message** because **HTTP (Hypertext Transfer Protocol)** is the standard for browser–webserver interactions. This request message asks for a file to be delivered. The webserver **server program** on the webserver locates the file and sends it back in an **HTTP response message** containing the requested file (or an error message to say why it could not be delivered).

Browsers and webserver application programs are networked applications, but they are certainly not the only ones. Figure 1-10, for example, shows Microsoft Excel acting as a client program. It is using the Open Database Connectivity Protocol (ODBC) to query a proprietary database. Other client/server networked applications include Dropbox and Skype. The key point is that the client program is not always a browser and the server application is not always a webserver program.

The client program is not always a browser, and the server application is not always a webserver program.

In general, the days of writing a program to run on a single machine are rapidly disappearing. Today, we usually write a program on one machine to work with a program on another machine. Ideally, networks would simply work transparently, making them irrelevant to programmers. However, reality often falls short of the ideal. Programmers who do not understand networking are strongly limited. So are database professionals, e-commerce professionals, and data analytics professionals.

Test Your Understanding

11. a) What are networked applications? b) Is the client always a browser? c) Is the server always a webserver?

The Job of the Source Host

During transmission, a **source host** sends an application message to a **destination host**. Let's look at that process in a little more detail.

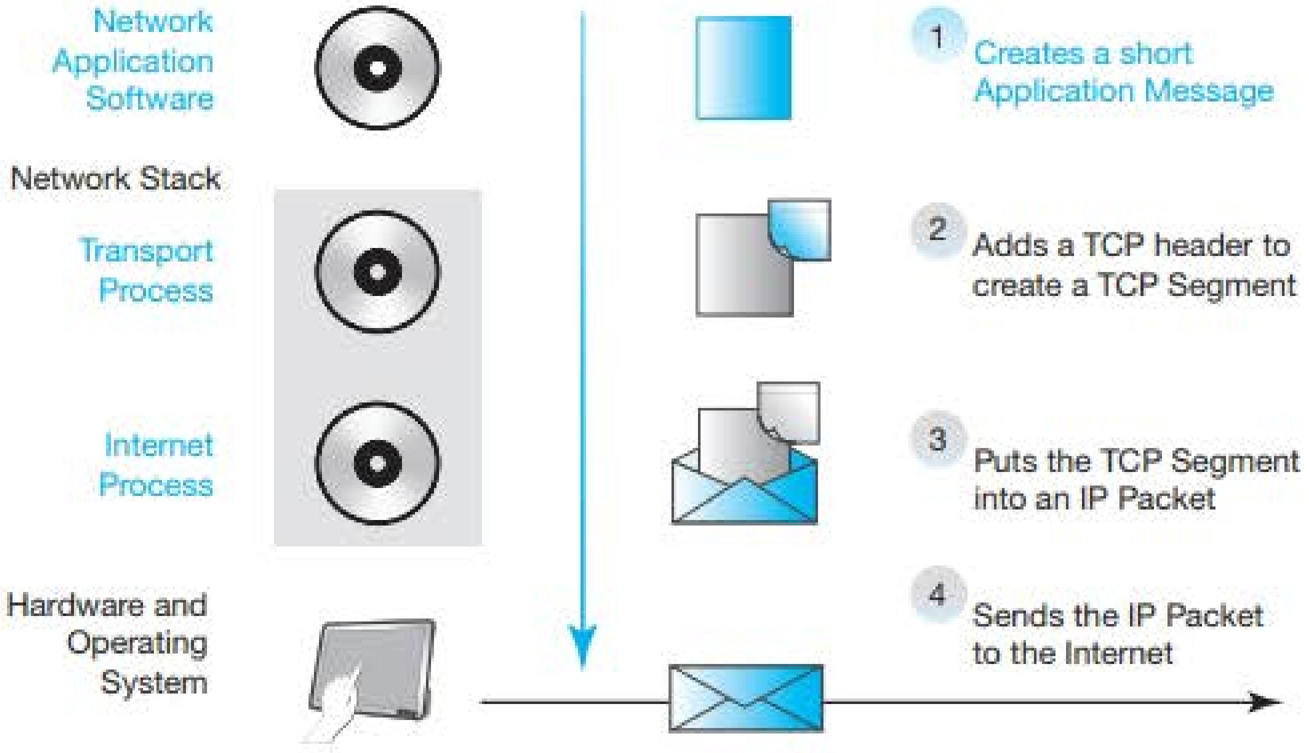
Network Application Software
Network Stack
Transport Process
Internet Process
Hardware and Operating System
1 Creates a short Application Message
2 Adds a TCP header to create a TCP Segment
3 Puts the TCP Segment into an IP Packet
4 Sends the IP Packet to the Internet

FIGURE 1-11 On the Source Host: Sending a Short Application Message

During transmission, a source host sends an application message to a destination host.

Sending Short Application Messages Figure 1-11 shows what the source host does when its application creates a message for the application on the destination host. IP packets have limited size. A short application message is one that is small enough to fit into a single packet. A single packet can be up to 65,536 bytes in size. Most are smaller.

IP packets have limited size. A short application message is one that is small enough to fit into a single packet

First, the application program creates the application message. This message is designed to be read by the application program on the destination host.

Second, the application program sends the application message to the network stack on the source host. The **network stack** is a small group of programs that will govern the submission of the application message to the Internet and the reception of incoming Internet messages.

- For short application messages, the **transport process** in the network adds a short **Transmission Control Protocol (TCP)** header to create a **TCP segment**. More on this later.[15]

[15] For many applications, the transport process creates UDP headers instead of TCP headers. We will see the distinction between UDP and TCP in Chapter 2. We use TCP in examples in this chapter because we want to introduce application message fragmentation. UDP requires that application messages be short enough to fit into a single packet. In contrast, TCP can handle application messages short enough to fit in a single packet or messages that must be sent in multiple packets.

- The transport process then passes the short application message down to the **internet process** in the network stack. The internet process adds an IP header. Effectively, it places the TCP segment in an envelope called a **packet**.

Point of Terminology:

We use Internet (with an uppercase I) for the global Internet that serves users and when Internet is in the name of a protocol (for instance, the Internet Protocol).

We use internet (with a lowercase i) for the internet process and other things.

The computer's hardware and operating system then submits the packet to the Internet. The Internet does the rest.

Short application messages fit into single packets.

The Final IP Packet Figure 1-12 shows the final IP packet. It begins with the application message. The transport process adds a TCP header. The Internet process adds an IP header. (We will see these headers in the next chapter.) This is the complete IP packet. When the packet is transmitted, the IP header is transmitted first, then the TCP header, and then the application message.

Sending Long Application Messages Things are a bit more complex for long application messages that are too long to fit into a single IP packet. (As just noted, the maximum packet size is 65,536 bytes.) In this case, Figure 1-13 shows that the transport process first fragments the application message into several fragments, placing each in a separate **TCP segment**. As Figure 1-12 shows, the TCP segment header has a sequence of numbers so the application message fragments can be put back in order at the other end.

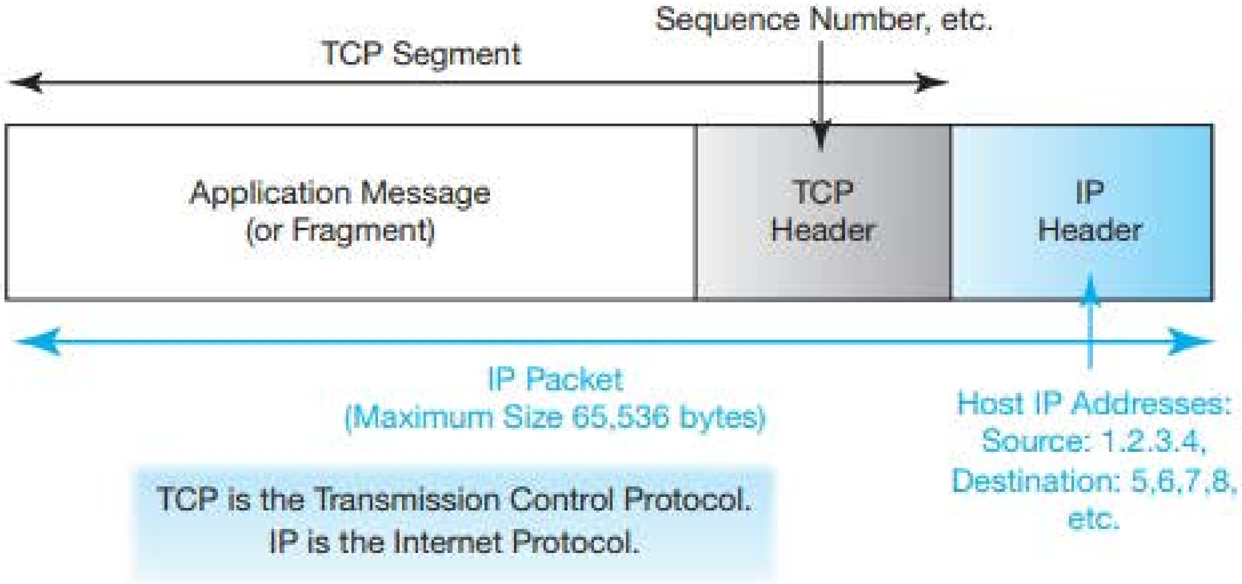

TCP is the Transmission Control Protocol.
IP is the Internet Protocol.

FIGURE 1-12 The Final IP Packet

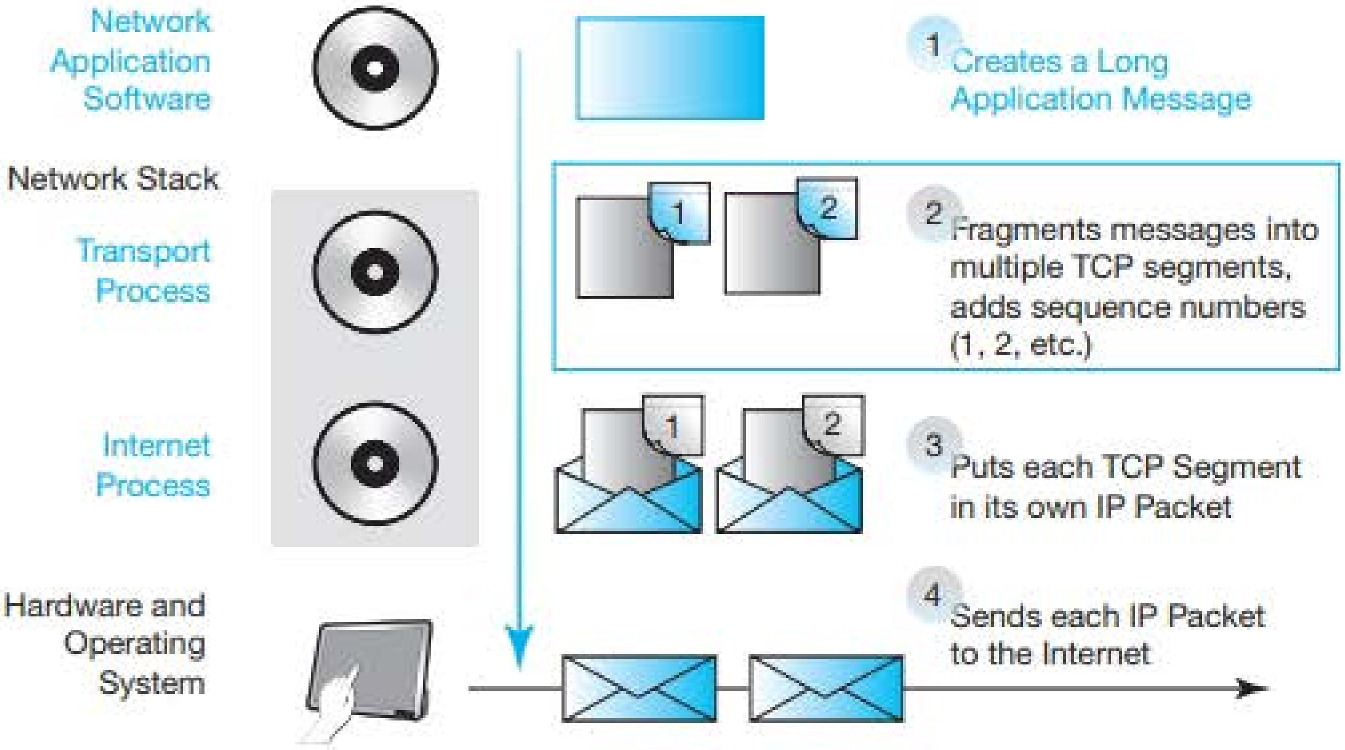
Network Application Software
Network Stack
Transport Process
Internet Process
Hardware and Operating System
1 Creates a Long Application Message
2 Fragments messages into multiple TCP segments, adds sequence numbers (1, 2, etc.)
1
2
3 Puts each TCP Segment in its own IP Packet
4 Sends each IP Packet to the Internet

FIGURE 1-13 On the Source Host: Sending a Longer Application Message

The transport process then sends each segment down to the internet process. The internet process places each TCP segment into a separate IP packet and submits the packets to the Internet for delivery.

Long application messages must be divided into smaller fragments, each of which is placed in its own TCP segment, which is placed in its own IP packet.

Test Your Understanding

12. a) What two processes does the network stack provide? b) What is the maximum size of an IP packet? c) What does the transport process do to the application message if it is short enough to fit in a single packet? d) If the application message is too long? e) What does the transport process add to the application message or fragment? f) What is the resulting message called? g) What does the internet process do with each TCP segment?

13. What are the three parts of an IP packet?

The Job of the Destination Host

Earlier, we looked at what happens on the source host. Figure 1-14 completes the picture by showing what happens on the destination host. The internet process pulls each TCP segment from its packet, reassembles the TCP segments in order, and passes the reassembled application message to the destination application program.

Freeing the Application Program from Networking Details Note that the application program is not involved with networking details. It merely creates application messages and receives application messages of any length. Segmentation and reassembly? Not its concern. Putting things in packets? Not its concern either. We say

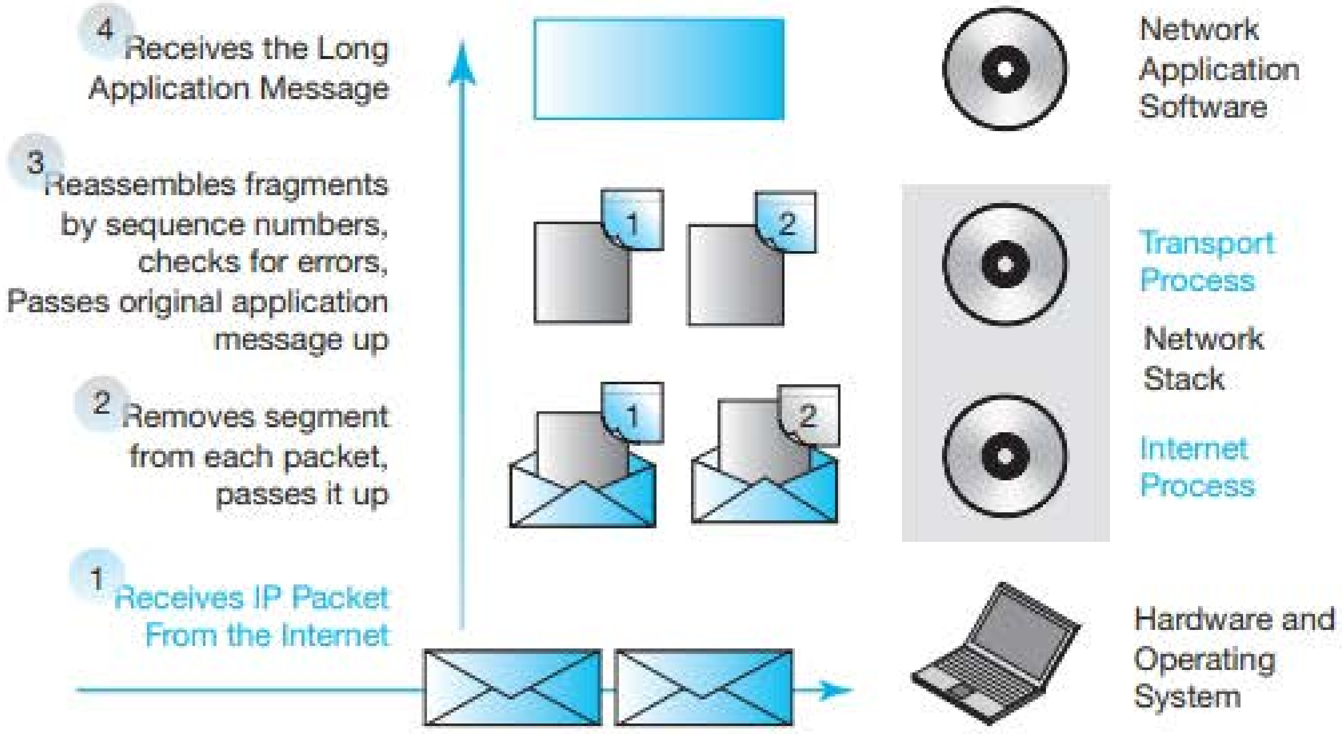
4 Receives the Long Application Message
3 Reassembles fragments by sequence numbers, checks for errors, Passes original application message up
2 Removes segment from each packet, passes it up
1 Receives IP Packet From the Internet
1
2
1
2
Network Application Software
Transport Process
Network Stack
Internet Process
Hardware and Operating System

FIGURE 1-14 On the Destination Host: Receiving a Long Application Message

that networking is **transparent** to application programs. This means that application programmers can focus on writing their applications, not worrying about how application messages will be delivered over a network.

Test Your Understanding

14. a) What does the internet process on the destination host do when a packet arrives for it? b) What does the transport process on the destination host do with multiple TCP segments from a single application message? (This answer is not short.)

INSIDE THE INTERNET

So far, we have been treating the Internet as an opaque cloud, focusing on what happens on the hosts and network applications that are *outside* the Internet. (The discussion was also a way to begin sneaking in broader concepts, such as the distinction between application functionality, transport functionality, and internet functionality.)

The Main Characters: IP Addresses, Packets, Routers, Data Links, and Routes

Now we finally look *inside* the Internet to see how it functions at a very broad level. Figure 1-15 shows the main elements we see inside the Internet. These are IP addresses, packets, routers, data links, and routes.

IP Addresses

We start with IP addresses because they are the key to understanding everything else. If you want to call Bob on his mobile phone, you need to know Bob's phone number. Similarly, hosts need addresses so that the Internet can deliver packets to the right hosts.

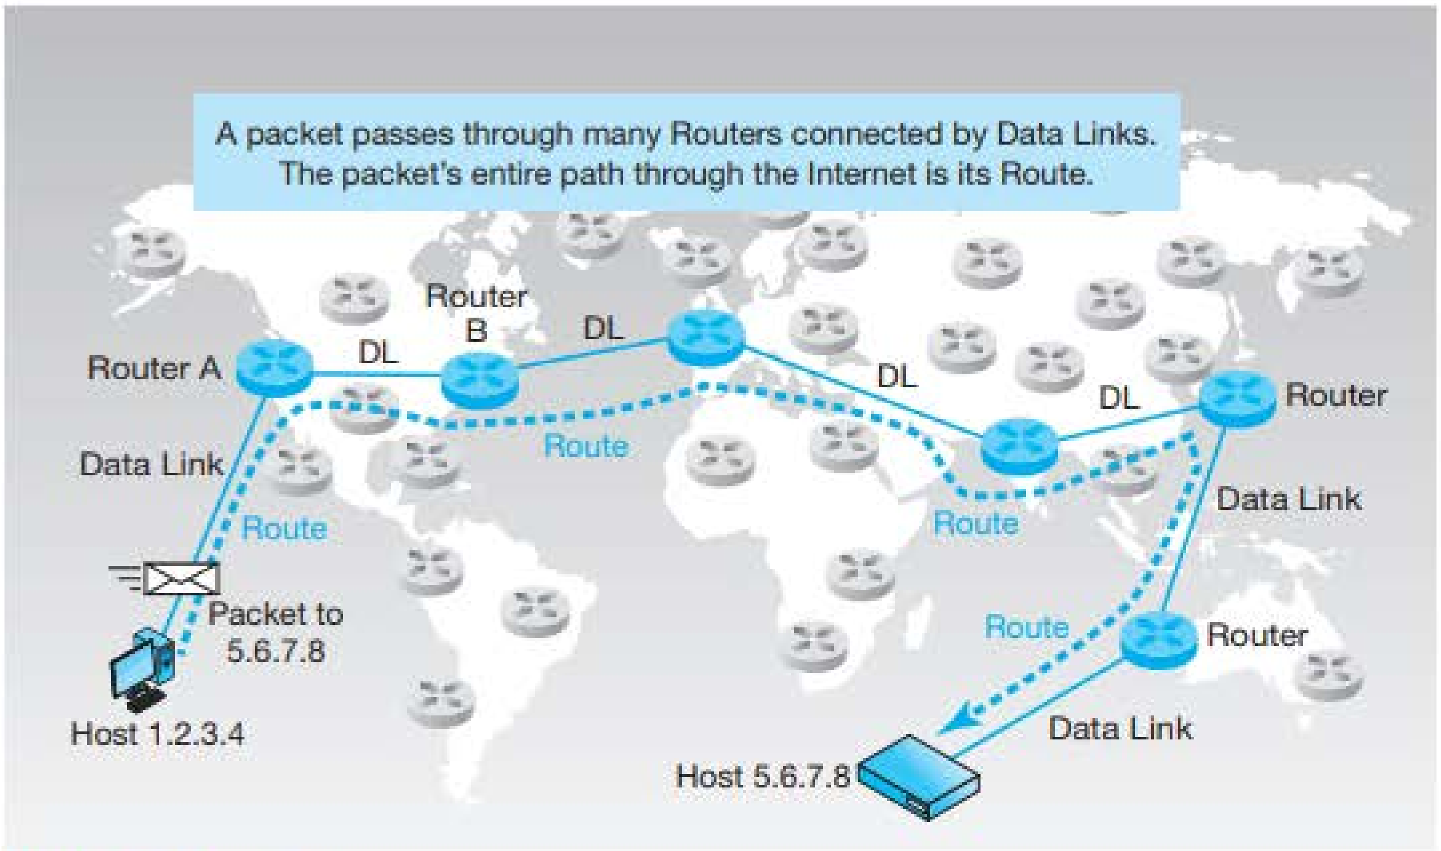

FIGURE 1-15 Inside the Internet

These are **Internet Protocol (IP) addresses**. The first generation of IP addresses were **IP Version 4 (IPv4) addresses**.[16] They were 32 bits (1s and 0s) long.

Routers and hosts have no problem reading and writing 32-bit strings. Humans have a lot of problems with them. As an aid to inferior biological entities such as ourselves, IPv4 addresses are usually written in **dotted decimal notation**. In this notation, they are written as four decimal integers separated by dots (periods). Each number represents a group of 8 bits. An example of an IPv4 address is 1.2.3.4. Another is 127.171.17.13.

Each host on the Internet needs an IP address to receive IP packets.

For human reading and writing, IPv4 addresses are shown as four decimal integers separated by dots; this is dotted decimal notation.

Binary to Decimal Figure 1-16 shows how to write 32-bit IPv4 addresses in dotted decimal notation (DDN).

- First, the 32 bits are divided into four 8-bit "segments" (not to be confused with TCP segments).
- Then, each 8-bit segment is treated as a binary number and converted into a decimal integer. For example, 00000000 is 0 in decimal, 00000001 is 1 in decimal, and 11111111 is 255 in decimal. You can use Excel's *bin2dec* function to do the conversion. Most advanced calculators will do it as well.
- Next, the four segment numbers are put together and separated by dots. Hence the name "dotted decimal notation."

[16] There were no Versions 0, 1, 2, or 3.

Devices Use 32-Bit IP Addresses Directly

32-bit IPv4 Address: 10101101000101101100001010101011

This is too difficult for humans to read and write.

Humans Write IPv4 Addresses in Dotted Decimal Notation

DDN is easier to read, write, and remember for inferior biological entities.

Start with 32-bit IPv4 Address.	10101101000101101100001010101011			
Divide it into four 8-bit segments.	10101101	00010110	11000010	10101011
Convert each segment to a decimal integer.*	173	22	194	171
Place dots between segments.	173.22.194.171			

**In Excel, bin2dec(10101101) = 173*

FIGURE 1-16 IPv4 Addresses and Dotted Decimal Notation for Humans

Decimal to Binary You can also reverse this process to go from dotted decimal notation back to native binary IP addresses. However, keep in mind that Excel dec2bin and other calculation approaches treat the result as a binary number rather than what they are—strings of 8 bits. For example, if you use dec2bin to covert 22 to binary, you will get the answer 10110. You must add three leading zeroes 00010110 to get 8 bits.

IP Version 6 Addresses We show IP addresses as IPv4 addresses in dotted decimal notation in figures and examples here. Newer **IP Version 6 (IPv6)**[17] addresses are becoming widespread today. IPv6 addresses are 128 bits long. As we will see in Chapter 8, writing IPv6 is more complicated than dotted decimal notation, and it seems best to avoid this added complexity until later.

Test Your Understanding

15. a) How many bits long are IPv4 addresses? b) Convert 00000001 00000010 00000000 11111111 to dotted decimal notation (spaces have been added). (Note: 00000001 is 1) c) Convert 5.6.0.255 to a 32-bit IP address (add spaces between groups of 8 bits). (Note: 5 is 0000101, not 101)

IP Packets

Figure 1-12 showed a final IP packet. Note that an IP header contains a **source IP address** and a **destination IP address**. These give the IP addresses of the source host (sender) and the destination host (receiver). Routers use a packet's IP destination addresses to deliver the packet to its destination.

Test Your Understanding

16. a) What are the three parts of an IP packet? (Yes, this is a repeat of an earlier question.) b) In which part will you find the source and destination IP addresses? c) Which of these addresses will routers use to deliver the IP packet?

[17] There was an IPv5, but it was never made an official standard.

Routers

When a host transmits a packet, it sends the packet to a router. A **router**[18] is like a railroad switch yard. It receives an arriving packet, then forwards it to another router closer to its destination host. An IP packet may travel through dozens of routers as it passes through the Internet to the destination host.

> *A router receives an arriving packet, then forwards it to another router closer to its destination host.*

Routing Figure 1-17 shows how routers work in slightly more detail. In the figure, an IP packet arrives at Router A. The packet is addressed to destination host 60.3.27.46. Router A must send it on to a router closer to the destination host. In the figure, the router has two choices. It may forward the packet to either Router B or Router D, which are the only routers it connects to that will move the packet closer to the destination host. A router's process for forwarding packets is called **routing**. A router's forwarding decision is called a **routing decision**.

> *A router's process for forwarding packets is called routing. A router's forwarding decision is called a routing decision.*

In Chapter 8, we will see how Router A makes its decision. For now, suffice it to say that Router A will make its decision intelligently, sending the packet back out in the best way for the packet to reach its destination host.

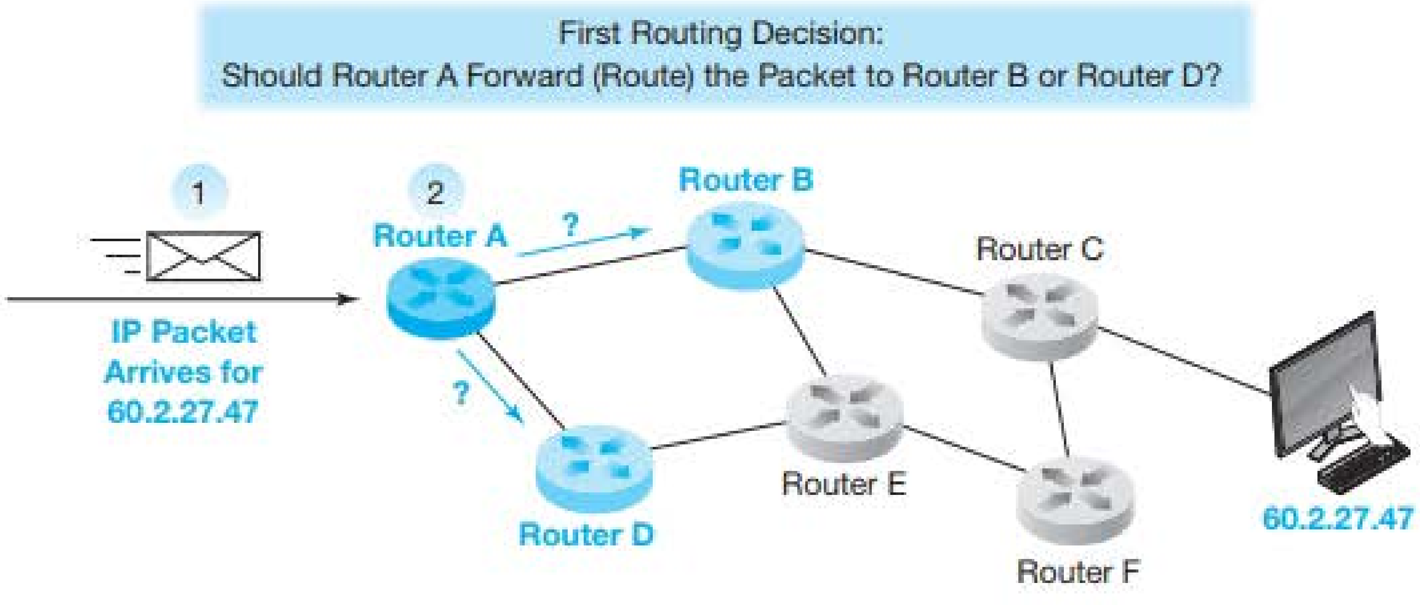
First Routing Decision:
Should Router A Forward (Route) the Packet to Router B or Router D?
1
2
Router A
?
Router B
Router C
IP Packet Arrives for 60.2.27.47
?
Router E
Router D
Router F
60.2.27.47

FIGURE 1-17 Routing (Router Forwarding)

[18] How do people pronounce "router?" It depends where you are from. Americans usually say "rowter," with the ow being pronounced like the ow in "now." Pretty much everybody else pronounces it "rooter."

On Router B Suppose that Router A decides to send the packet to Router B. Router A will then transmit the packet to Router B. When Router B receives the packet, it must make its own routing decision. Router B also has two choices. It can route the packet either to Router C or Router E. You may guess that it will forward the packet to Router C because the packet will then be only one more hop away from the destination host. Sending it to Router E will require two more hops. In practice, however, a router takes many things into account when it makes its routing decision, not just the number of hops.

Test Your Understanding

17. a) What does a router do when an IP packet arrives? b) What is router forwarding called? c) In Figure 1-17, suppose that 60.3.27.47[19] transmits a packet to 128.171.17.13. When Router C receives the packet, what will be its routing choices?

Data Links and Routes

Figure 1-15 shows that **data links** are transmission links that carry packets between *pairs of routers*. Packets travel over these data links to move between routers. Note the term "link" instead of "line." Data links often use wireless transmission instead of physical lines, so the neutral (and vague) term *link* is used.

A data link is the transmission path of an IP packet between two routers.

We also need a name for the entire path a packet takes between the source host and the destination host, across multiple routers and transmission links. It is called a **route**. It is very easy to confuse data links and routes, but their distinction pervades Internet thinking, and you need to distinguish between them clearly in your mind.

The route is the packet's entire path between the source host and the destination host.

Test Your Understanding

18. a) Distinguish between data links and routes. b) In Figure 1-15, how many data links will there be when the packet travels to Host 5.6.7.8? c) How many routes will there be? d) In general, when a source host sends a packet to a destination host, will there probably be more data links or routes along the way? Explain. (The answer is not in the text.)

[19] A professor at Pomona College "proved" that all numbers are equal to 47. This did not catch on in mathematics. However, Pomona College has produced many writers who tend to have an affinity for the number. This is especially obvious in science fiction. Nearly every episode of the *Star Trek* series has the number 47 in it. The second author of this book went to Pomona College, but the first author is solely to blame for the frequent use of 47 in this book.

The Transport and Internet Processes in the Network Stack

When the designers of the Internet considered moving packets over their budding creation, they knew that they faced two conflicting requirements.

- First, as we will see in Chapter 8, the routers would have to do considerable work on each packet. To keep router costs reasonable, this work should be limited as much as possible.
- Second, to provide adequate quality of service, the Internet would have to provide error detection and correction, so that application programs got "clean" data. In addition, packets would sometimes arrive out of order. A method would be needed to put them back in order. Also, most application messages would be fragmented to fit in packets. These would each create significant cost. This was particularly true for error detection and correction, which involves considerable mathematical processing.

To meet these conflicting requirements, the Internet's designers decided to divide the work of Internet transmission into two parts. Figure 1-18 shows how they did this.[20]

- They would create a standard that would be used for the source host to transmit a packet to the Internet, to move packets between routers, and for the final router to

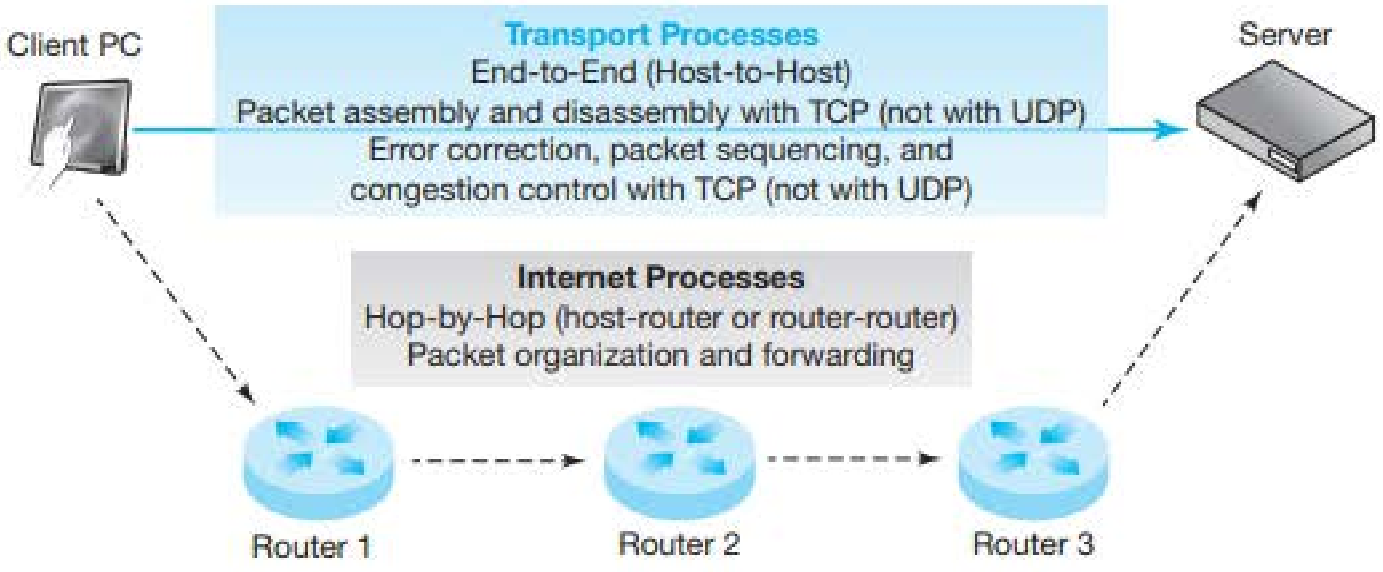
Client PC
Transport Processes
End-to-End (Host-to-Host)
Packet assembly and disassembly with TCP (not with UDP)
Error correction, packet sequencing, and
congestion control with TCP (not with UDP)
Server
Internet Processes
Hop-by-Hop (host-router or router-router)
Packet organization and forwarding
Router 1
Router 2
Router 3

FIGURE 1-18 The Transport and Internet Processes

[20] A historical note may aid your understanding. (Or may not. That is why it is in a footnote.) Initially, there was only a single Internet transmission standard, the Transmission Control Protocol (TCP). It handled both what we now call transport and internet functionality. Before the Internet was finalized, however, the IETF decided that the standard was becoming too complex, so they divided the standard into two smaller parts. The Internet Protocol was created to govern internet matters, and TCP was restricted to transport matters. This division also allowed a second transport standard to be created, the User Datagram Protocol, which we will see in Chapter 2 and following chapters.

In the next chapter, we will see that error correction is done only once, by the transport processes on the source and destination hosts. If it were done by the internet process, it would have to be done on each router hop along the way. That would slow delivery while placing a greater processing burden on each router, raising router prices substantially. There are a lot of routers on most routes, so the internet process in general is stripped down to do as little as possible with each packet while still getting it across the Internet.

deliver the packet to the destination host. This would be the Internet Protocol (IP). The IP would have to be executed on every router along the route, so it was kept as simple as possible.

- The Internet Protocol would not handle difficult work like error correction. That would be done by a transport protocol, such as the Transmission Control Protocol (TCP). Note in Figure 1-18 that unlike IP, which would be executed on each hop, transport protocols would only be done on the source and destination hosts. This meant that heavy processes such as error correction would only have to be handled once, on the two hosts.

Overall, several internet processes on the source host, the destination host, and intermediate routers are involved in IP packet transmission, but only the two transport processes on the two hosts are active.

Test Your Understanding

19. a) There are six routers between the source and destination host. How many transport processes will be involved? Explain. b) How many internet processes will be involved? Explain.

Supervisory Standards: Beyond TCP and IP

We have seen that packet transmission among routers is governed by the Internet Protocol. The transport process was governed by the Transmission Control Protocol (TCP) in our examples. In the next chapter, we will see that the transport process also has an alternative protocol, the User Datagram Protocol (UDP). These three standards govern most Internet activity because they are all that is necessary to deliver a packet, and delivering packets is the main work of the Internet.

Supervisory Protocols However, from the beginning the Internet was created to be a worldwide network. This required the creation of **supervisory protocols** beyond the IP, TCP, and UDP delivery protocols. To give you a sense of what supervisory protocols do, let's look at two supervisory protocols that users deal with extensively, DHCP and DNS.

IP, TCP, and UDP are sufficient to deliver IP packets between hosts, which is the main job of the Internet. However, these three protocols must be supplemented by many supervisory protocols to do the additional work that is needed beyond what IP, TCP, and UDP do.

Test Your Understanding

20. Why does the Internet need supervisory protocols?

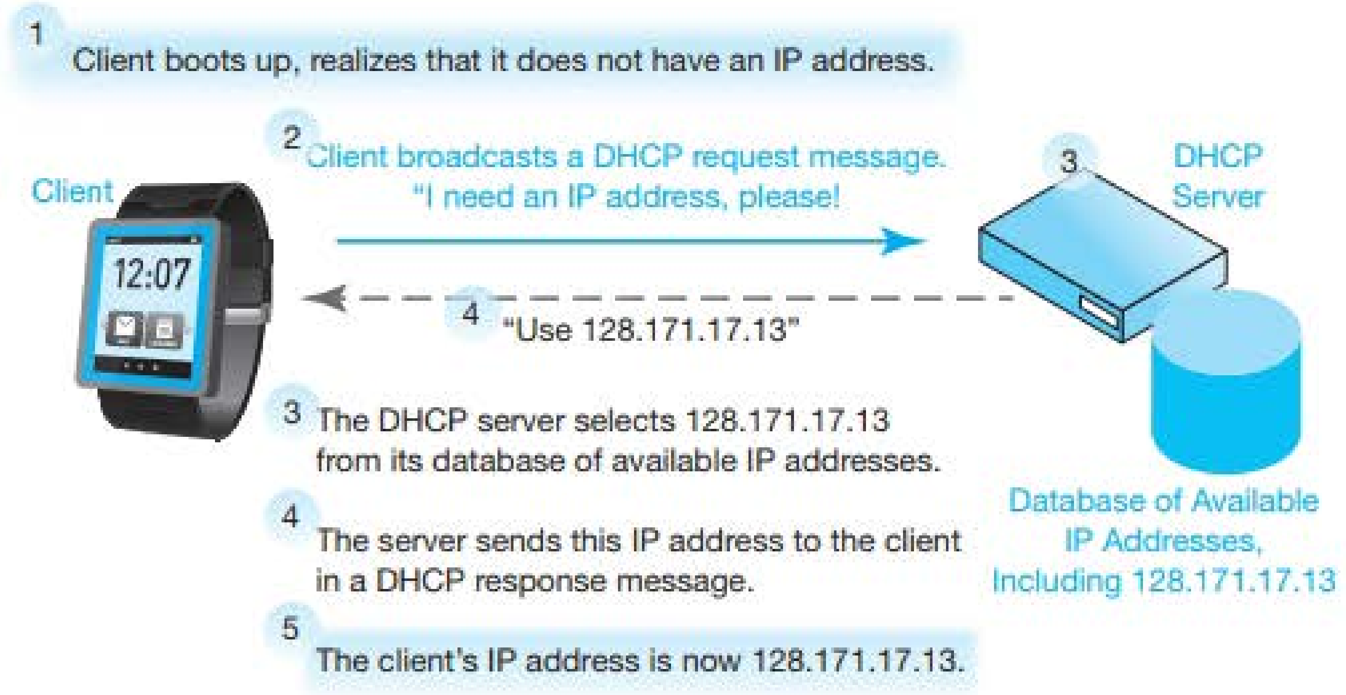
1 Client boots up, realizes that it does not have an IP address.
2 Client broadcasts a DHCP request message. "I need an IP address, please!
Client
12:07
3 DHCP Server
4 "Use 128.171.17.13"
3 The DHCP server selects 128.171.17.13 from its database of available IP addresses.
4 The server sends this IP address to the client in a DHCP response message.
Database of Available IP Addresses, Including 128.171.17.13
5 The client's IP address is now 128.171.17.13.

FIGURE 1-19 Dynamic Host Configuration Protocol (DHCP)

Dynamic Host Configuration Protocol (DHCP) The host you use to surf the Web or do other tasks needs an IP address. How does it get its address? The answer is that client hosts get their IP addresses using the **Dynamic Host Configuration Protocol (DHCP)**. As Figure 1-19 illustrates, when a client device boots up, it realizes that it does not have an IP address. It broadcasts a DHCP request message to its local DHCP server.[21] This message asks for a temporary IP address to use. The server finds an available IP address in its database and responds by sending the client a DHCP response message that includes the IP address. For subsequent packets sent by the client, this is the packet's source IP address. This type of address is called a **dynamic IP address**.

When the client shuts down, it forgets the IP address. The next time the client boots up, it contacts the DHCP server for a new IP address to use. It typically receives a different IP address each time it does this.

DHCP servers typically give a client a different IP address each time it boots up.

What about servers? Servers need stable IP addresses, which are called **static IP addresses.** (Imagine trying to shop at a business that keeps moving so that it has a different address each day. Hmm, sounds a bit illegal.) A network technician types the static IP address into the server host's configuration file, and DHCP is not used at all.

[21] Broadcasting is necessary because the client does not know anything about the network, including the IP address of the local DHCP server. To broadcast the DHCP request message, the client makes the destination IP address thirty-two 1s. When a router receives a packet with an all-1s destination IP address, it broadcasts the IP address to all nearby hosts. All hosts read all broadcast packets. Only the DHCP server responds. If more than one DHCP server responds, the client selects one of them.

What source IP address does a client host use to send the packet containing the DHCP request message? It does not have one yet, so it places thirty-two 0s in the source IP address field of the packet.

Test Your Understanding

21. a) What type of host gets a dynamic IP address? b) What type of host gets a static IP address? c) Why is a static IP address needed for this type of host? d) Does a DHCP server give a host the same IP address each time?

Domain Name System (DNS) IPv4 addresses are difficult to write and remember, even in dotted decimal notation. In Chapter 8, we will see that iPv6 addresses are even longer, and it is rare to write one of these addresses correctly in the first attempt, much less remember it. To address human limitations, the Internet allows host owners to create **host names** for their servers. In Figure 1-20, the host name of server 128.171.17.13 is Voyager.shilder.hawaii.edu. This is still long, but it is far easier to remember and write. When you use a host, you probably know its host name. You rarely know its IP address.

However, routers can only work with IP addresses. They know nothing about host names. If you type in a host name, your computer needs to resolve it, that is, determine the IP address associated with that host name.

(1) In Figure 1-20, a host wishes to send packets to Voyager.shidler.hawaii.edu. The host wishing to do so is the originating host. Voyager is the target host.

(2) To find the host's IP address, the originating host sends a DNS request message to a Domain Name System (DNS) server. This message gives the target host's host name and asks for its IP address.

(3) The DNS host looks up Voyager.shidler.hawaii.edu in its DNS Table. In notes that the IP address for Voyager is 128.171.17.13.

(4) The DNS server sends back a DNS response message to the originating host. This response message gives the IP address of Voyager.

(5) Finally, the originating host can send packets to Voyager by addressing them to 127.171.17.13. Now that it knows the IP address, it has no more need for the DNS

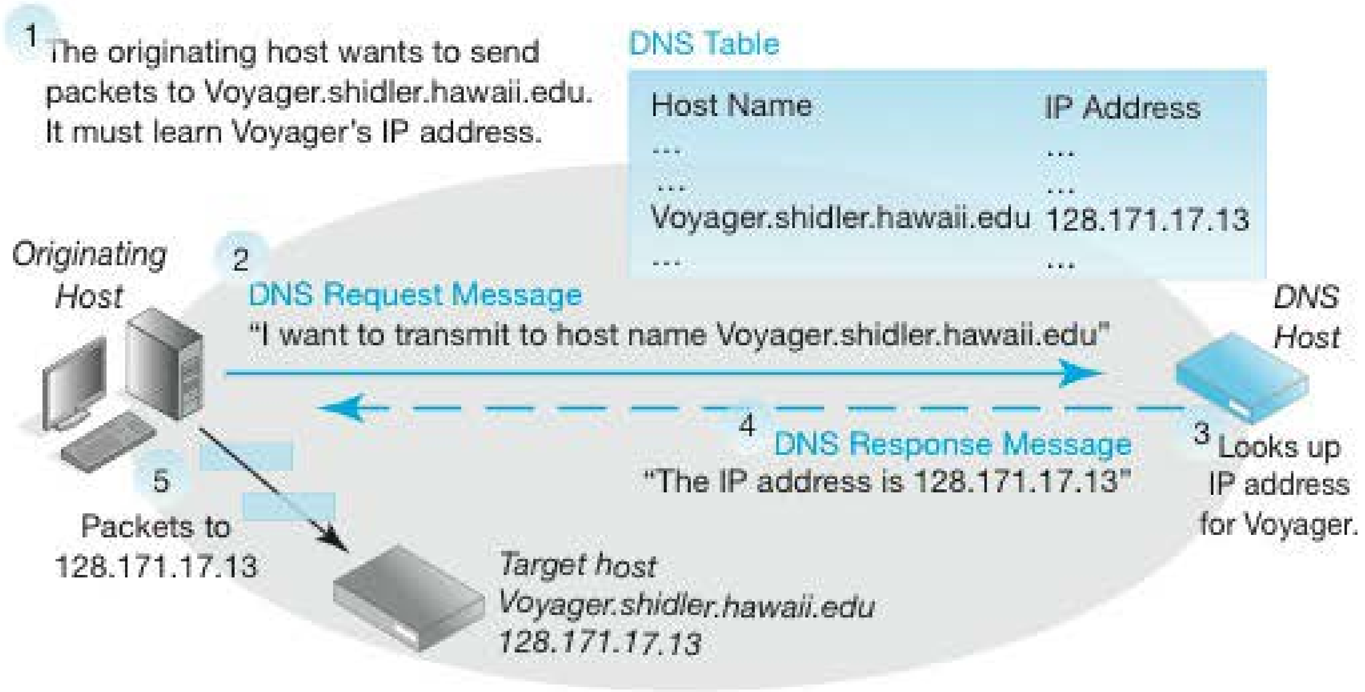
1 The originating host wants to send packets to Voyager.shidler.hawaii.edu. It must learn Voyager's IP address.
DNS Table
Originating Host
2 DNS Request Message
"I want to transmit to host name Voyager.shidler.hawaii.edu"
DNS Host
4 DNS Response Message
"The IP address is 128.171.17.13"
3 Looks up IP address for Voyager.
5 Packets to 128.171.17.13
Target host Voyager.shidler.hawaii.edu 128.171.17.13

Host Name	IP Address
...	...
...	...
Voyager.shidler.hawaii.edu	128.171.17.13
...	...

FIGURE 1-20 Domain Name System (DNS)

host. The originating host continues to send packets to the target host without subsequent calls to the DNS server.

Although not shown in the figure, when the originating host learns the IP address of a host name, it stores this information in its local DNS cache. If it wants to reach Voyager.shidler.hawaii.edu a few days later, it looks up the IP address from its DNS cache. There is no need to use the DNS server.

Test Your Understanding

22. a) Distinguish between the originating host, the DNS server, and the target host. b) What is the purpose of a DNS lookup? c) Does the originating host need to contact the DNS host each time it sends a packet to the target host? Explain.

SINGLE NETWORKS, DATA LINKS, AND PHYSICAL LINKS

We saw earlier that transmission *data links* connect hosts to routers and routers to other routers. We will now see that although data links sound simple, the way they provide these connections is sometimes complex.

Point-to-Point Single Networks

To understand data links, you need to understand a concept called the single network. A **single network** is a network that uses a single set of standards for all devices. There are many single network standards, and they are deeply incompatible. If Host A is on one single network and Host B is on another of a different type, they cannot communicate.[22]

A single network is a network that uses a single set of standards for all devices.

If one host is on one single network and another is on a single network of a different type, they cannot communicate.

Figure 1-21 shows the simplest type of single network technology. This is a **point-to-point network**. It works on a direct point-to-point physical connection between two hosts. Not much of a network, you are probably saying. This is true, but its simplicity makes it a good place to begin talking about single network standards. In addition, it is used in many connections between pairs of routers on the Internet.

At the heart of the point-to-point network is the direct physical connection. This is defined by a **physical standard**. A physical standard covers three things: the transmission medium (optical fiber, radio transmission, etc.), a physical connector on each device, and how 1s and 0s are transmitted over this physical link.

[22] In fact, even if two networks use the same single network standard, they still may not be able to communicate because the same single network address may be used in both single networks.

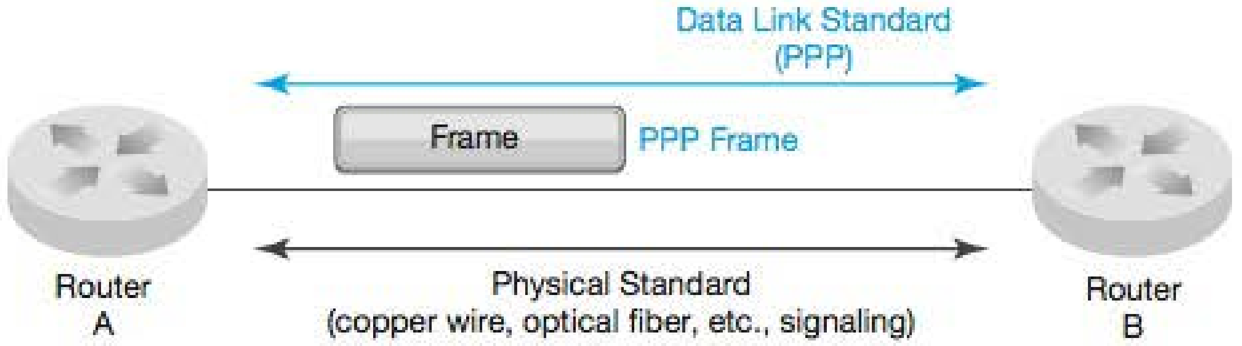
Data Link Standard (PPP)
Frame
PPP Frame
Router A
Physical Standard (copper wire, optical fiber, etc., signaling)
Router B

FIGURE 1-21 Point-to-Point Single Network Using the Point-to-Point Protocol (PPP) Data Link Standard

A physical standard covers three things: the transmission medium (optical fiber, radio transmission, etc.), a physical connector on each device, and how 1s and 0s are transmitted over this physical link.

A single network also needs standards for data links. For data links, the bits of each message are organized into a message called a **frame**. It is not a packet. The data link standard governs how the frame is organized. In addition, data link standards govern how switches, access points, and other single network forwarding devices forward frames. We will see more about data link standards in the next subsection when we look at Ethernet switched networks.

A message in a single network is a frame, not a packet.

The data link standard governs how the frame is organized.

The data link standard in this point-to-point network is the appropriately named **Point-to-Point Protocol (PPP)**. There are other data link standards for point-to-point networks, but PPP dominates. In particular, PPP is almost always used when point-to-point networks are used to connect pairs of routers.

The Point-to-Point Protocol (PPP) is the most common data link protocol for point-to-point single networks.

Test Your Understanding

23. a) Distinguish between physical links and data links. b) In a point-to-point single network, how many physical links will there be when a packet is transmitted? c) How many data links?

Ethernet Single Networks

Another widely used single network standard is **Ethernet**, which was created for switched local area networks. Figure 1-22 illustrates an Ethernet LAN with three switches, two routers, two hosts, and six physical links between switches, routers, and hosts when Router A sends a frame to Router B through this network.

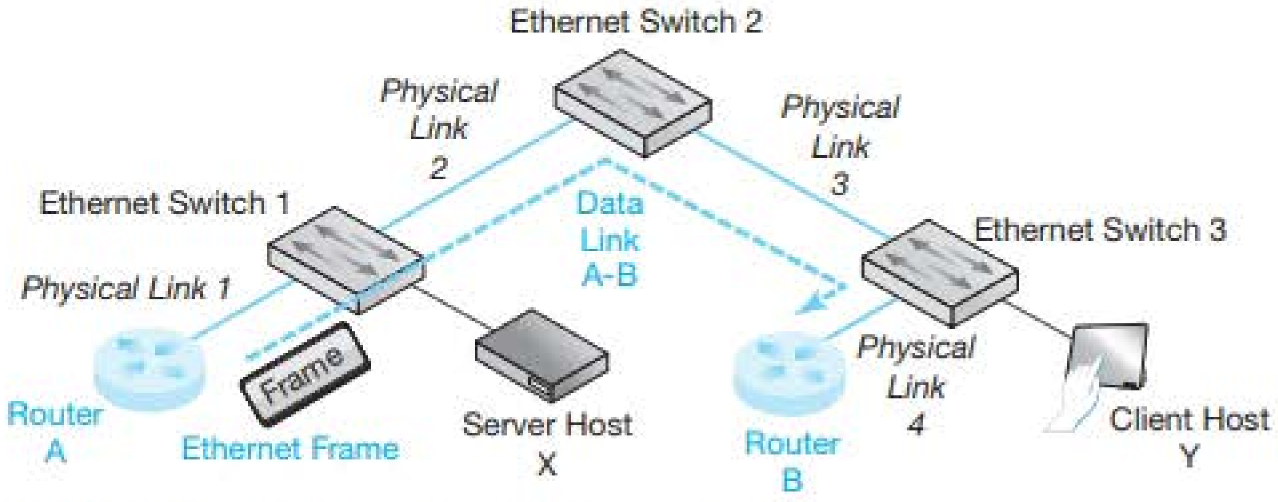

FIGURE 1-22 Ethernet Switched Single Network

Ethernet Frames and Data Links In the figure, this small Ethernet network connects two routers, Router A and Router B.

- Router A sends an **Ethernet frame** to Router B. The router transmits this frame over Physical Link 1 to Ethernet Switch 1.
- Switch 1 forwards (**switches**) the frame over Physical Link 2 to Ethernet Switch 2.
- Switch 2 forwards the frame over Physical Link 3 to Switch 3.
- This final switch forwards the frame over a fourth physical link. The frame then arrives at Router B.

The path that the frame travels through this single network is a **data link**. In fact, a frame's path from the source device to the destination device through a single network of any type is called its data link.

A frame's path from the source device to the destination device through a single network of any type is called its data link.

Ethernet Physical Links (versus PPP Physical Links) In PPP, there always is a single physical link and a single data link when a frame is transmitted. In Ethernet, there is also a single data link, but there usually are multiple physical links.

By definition, there is always a single data link when a source device sends a frame to a destination device through a single network.

The number of physical links the frame travels over ranges from one with PPP to many for Ethernet.

Test Your Understanding

24. a) In Figure 1-22, how many physical links will there be when Router A sends a Packet to Router B? (Answer: 4) b) How many data links will there be? (Answer: 1) c) When Client Host Y sends a packet to Router B, how many physical links will there be? d) Data links? e) When Client Host Y sends a packet to Server Host X, how many physical links will there be? f) Data links?

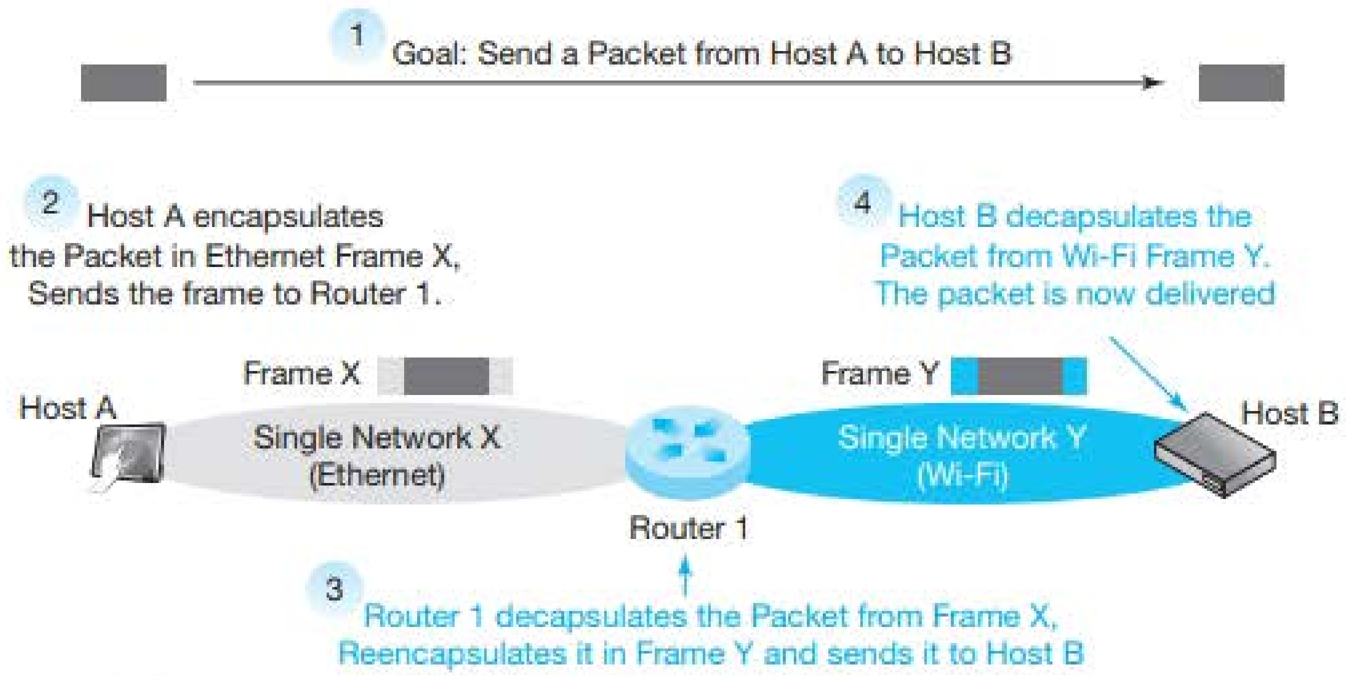
1 Goal: Send a Packet from Host A to Host B
2 Host A encapsulates the Packet in Ethernet Frame X, Sends the frame to Router 1.
4 Host B decapsulates the Packet from Wi-Fi Frame Y. The packet is now delivered
Frame X
Frame Y
Host A
Single Network X (Ethernet)
Single Network Y (Wi-Fi)
Host B
Router 1
3 Router 1 decapsulates the Packet from Frame X, Reencapsulates it in Frame Y and sends it to Host B

FIGURE 1-23 Packets Are Carried Inside Frames in Single Networks

Frames and Packets

In single networks, messages are *frames*. On the Internet, they are *packets*. These two concepts are not separate. They are deeply intertwined, and the way they are related is the key to how the Internet functions.

Figure 1-23 shows two single networks. Single Network X is an Ethernet network. Network Y is a Wi-Fi network. These two networks use different frame forwarding methods. (A Wi-Fi network uses an access point to forward Wi-Fi frames within the network.) A single frame could not travel from the source host to the destination host across these two very different network technologies.

This is where packets come in. A single packet must travel all the way from the source host (Host A) to the destination host (Host B). However, this packet must always be carried inside a frame when it travels through a single network. Single networks understand frames and what to do with them. They have no idea what a packet is.

A packet must always be carried inside a frame when it passes through a single network. Single networks understand frames and what to do with them. They have no idea what a packet is.

- In Network X, the source host places the packet in an Ethernet frame (Frame X) and sends this frame to Router 1.
- The router takes the packet out of Frame X. It places the packet in a new frame, Frame Y. This is a Wi-Fi frame. It transmits this Wi-Fi frame containing the packet to the destination host over Network Y.
- The destination host takes the packet out of Frame Y. The packet has reached its destination.

In this example, there were only two single networks. There was a single packet (there always is a single packet), and there were two frames. What if the packet

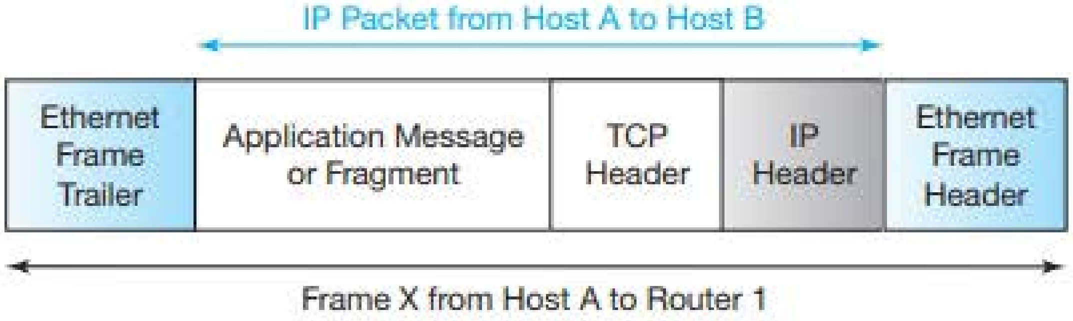
IP Packet from Host A to Host B
Ethernet Frame Trailer
Application Message or Fragment
TCP Header
IP Header
Ethernet Frame Header
Frame X from Host A to Router 1

FIGURE 1-24 Packet Encapsulated within the Ethernet Frame That Host A Sends to Router 1 (in Figure 1-23)

had to travel through 100 networks? There would still be a single packet (there is always a single packet). (Note the statement in parenthesis. It's a big deal.) However, the packet would be carried in 100 different frames along the way, one in each network.[23] This process of encapsulating the packet into a frame in each single network means that the Internet can contain millions of single networks with many different technologies.[24]

The encapsulation of packets inside frames is one of the central concepts of how the Internet works. To reinforce this, Figure 1-24 illustrates the Ethernet frame that Host A sends to Router 1 (Frame X). As we saw earlier, a packet contains an IP header, a TCP header (or UDP header), and an application message or fragment of an application message. The Ethernet frame begins with an Ethernet Header and ends with an Ethernet Trailer. We will see these in Chapter 5.

Test Your Understanding

25. a) Are packets carried inside frames, or are frames carried inside packets? b) A host sends a packet to another host. There are ten single networks along the way. How many hosts will there be? c) How many data links? d) How many routes? e) How many frames? f) How many packets? g) To what device will the first host send a frame? h) To what device will the final router send a frame?

[23] A historical note may help you understand why packets are carried inside a frame. Initially, frames carried application messages, pure and simple. There were no packets. The genius of Cerf and Kahn, who created the principles behind the Internet, was realizing that they could lie. When the source host transmitted Frame X, it expected the frame to go to another host to deliver its application message. Cerf and Kahn saw that they could place a router there instead. The router would pretend to be a host. However, the router knew that the frame contained a packet, not an application message. The router decapsulated the packet, put it into a frame on Network Y, and sent the frame on to the destination host. Note that this required the router to act like a host on Network X and a different host on Network Y. This is a lot of lying, but routers are shameless that way.

For this to work, hosts have to be in on the deception. Each has a network stack of software that intercepts the application message to be put into a frame, puts it into a packet (or several packets), and passes the packet to the data link process that handles frames. The network stack on the destination host reverses the process. The data link processes have no idea that what they receive is a packet instead of an application message.

[24] What if the two single networks in Figure 1-23 are both Ethernet networks? Will the single frame simply be passed on? The answer is no. Even if a packet travels through a hundred single networks using the same technology, the whole decapsulated-encapsulate-send process will be used on each router. It is simpler to have one rule that is always followed than to have exceptions for various pairs of single networks that follow the same standards.

Single Network Addresses

Packets are delivered to IP addresses. The source and destination IP addresses are placed in the packet header. Frames are delivered to **data link addresses** within a single network. For instance, Ethernet frames are delivered to Ethernet addresses. The source and destination Ethernet addresses are placed in the frame header.

Ethernet data link addresses follow the **Extended Unique Identifier-48 (EUI-48)** standard. An identifier is an address. These addresses must be unique. Extended means, well, never mind. You may have heard that Ethernet uses **Media Access Control (MAC) addresses.** It used to. Recently, the name MAC was changed to EUI-48. There was probably a good reason for this. In any case, you need to know what Ethernet addresses are called today. By the way, Wi-Fi also uses EUI-48 addresses. And yes, they too used to be called MAC addresses.

IP addresses are 32 bits long and are written for humans as four integers separated by dots. An example might be 1.2.3.4. EUI-48 addresses are 48 bits long. As we will see in Chapter 5, Ethernet addresses are written for humans to look something like A1-BB-DE-19-C3-4F.

Each host is on both a single network and the Internet. Therefore, each host needs to have two addresses. For a host on an Ethernet network, its address is its EUI-48 address. Its address on the Internet is its IP address. For a router, if the router connects to two networks, it will have a different EUI-48 address on each network it connects to.

Figure 1-25 adds addresses to Figure 1-23. Host B has the IP address 5.6.7.8. The packet that Host A sends to Host B is addressed to 5.6.7.8.

The frame that travels through Single Network X goes as far at the router. The frame's destination data link address is therefore the data link address of the router on Network X. This is the EUI-48 address A1-BB-DE-19-C3-4F. The packet it carries, again, is addressed to the IP address 5.6.7.8.

In turn, the frame that the router sends to the destination host is addressed to the EUI-48 address of the destination host. This is B2-23-FF-9F-CA-DE. The packet is still addressed to the IP address 5.6.7.8, the IP address of the destination host.

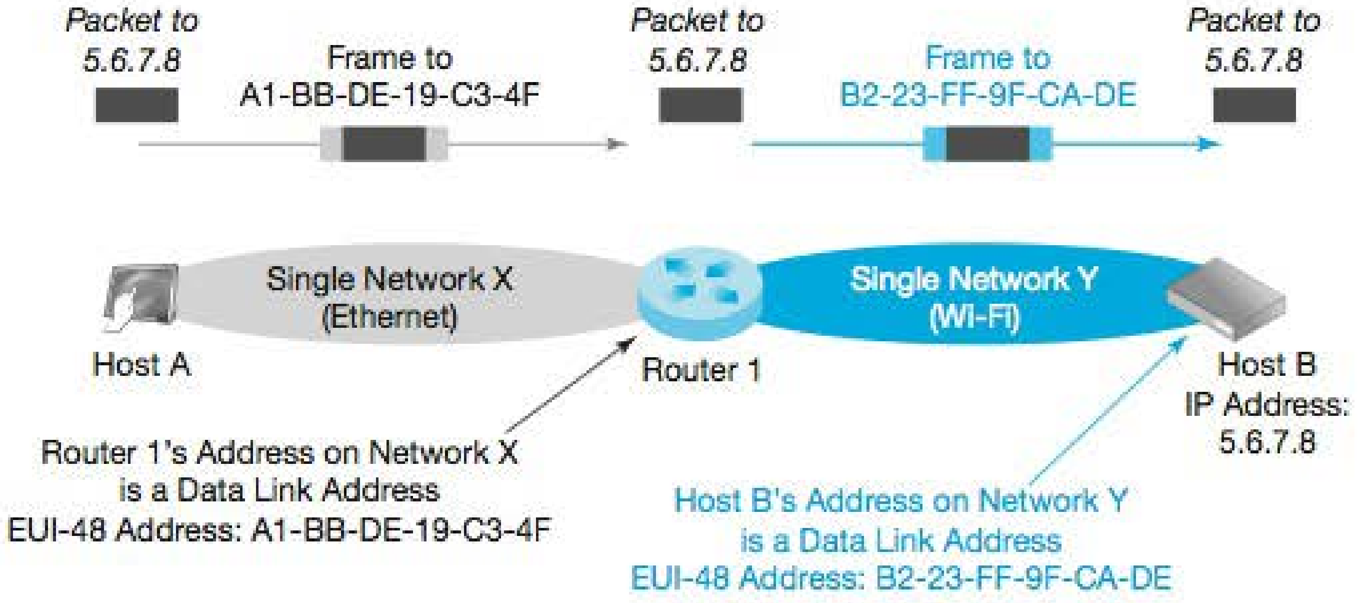

FIGURE 1-25 Packet Transmission Through Two Single Networks with Addresses Added (Based on Figure 1-23)

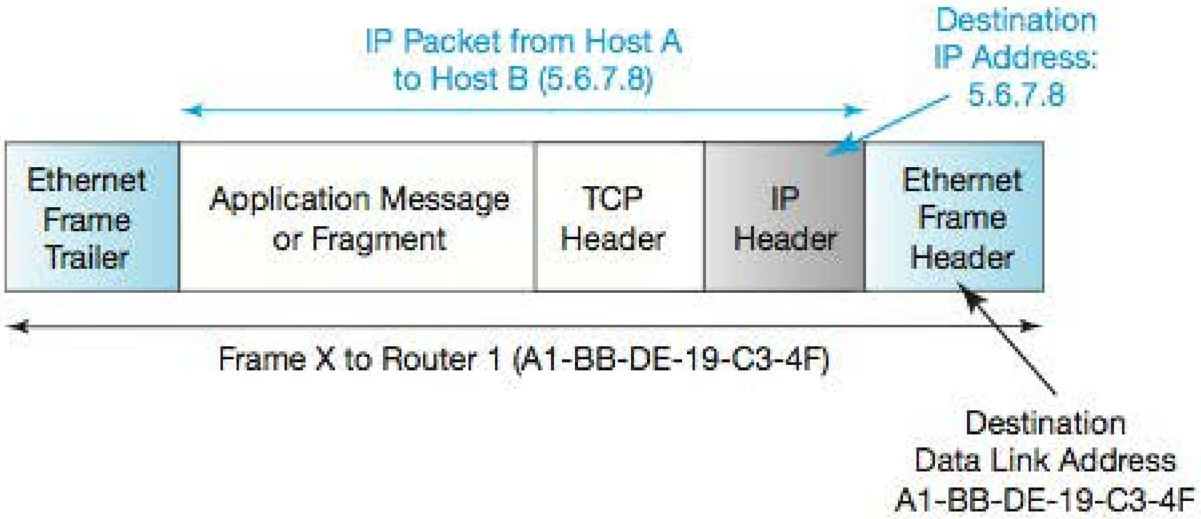
IP Packet from Host A to Host B (5.6.7.8)
Destination IP Address: 5.6.7.8
Ethernet Frame Trailer
Application Message or Fragment
TCP Header
IP Header
Ethernet Frame Header
Frame X to Router 1 (A1-BB-DE-19-C3-4F)
Destination Data Link Address A1-BB-DE-19-C3-4F

FIGURE 1-26 Frame and IP Header Showing Data Link (EUI-48) and IP Destination Addresses for the Frame Sent from Host A to the Router

Test Your Understanding

26. a) Are all data link addresses EUI-48 addresses? b) In which header are source and destination IP addresses found? c) In which header are source and destination data link addresses found? d) What kind of data link address do Ethernet networks use? e) What kind of data link address do Wi-Fi networks use? f) Why do hosts need two addresses?

INTERNET ROUTERS AND PERSONAL ACCESS ROUTERS

It is common for students to confuse Internet core routers, corporate access points, and residential access routers, the last of which contains both a trivial router and a consumer-grade access point. Figure 1-27 contrasts these three important but easy-to-confuse devices.

Internet Core Routers

The routers in the core (central part) of the Internet are designed to fit into standard equipment racks, but they are not merely 1U or 2U tall. **Internet core routers** range from the height of dorm room refrigerators to full-size refrigerators.[25] That isn't terrifically large, but they are powerhouses that can route high volumes of traffic and do complex routing to deliver packets along different routes to large numbers of destinations. They are also remotely manageable.

Residential Access Router

We call them **residential access routers**, but these little boxes are multifunction devices with surprising utility. They contain an Internet switch, a DHCP server, at least a simple firewall, and a limited consumer-grade access point. Routing is one of

[25] The routers at the edge of the Internet are smaller but vary considerably in size. A branch office router, for example, may indeed be only 1U or 2U tall.

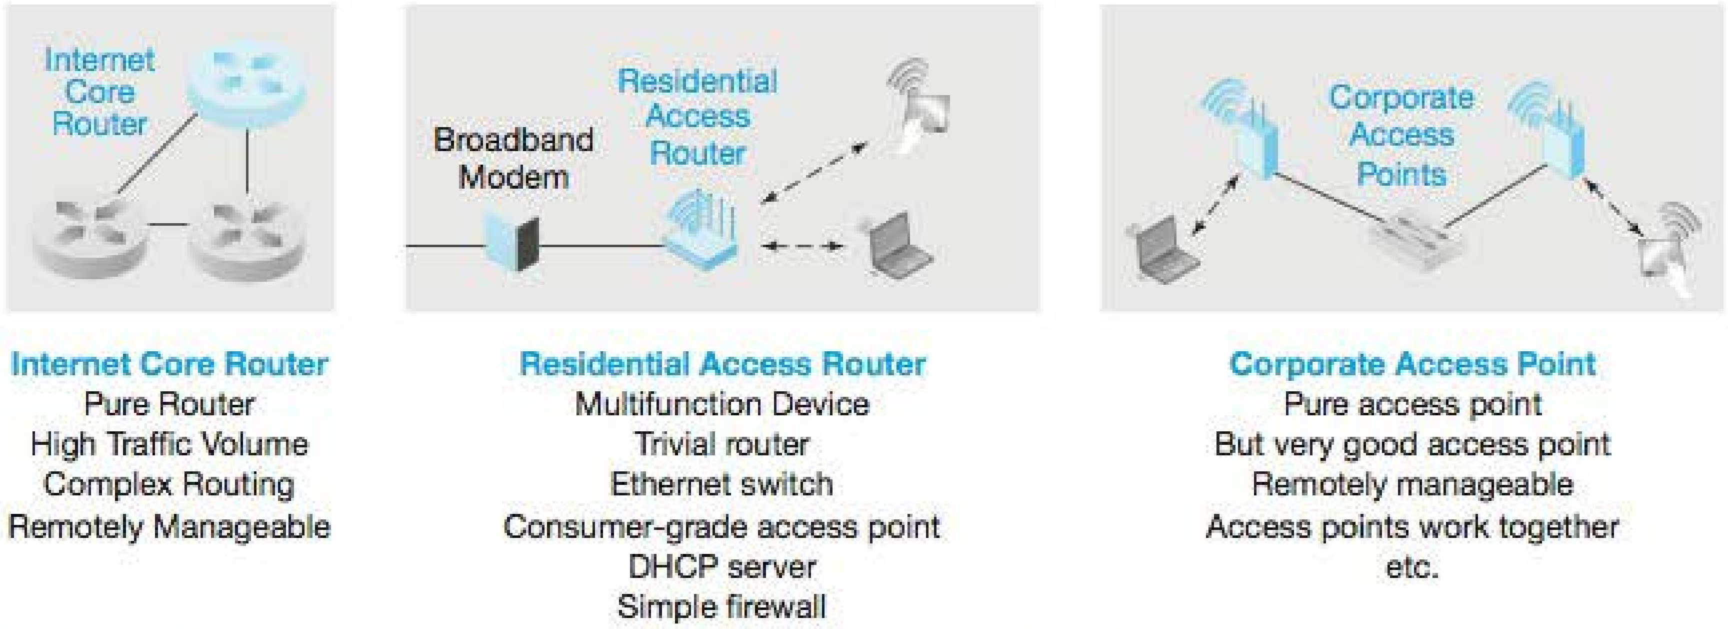

FIGURE 1-27 Internet Core Routers, Residential Access Routers, and Corporate Access Points

their functions, but it is the most trivial. Everything coming from the access router's connection to your devices is sent out to the Internet, and everything from the outside is sent inside.

Corporate Access Point

Corporate access points are usually smaller than home access routers, but they are pure access points, and they are very good *commercial-grade* access points with several capabilities that can be configured remotely. For instance, the network administrator can adjust the relative power of nearby access points to adjust for changing numbers of wireless devices throughout a section of a building. We will see in Chapter 7 that they also participate in collecting data for and implementing network security. Larger corporate access points have multiple radios and antennas, each of which can focus on a different direction, allowing more devices to share their area of service.

Test Your Understanding

27. a) Compare Internet core routers with home access routers in terms of functionality. b) Compare them in terms of routing complexity. c) Compare corporate access points and Internet access routers with wireless access point capabilities.

WHERE TO NEXT?

In this chapter, we looked at core Internet (and networking) concepts and principles.

- Chapter 2 looks in more depth at standards.
- Chapter 3 will teach you the core elements of network management. (These things do not manage themselves.)
- Chapter 4 looks at security tools and concepts. We do not put security off to the fourth chapter to indicate that it is unimportant. We do so because it is impossible

to discuss security until you understand the networking concepts and tools in the first three chapters. Your teacher may cover Appendix after Chapter 4. This deals with security management. This may be covered in other courses instead.

With this information about concepts, tools, and management, you will then apply your knowledge to specific standards and technologies.

- Chapter 5 takes a 360-degree view of Ethernet switched LAN standards, discussing them in terms of standards, technology, management, and security.
- Chapters 6 and 7 do the same for local wireless technologies, such as Wi-Fi.
- Chapters 8 and 9 look at the TCP/IP Internet standards in the same integrated way. You may wonder why we do not start with the Internet first. The answer is a pragmatic one. Ethernet and Wi-Fi are very familiar to you and are simpler than the Internet. Learning them first will give you a stronger set of base information to take on the Internet, which was designed from the ground up to be a full worldwide network.
- Chapter 10 takes you outside the local environment to discuss wide area networking. You might say, "Hey, isn't that the Internet?" To a large degree it is, but corporations cannot rely on the Internet completely because the Internet is only a best effort network, and corporations need to have tighter control over long-distance performance. In addition, there is the matter of accessing the ISP. You are probably generally familiar with mobile telephony, cable modems, and ADSL, and we will look at them in more depth. Corporations use another access technology for most of their connections; this is leased lines.
- Finally, Chapter 11 takes us to applications. Many teachers skip this chapter because their programs cover applications in other chapters. Chapter 11 generally looks at applications from a networking perspective rather than on the great things they can do. (You already know about that.)

END-OF-CHAPTER QUESTIONS

Thought Questions

1-1. In Figure 1-28, when Host A transmits a packet to Host B, how many physical links, data links, and routes will there be along the way? How many packets and frames? How many switches and routers? (Hint: The answers are in the figure, but work it out yourself.)

1-2. Repeat for Host C sending a packet to Host E.

1-3. Repeat for Host A to Host C.

1-4. Repeat for Host E and Router 3.

1-5. Repeat for Router 1 and Router 3.

1-6. Repeat for Router 1 and Router 2.

Perspective Questions

1-7. What was the most surprising thing you learned in this chapter?

1-8. What was the most difficult thing in this chapter for you? Why was it difficult?

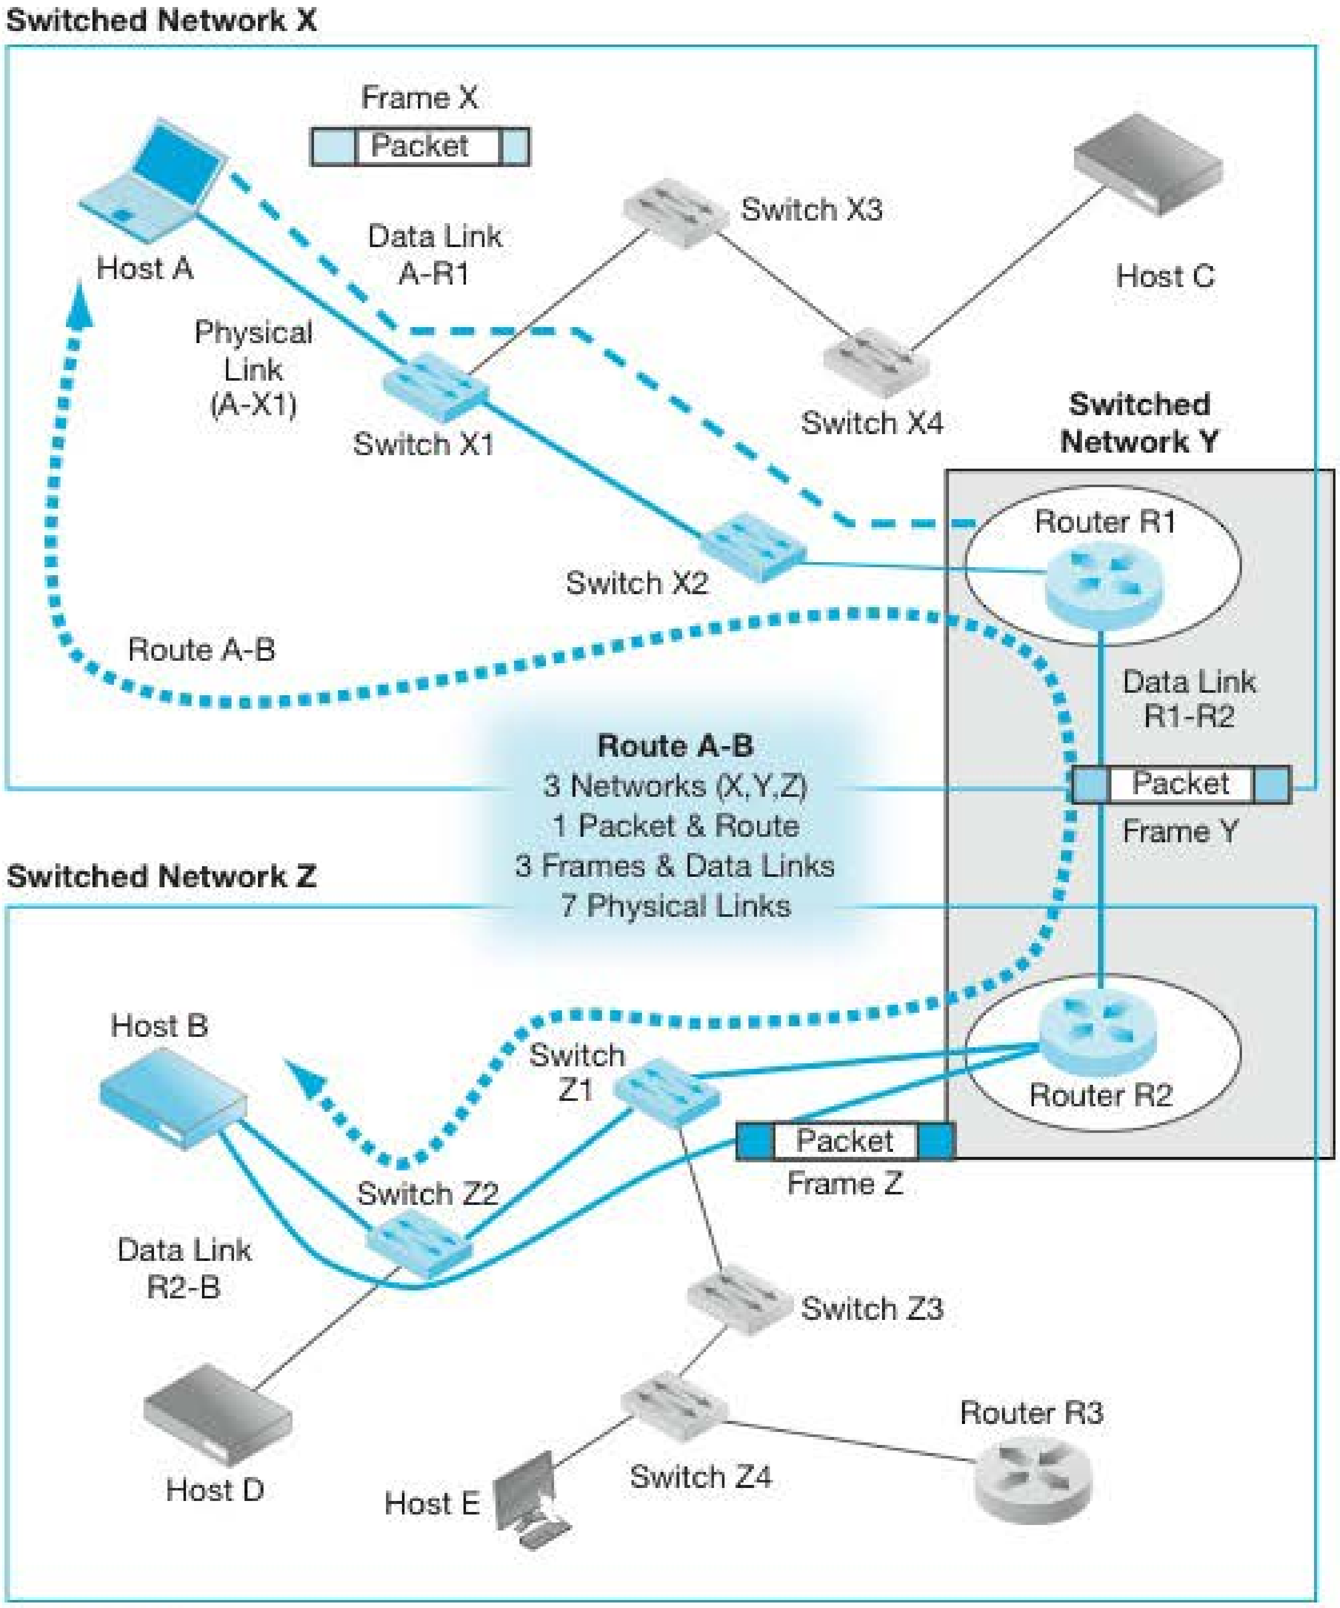

FIGURE 1-28 An Exercise in Physical, Links, Data Links, and Routes

Chapter **1a**

Hands-On: A Few Internet Tools

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Test your Internet connection speed.
- Look up a host's IP address by querying a DNS server.
- Use ping and traceroute to diagnose an Internet connection.

HANDS-ON EXERCISES

1. How fast is your Internet connection? See with one of the following websites. If you are asked to download a program for the test or run a program to see why your computer is running slowly, do not do so. Sites offering speed testing include www.zdnet.com/broadband-speedtest/, testymy.net, testinternetspeed.org, www.speedtest.net, and www.speakeasy.net/speedtest/. Report download speed, your upload speed, and your access technology (home DSL connection, school lab, 3G mobile phone, 4G mobile phone, etc.). Use two of these tools on one device or a single tool on two different devices.
2. Look up the IP address for panko.com. Tools for doing DNS lookups include ping.eu and networktools.com. If you are asked to download a program for the test or run a program to check your computer, do not do so. What result do you get?
3. Ping looks up whether an IP address or host name represents an active host and what delay there is in reaching the host. Traceroute is similar but shows all routers along the way. Tools for pinging and traceroute include ping.eu and networktools. com. If you are asked to download a program for the test or run a program to check your computer, do not do so. Try ping and traceroute for panko.com. If ping fails or if traceroute cannot get all the way to the host, a firewall may be prohibiting ping and traceroute. What results do you get?
4. Repeat the previous question for yahoo.com.

Chapter **2**

Network Standards

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Explain how Internet standards are made and why this approach is valuable.
- Provide the definitions of network standards and protocols; articulate their importance.
- Explain the OSI, TCP/IP, and Hybrid TCP/IP-OSI architectures and their standards agencies.
- Explain the purpose of each standards layer in the Hybrid TCP/IP-OSI architecture, what is standardized at each layer, and which standards agency dominates standards at each layer.
- Explain message ordering in general and in HTTP and TCP.
- Explain message syntax in general and in IP packets, TCP segments, UDP datagrams, and Ethernet frames.
- Demonstrate how application programs encode alphanumeric, decimal, and alternative data into bits (1s and 0s) before passing their messages to the transport layer.

HOW INTERNET STANDARDS COME TO BE

> *Those who love sausage and revere the law should never see either being made.*
>
> *Attributed to German Chancellor Otto von Bismarck*

Standards are detailed and precise. You might expect that standards creation would be orderly and precise as well. For most standards agencies, this is true. For the Internet, things are a little different.

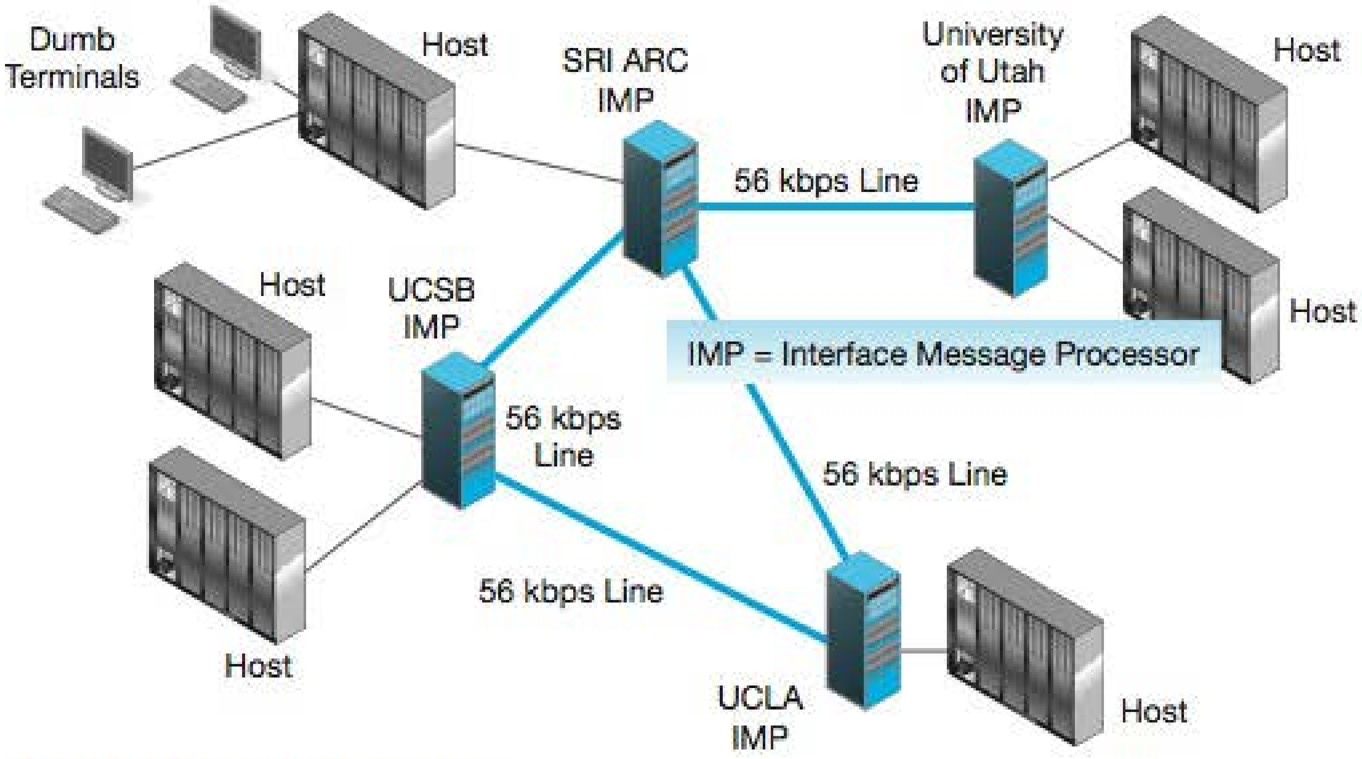

FIGURE 2-1 The Early ARPANET

The ARPANET The Internet grew out of the ARPANET research network funded by the Defense Advanced Research Projects Agency (DARPA).[1] DARPA funded it to explore the then-new technology of packet switching (what we would now call frame switching). Figure 2-1 shows that when the **ARPANET** began in 1969, it had four sites: UCLA, the Stanford Research Institute's Augmentation Research Center, UCSB, and the University of Utah. Each site had a switch called an interface message processor (IMP). IMPs exchanged packets (what we now call frames) through 56 kbps lines, which seemed blazingly fast at the time.

The Need for Standards Bolt, Beranek, and Newman (BBN) built the IMPs and designed protocols for IMPs to exchange messages (the blue lines and IMPs in Figure 2-1). That was all they did. At meetings during the ARPANET's development phase, researchers from the four sites met with BBN to discuss the network. They realized that the ARPANET would be useless without many additional standards. There had to be standards for hosts to communicate with their IMPs. Far more fundamentally, there had to be application standards if the network was to be useful.

The Network Working Group and Requests for Comments (RFCs) Knowing the importance of standards, and knowing that a vacuum existed in standards setting, the participants decided to do it themselves. They called their small team the **Network Working Group** and asked others to join them. When they came up with a

[1] Was it ARPA or DARPA? It depends on the year. It was born ARPA in 1958. In 1972, it became DARPA to emphasize its status as a Department of Defense agency. In 1993, it went back to ARPA. Then it went back to DARPA in 1996. Source: DARPA, "ARPA-DARPA: The Name Chronicles," undated, last viewed in August 2009. http://www.darpa.gov/arpa-darpa.html.

standard, they did not call it a standard because they felt that they lacked the authority to do so. Steve Crocker, who led the group and wrote the first document, called it a **Request for Comments (RFCs)**. Today, new standards are still RFCs, as are other types of documents.[2]

Group members quickly developed key application standards. In 1971, Ray Tomlinson realized that e-mail could work across sites. He was already working on e-mail for users of a single host. Mail systems on single hosts used usernames as addresses for delivering mail. Tomlinson saw that an ARPANET address would have to include both a username and the host name. Looking at his keyboard, he saw that the @ sign did not seem to be used very much.[3] He assigned it to separate the username from the host name. (The first author was Ra3y@Office1.) It took him a weekend to write the software. E-mail quickly dominated use of the ARPANET.

Internet Engineering Task Force Born in the late 1960s, the Network Working Group reflected its times. There was a strong focus on egalitarian participation and the recognition of technical merit. A few years later, the **Internet Engineering Task Force (IETF)** took over Internet standards development. Like the Network Working Group, the IETF has no formal membership. Anyone can participate in the IETF Working Groups that develop individual standards in specific areas.

Describing how the IETF works, Dave Clark wrote, "We reject: kings, presidents, and voting. We believe in: rough consensus and running code."[4] Rejecting kings and presidents refers to the IETF's strong egalitarian culture. In general, anyone with a good idea stands a fair chance of being heard. By not suppressing new ideas, this culture accounts for much of the rapid development pace of Internet standards. The rejection of voting and going forward if there was rough consensus also made the IETF action-oriented.[5]

The importance of "running code" is not as obvious. Most standards agencies develop full standards before devices and software are ever built. When vendors implement these standards, they often find unforeseen ambiguities and even contradictions. When they build their products to these standards, they often find that their products do not work with products from different vendors who supposedly follow the same standard. In addition, committees tend to design standards that are so complex that products take extensive resources to develop and are therefore expensive and slow to

[2] All standards are RFCs, but not all RFCs are standards. Even for standards-track RFCs, there are proposed standards, draft standards, and Internet standards. Only RFCs that are Internet standards are official standards. RFCs also can be listed as best current practice, informational, experimental, historical, and even unknown. How do you know which RFCs are current Internet standards? The IETF occasionally publishes an RFC that lists them. Wikipedia has a listing as well, although it should not be accepted as definitive without the official list in the relevant RFC.

[3] Personal communication with Ray Tomlinson, May 1986.

[4] Dave Clark, "A Cloudy Crystal Ball—Visions of the Future," in *Proceedings of the Twenty-Fourth Internet Engineering Task Force* (Cambridge, Mass.: Massachusetts Institute of Technology NEARnet, July 13–17, 1992), 539–43.

[5] At meetings, the audience is asked to hum on agreement. Humming allows more anonymity than voice or hand voting, and it is probably less precise. Unlike traditional voting, the item being hummed is dropped or sent back for more work unless there is strong consensus for going further.

IN MORE DEPTH

April 1 and RFCs

The IETF has a sense of whimsy. In the United States and some other countries, April 1 is April Fool's Day—a day to play jokes on people by telling them something completely false. A robust tradition in the IETF is the publishing of a facetious RFC or two on April Fool's Days. One of the most popular is RFC 2549, IP over Avian Carriers. Written in 1990, this RFC describes how to transmit IP packets using carrier pigeons. This RFC was updated twice, in 1999 (to add quality of service) and in 2011 (so that the protocol will work with the new IPv6 protocol). Another April 1 RFC warned of a serious authentication problem at IETF meetings. There were so many heavily bearded guys that it was impossible to tell them apart. RFC 3093 introduced the Firewall Enhancement Protocol, which allows all traffic to pass through firewalls while leaving the firewall in place (and useless). An April 1 RFC from 1998, the Hyper Text Coffee Pot Control Protocol, was created as RFC 2324. In justifying the HTCPCP, the RFC said that "there is a strong, dark, rich requirement for a protocol designed espressoly for the brewing of coffee." One limitation in the protocol was that decaf coffee was explicitly excluded. The explanation was, "What's the point?" Although this RFC was a joke, a serious protocol on remote brewing will almost certainly be seen as the Internet of Things unfolds.

come to market. In the IETF, almost all standards are created based on running demonstration systems. Experience identifies unforeseen problems and solves them before standards are made.

More subtly, demonstration code is simple. This leads to *simple standards*. Many IETF RFCs even have "simple" in their name; for instance, the Simple Mail Transfer Protocol standardizes communication among mail servers. We will see in Chapter 3 that the Simple Network Management Protocol is now the core tool for remotely managing network resources. Simple products emerge quickly, so while OSI development plodded along slowly, simple TCP/IP products appeared fast, at low prices. As something of an insult (although it was not intended to be), the IETF sometimes took bloated OSI standards and created simpler versions of them. These simplified versions often became dominant. Over time, simple IETF standards usually evolve to becoming full-featured, but each step along the way is based on real-world experience.

Test Your Understanding

1. a) What are IETF standards called? (Spell out the name and give the acronym.) b) What factors in the Internet's informal development process lead to rapid standards development and low-cost products?

INTRODUCTION

In Chapter 1, you saw a handful of standards. In the rest of this book, you will see many more. Fortunately, if you master some core standards concepts, you will be able to see a new standard and immediately understand a lot about it. (If you do not master these concepts, the standards you see will become confusing masses of detail.) This chapter covers these core standards concepts.

Standard = Protocol

In this book, we use the terms *standard* and **protocol** *to mean the same thing*. In fact, standards often have *protocol* in their names. Important examples are the Hypertext Transfer Protocol that governs communication on the World Wide Web, the Internet Protocol, the Transmission Control Protocol, and the User Datagram Protocol.

In this text, we use the terms standard and protocol to mean the same thing.

What Are Network Standards?

Network standards are rules of operation that specify how two hardware or software processes work together by exchanging messages. As Figure 2-2 illustrates, network standards govern the exchange of messages between two hardware or software processes. To give a human analogy, in the authors' classes, the standard language is American English. Not all of these students are native English speakers, but we are able to communicate because we use a standard language.

Network standards are rules of operation that specify how two hardware or software processes work together by exchanging messages.

The Importance of Standards

Figure 2-2 notes that network standards allow products from different vendors to **interoperate** (work together effectively). The client program is from Apple, and the server program is from Microsoft. These companies often dislike each other, but their products work together because they exchange messages using the Hypertext Transfer Protocol (HTTP) network standard.

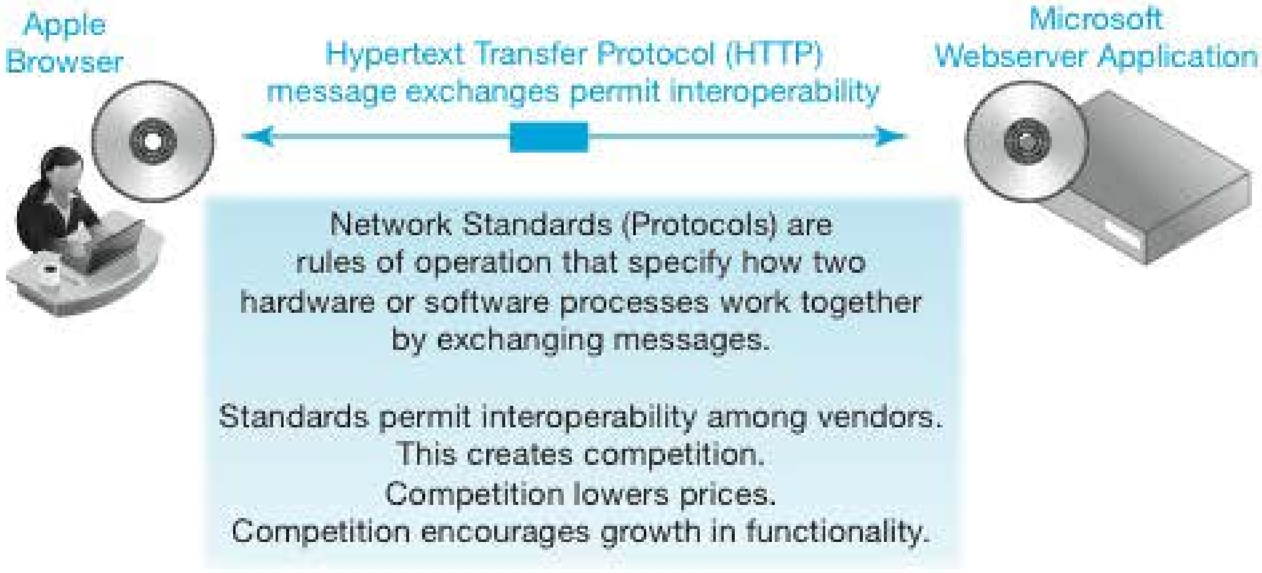
Apple
Browser
Hypertext Transfer Protocol (HTTP)
message exchanges permit interoperability
Microsoft
Webserver Application

Network Standards (Protocols) are rules of operation that specify how two hardware or software processes work together by exchanging messages.

Standards permit interoperability among vendors.
This creates competition.
Competition lowers prices.
Competition encourages growth in functionality.

FIGURE 2-2 Network Standards

To interoperate is to work together effectively.

Standards are important for three reasons.

- Standards increase competition. With network standards, it is impossible for any company to maintain a monopoly by closing out competitors.
- With no monopolies, competition drives down prices.
- Standards also spur companies to add new features to their products. Adding new features prevents their products from being undifferentiated *commodities* that can only compete on price. These new features often appear added to the next version of the standard, requiring a new round of innovations to create competitive advantages.

Network standards are the key to networking in general. To work in networking, you need to understand individual standards so that you can design networks, set up network components, and troubleshoot problems. Learning networking is heavily about learning standards.

Test Your Understanding

2. a) Distinguish between standards and protocols. b) What is a network standard? c) What is interoperability? d) What are the benefits of standards?

CREATING STANDARDS

Standards are developed by **standards agencies**. At the beginning of this chapter, we looked briefly at one important standards agency, the IETF. Now we look more broadly at standards agencies and their standards architectures, including the hybrid TCP/IP-OSI standards architecture that most organizations actually use today.

Standards agencies are organizations that create standards.

Standards Agencies

It would make things simpler if there were only a single standards agency in networking, but there are many. Two are broadly important.

- Again, Internet standards come from the Internet Engineering Task Force (IETF). These standards are used especially by internet processes, transport processes, and Internet supervisory standards.
- There is also another important pair of standards agencies, the **International Organization for Standardization (ISO)**[6] and the **International Telecommunications**

[6] No, the standard acronym and the standard name in English do not match. In fact, "ISO" does not translate into the organization's standard name in any language. ISO is based on the Greek word for *true*. ISO separately standardizes its name in every language. Try not to think about this too much. It will hurt your head. A lot of things do that in standards.

Union–Telecommunications Standards Sector (ITU-T).[7] Their collaboration began well before the IETF started. ISO and ITU-T create a variety of network standards, especially for physical and data link processes. We will use the abbreviations for these two organizations in this book.

Test Your Understanding

3. a) What standards agency creates Internet standards? b) What other two standards agencies work together to create network standards? c) Which standards agency(ies) is(are) especially important for internet processes? d) For physical transmission processes? e) For data link processes? f) For transport processes? g) For Internet supervisory processes?

Standards Architectures

When we are faced with big jobs, we naturally break them into smaller pieces that will collectively get the job done. We then assign individual parts to people with the most relevant skills.

Similarly, a standards agency begins its work by creating a standards architecture. **Standards architectures** specify everything needed for two different programs on two different hosts on different single networks to interoperate. Standards architectures are collectively exhaustive.[8]

Standards architectures specify everything needed for two applications on two hosts on different single networks to interoperate.

Layering In network standards architecture, the overall architecture is divided into **layers**. Collectively, the layers in a standards architecture specify everything that must be standardized for two different application programs on two different hosts on two different networks to interoperate.

Internet Engineering Task Force (IETF)	ISO and ITU-T
Standards for the Internet, especially internet processes, transport processes, and Internet supervisory standards	A variety of network standards, especially for physical and data link processes

FIGURE 2-3 Major Standards Agencies in Networking (Study Figure)

[7] No, the name and abbreviation do not make sense. Again.

[8] Standards architectures are created within a year or two after a standards agency forms. From then on, the job is to create a series of individual standards at each layer.

Collectively, the layers in a standards architecture specify everything that must be standardized for two different application programs on two different hosts on two different networks to interoperate.

Each layer provides services to the next-higher layer. Consider an analogy, driving between two locations. The lowest layer is the road. It provides services to the next higher layer, the wheels. Specifically, the road provides a supportive and adequately smooth surface for the wheels to work on. The wheels, in turn, support the car's body. The body supports the driver (Figure 2-4).

Each layer in a standards architecture provides services to the next-higher-layer.

Specialization in Design Layering permits specialization in design. Road designers do not have to worry about tires, car bodies, or drivers (at least at a low level). Instead, they can focus on soil analysis, the strength of paving materials, and things of that ilk. Wheel designers, in turn, can specialize in wheel tensile strength, wear for different tire compounds, and similar things.

Changing a Single Layer If layering is done well, a change can be made at one layer without requiring other layers to change. For instance, if a car is given auto parking ability, the driver can ignore it, and there is no impact on wheel or road standards. However, if one layer is improved, the layer above it can be improved if desired. In this example, the driver can decide to do auto park. Because changes can be made in different layers at different times, there is no need to change everything every time there is a change at one layer.

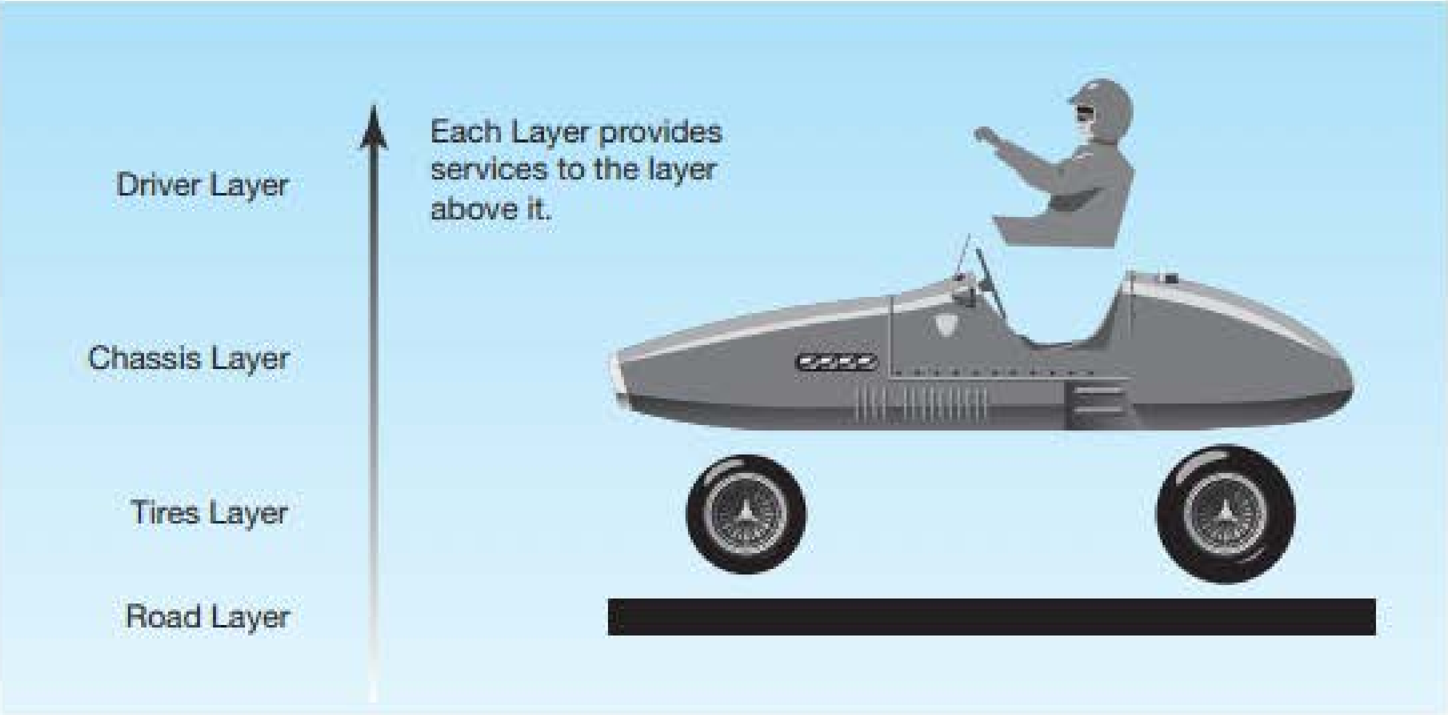

FIGURE 2-4 Layering in Automobile Travel

Specialization in Design

For the road layer, soil analysis, strength of paving materials, etc.

For the wheels layer, tensile strength, wear for different compounds, etc.

The Ability to Change One Layer While Not Changing Others

If add auto parking at the car body level, need not adopt it at the driver level

However, the driver layer can change to take advantage of it if desired

Upgrade layers as desired

It would be too expensive to upgrade all standards every time a standard changed

FIGURE 2-5 Layering Benefits (Study Figure)

Test Your Understanding

4. a) Why do standards architectures have multiple layers? b) To what does a standards layer provide services? c) If you change a standard at one layer, do standards at other layers need to be changed? d) Why may it be advantageous to change a standard if the standard at the layer below it is upgraded?

The OSI Standards Architecture

As Figure 2-6 shows, different standards agencies have different standards architectures. For example, ISO and ITU-T created the **Reference Model of Open Systems Interconnection**. Thankfully, this nearly unpronounceable name is always shortened to **OSI**. Not thankfully, OSI the architecture is easy to confuse with ISO the organization.[9] The first column shows that the OSI architecture has seven layers.

OSI Architecture (ITU-T and ISO) Layer Number and Name	TCP/IP Architecture (IETF) Layer Number and Name	Hybrid TCP/IP–OSI Standards Architecture Layer Number and Name	Standards Come Predominantly From
7. Application	4. Application	5. Application	Various Standards Architectures
6. Presentation			
5. Session			
4. Transport	3. Transport	4. Transport	TCP/IP (IETF)
3. Network	2. Internet	3. Internet	TCP/IP (IETF)
2. Data Link	1. Subnet Access Protocol	2. Data Link	OSI (ITU-T and ISO)
1. Physical		1. Physical	OSI (ITU-T and ISO)

FIGURE 2-6 Standards Agencies and Layered Standards Architectures

[9] I believe that this was done deliberately to confuse students.

OSI the architecture is easy to confuse with ISO the organization.

Note that the bottom two OSI layers have names that should be familiar to you. The physical layer is for standards that deal with physical processes—transmission media, connectors, and signaling. The data link layer standardizes data link processes—frames, switches, wireless access points, and data links.

By the late 1970s, quite a few OSI standards at these two layers were solid. In general, OSI standards quickly dominated at the physical and data link layers. Above the data link layer, however, ISO and ITU-T ran into trouble. They did not have clear ideas about internetworking and took time to develop their understanding in this area. This left the door open for the architecture created by the IETF.

The TCP/IP Standards Architecture

In contrast, the IETF began with a laser focus on internetworking. It knew that internet and transport layer standards were needed to build the Internet. The IETF standards architecture is named after its two main initial standards. It is called **TCP/IP.**[10] This makes sense, but it can cause some confusion. TCP/IP is the architecture, and TCP and IP are standards within the architecture.

TCP/IP is the architecture, and TCP and IP are standards within the architecture.

Good single-network standards were already available from ISO and ITU-T, so the IETF simply decided to use them.[11] For the physical and data link layers, the TCP/IP architecture specifies the Subnet Access Protocol (SNAP). This basically says, "Use OSI standards here." **Subnet** is the IETF's name for a single network. Its job is to create a data link between hosts and routers and between routers and other routers.

Above the Subnet Access Protocol is the **internet layer**. This is for internet processes, including packets, routers, and routes. The transport layer, then, is for transport processes. This includes application message fragmentation and reassembly. As we will see in this chapter, transport layer standards also do error correction.

When Do We Capitalize "Internet?"

When we refer to the global Internet, we capitalize the name. However, we do not capitalize internet when referring to the internet layer or when we refer to an internet other than the global internet. Yes, there are some.

[10] The IETF architecture also has an official name. However, it is almost never used. It is kind of like Voldemort.

[11] Occasionally, the IETF creates data link layer standards. Most notably, it created the point-to-point protocol that we saw in Chapter 1 and also will see later in this chapter. Point-to-point OSI standards did not have all of the functionality needed to directly connect two routers.

When we refer to the global Internet, we capitalize the name. However, we do not capitalize internet when referring to the internet layer or when we refer to an internet other than the global internet.

The Hybrid TCP/IP–OSI Standards Architecture

Real organizations care nothing about architectural purity. They just want to get their work done. As Figure 2-6 shows, what most firms *actually* use is a **Hybrid TCP/IP–OSI Architecture**, which combines OSI standards at the physical and data link layers with TCP/IP standards at the internet and transport layers.

The hybrid TCP/IP–OSI Architecture combines OSI standards at the physical and data link layers with TCP/IP standards at the internet and transport layers and (usually) standards from any architecture at the application layer.

Application Standards At the application layer, things are messier. No standards agency dominates at this layer, although both IETF and OSI standards are popular at this layer. Adding to the confusion at this layer, the IETF and ISO frequently work together to create application layer standards.

Fortunately, it normally does not matter what standards agencies create application protocols. Most application layer standards can work with IETF standards at the transport layer. Consequently, companies that use the Hybrid TCP/IP–OSI architecture have no problem using applications from different standards agencies and architectures.

Most application layer standards can work with IETF standards at the transport layer. Consequently, companies that use the Hybrid TCP/IP–OSI architecture have no problem using applications from different standards agencies.

Test Your Understanding

5. a) What are the standards agencies for OSI? Just give the abbreviations. b) Distinguish between ISO and OSI. c) What is the standards agency for TCP/IP? (Give both the name and the abbreviation.) d) What standards architecture do most organizations actually use in practice? e) At which layers of this architecture are IETF standards dominant? f) At which layers are ISO and ITU-T standards dominant? g) Why does it usually not matter what standards agency creates an application layer standard?

The Five Layers Figure 2-7 recaps the five layers of the Hybrid TCP/IP–OSI architecture. The first column looks at standards more broadly, grouping them into three broad functions. These are application program interoperability, transmission across an internet, and transmission across a single network.

Broad Function	Layer	Name	Specific Function
Interoperability of application programs	5	Application	Application layer standards govern how two applications work with each other, even if they are from different vendors.
Transmission across an internet	4	Transport	Transport layer standards govern aspects of end-to-end communication between two end hosts that are not handled by the internet layer, including reliability and application message fragmentation. These standards allow hosts to work together even if the two computers are from different vendors or have different internal designs.
	3	Internet	Internet link layer standards govern the transmission of packets across an internet—typically by sending them through several routers along a route. Hosts and routers can be from different vendors. Internet layer standards govern packet organization and routing.
Transmission across a single network	2	Data Link	Data link layer standards govern the transmission of frames across a single switched network—typically by sending them through several switches along the data link. Data link layer standards also govern frame organization, timing constraints, and reliability. As in all other layers, the devices can come from different vendors.
	1	Physical	Physical layer standards govern transmission between adjacent devices connected by a transmission medium, regardless of who the two vendors are.

FIGURE 2-7 Layers Recap

Test Your Understanding

6. a) What layer or layers govern(s) transmission media? b) Application programs? c) Transmission through a single network? d) Transmission through the Internet? e) Application message fragmentation?

Repeated Concepts at Layers 2 and 3 A common source of confusion is that concepts are repeated at the data link and internet layers but with different terminology. This occurs because internetworking required the creation of a second layer of forwarding standards to those used for transmission through single networks. Figure 2-8 shows how terminology differs between the data link and internet layers.

Packets Are Carried Inside Frames Recall that packets are carried inside frames. When a source host sends a packet to a destination host, the packet travels within a frame in each network along the way. If there are 19 single networks on the route between the source and destination hosts, a single packet will travel in 19 different frames.

	Layer 2	Layer 3
Layer Name	Data Link	Internet
Message	Frame	Packet
Forwarding Device	Switch	Router
Forward Occurs Within	A Single Network	The Internet as a whole
Path of Messages' Travel	Data Link	Route
Destination Address in Header	Data Link Layer (DLL) address; often, but not always EUI-48 addresses	IP addresses

FIGURE 2-8 Repeated Concepts and Different Terminology in Layers 2 and 3

Test Your Understanding

7. a) At what layer will you find standards for routers? b) Wireless access points? c) Packets? d) Switches? e) Frames? f) IP addresses? g) Routes? h) EUI-48 addresses? i) Data links?
8. a) If two hosts are connected by five networks, how many packets will there be when one host sends a packet to the other host? (Hint: Draw a picture.) b) How many frames? c) How many routers? d) If every host and router connects with a point-to-point connection, how many physical links will there be?

MESSAGE ORDERING (PLUS RELIABILITY AND CONNECTION ORIENTATION) IN STANDARDS

One thing that standards govern is **message order**, which is a fancy way of saying that they govern when each of the two processes may transmit messages. Writing programs on different machines that must work together is nearly impossible without firm control over the order in which processes may send messages. (In classes, you may not talk any time you want.)

Simple Message Ordering in HTTP

Figure 2-9 illustrates an HTTP request–response cycle. The client sends a request, and the server sends a response. The cycle is *always* initiated by the client, *never* by the server. The server cannot transmit unless the client has sent it an HTTP request message. This is a very simple type of message ordering.

In an HTTP request–response cycle, the cycle is always initiated by the client, never by the server.

Although HTTP message order is very simple, there are two things to note. Networking professionals categorize it as a **connectionless protocol**. This means that there

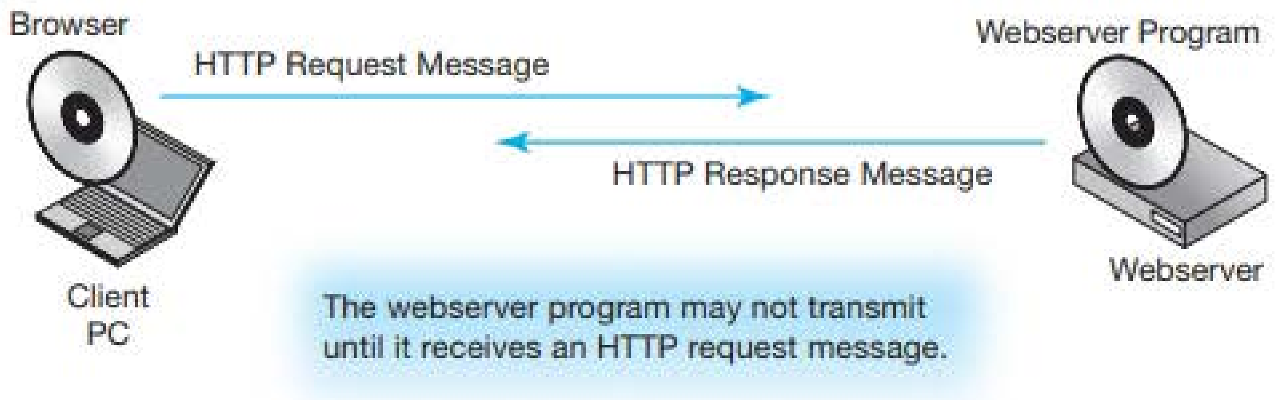

FIGURE 2-9 Simple HTTP Request–Response Cycle That Is Connectionless and Unreliable

is no need to have some sort of live connection before transmitting. The client may send a packet any time it wishes. In addition, HTTP is **unreliable**. There is no provision for the retransmission of lost or damaged messages. It is like sending a text message.

Test Your Understanding

9. a) In HTTP, which application program initiates an interaction? b) Is HTTP a connectionless protocol? c) Is HTTP a reliable protocol?

Message Ordering and Reliability in TCP at the Transport Layer

Many protocols have much more complex rules for message ordering. We look at Transmission Control Protocol (TCP) at the transport layer to see an example of this complexity.

Connections Figure 2-10 shows the transport layer processes on Host A and Host B. They are communicating via HTTP at the application layer. The Hypertext Transfer Protocol requires the use of TCP at the transport layer. The figure shows a sample communication session, which is called a **connection** because before the two sides begin to communicate, they first agree that they will communicate. At the end, they formally stop communicating. This is like talking on a phone. At the beginning, there is at least a tacit agreement that both sides are willing to talk. At the end of a telephone call, both sides usually agree to end the conversation. (Just hanging up is considered rude.) Technically, we say that TCP is a **connection-oriented protocol**.

In a connection-oriented protocol, the two sides first agree that they will communicate and formally stop communicating at the end.

TCP Segments In TCP, messages are called **TCP segments** because each carries a segment (fragment) of an application message if the message is long. We will see that it can also be a control segment that does not carry application data.

TCP messages are called TCP segments.

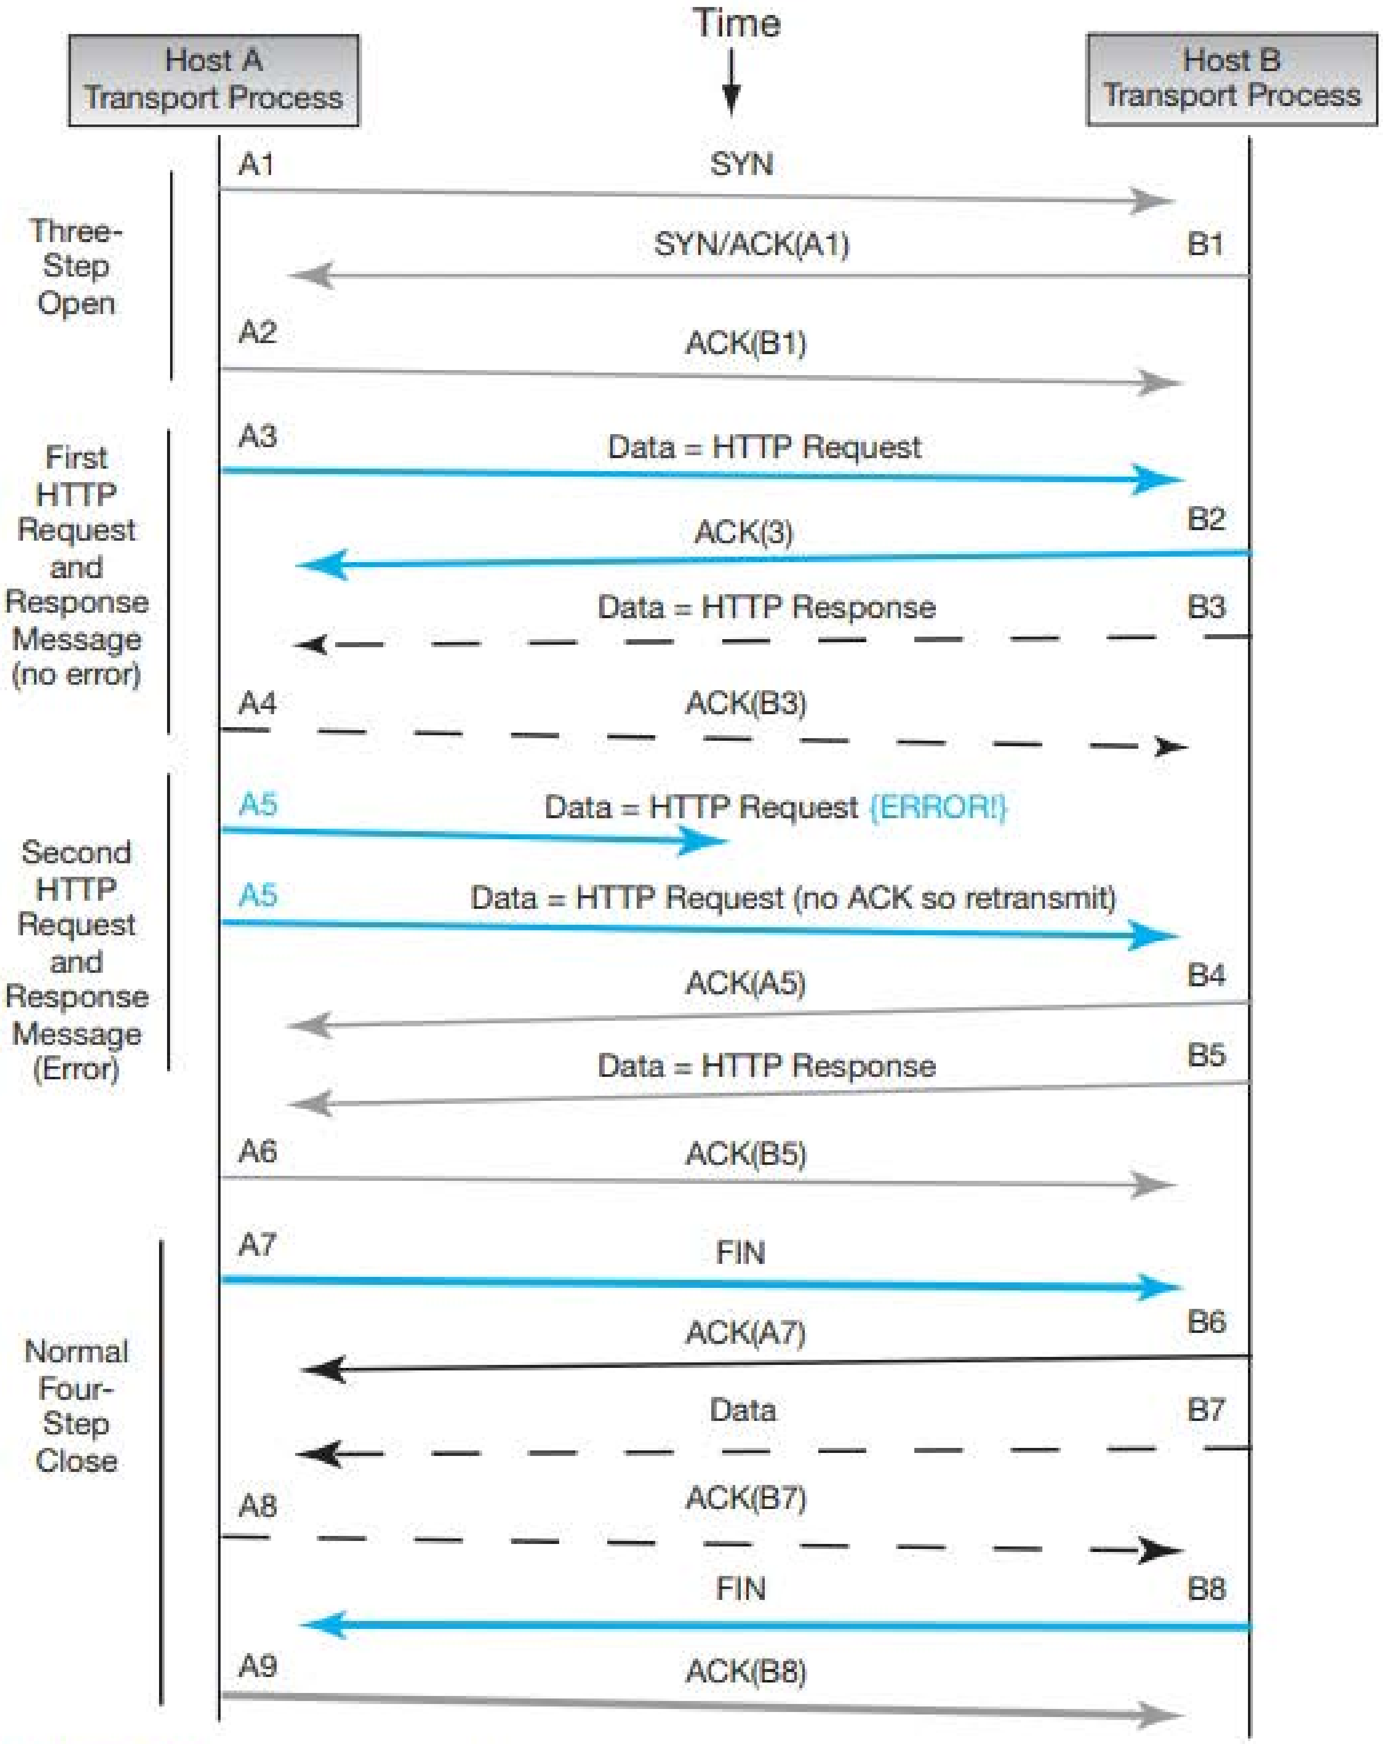

FIGURE 2-10 More Complex TCP Session with a Connection and Reliability

The Three-Step Opening The communication begins with a three-step opening handshake to establish a connection.

- Host A, which is the client in the HTTP exchange, initiates the communication. It transmits a TCP SYN (synchronization) segment to Host B. This indicates that Host A wishes to begin a connection.
- Host B sends back a TCP SYN/ACK segment. The SYN indicates that it also is willing to begin the communication. The ACK part is an acknowledgment of

Host A's SYN message (A1). In TCP, all segments are acknowledged, with the primary exception of pure ACKs. (If pure ACKs had to be acknowledged, there would be an endless series of ACKs.)

- Host A sends back a pure TCP ACK segment. This acknowledges Host B's SYN/ACK segment.

In TCP, all segments are acknowledged, with the primary exception of pure ACKs.

Control Segments These three TCP segments are noteworthy because they do not contain data. They are pure TCP headers, as we will see later. They are control segments.

Test Your Understanding

10. a) What do we call TCP messages? b) Describe the three-step opening in TCP. c) Is every TCP segment acknowledged? d) What is noteworthy about control segments?

Sequence Numbers In a connection-oriented protocol, each message is given a sequence number, which specifies the order in which it was sent. This allows the receiver to ensure that no message is missing and allows the receiving process to deal with duplicate segments. (It simply discards duplicates.)

In a connection-oriented protocol, each message is given a sequence number, which specifies the order in which it was sent.

Sequence numbers in TCP are important because application message fragments are delivered in separate packets. Sequence numbers allow the receiver to place the segments in order and reassemble them.

Note in Figure 2-10 that each side numbers its own sequence numbers. For simplicity, we have called Host A's sequence numbers A1, A2, A3, and so forth. We have done the same with Host B's messages. So Host A's SYN segment is A1, Host B's SYN/ACK is B1, and Host A's acknowledgment of the SYN/ACK is A2.[12]

Test Your Understanding

11. a) Is TCP connection-oriented or connectionless? b) What benefits do sequence numbers bring? c) How many segments does each side transmit?

Carrying Application Data The next four segments (A3, B2, B3, and A4) constitute an HTTP request–response cycle.

[12] Actually, sequence numbers increase with successive segments but in a complex way instead of increasing by one each time.

- A3 carries an HTTP request message.
- B2 is an ACK of A3.
- B3 carries the HTTP response message.
- A4 acknowledges the receipt of B3.

Usually, HTTP request messages are small enough to fit in a single TCP segment. However, most HTTP responses contain files that must be segmented and sent in a number of TCP segments. This does not change the basic picture, however. There would simply be several more exchanges like B3 and A4.

Reliability TCP is a reliable protocol. This means that it corrects errors. The second HTTP request–response cycle demonstrates how HTTP handles an error.

A reliable protocol corrects errors.

- Segment A5 is sent, but it never reaches Host B. An error has occurred.
- Host B does not send an acknowledgment, because as just noted, ACKs are only sent when a segment is received correctly.
- Host A realizes that A5 has not been acknowledged. It retransmits A5. Note that it has the same sequence number as the TCP segment that had an error during transmission. This allows the receiver to put it in order in case other segments had been transmitted before the retransmission.
- This time, the segment arrives correctly at Host B. Host B sends B4, which is an acknowledgment of A5.
- Finally, Host B sends an HTTP response message (B5) and receives an ACK (A6). Again, sending an HTTP response message tends to take several TCP data/acknowledgment cycles.

In this example, Segment A5 never reached the receiving transport process. There would be no way to acknowledge it in this case. What would have happened if A5 had reached the transport process but was merely damaged during transmission? In this case, the receiving transport process would discard the segment. It would not send an ACK. Note that there is a simple rule for ACKs. Unless a transport process receives a segment correctly, it does not send an acknowledgment.

There is a simple rule for ACKs. Unless a transport process receives a segment correctly, it does not send an acknowledgment.

Test Your Understanding

12. a) What kind of message does the destination host send if it receives an error-free segment? b) What kind of message does the destination host send if it does not receive a segment during a TCP connection? c) What kind of message does the destination host send if it receives a segment that has an error during a TCP connection? d) Under what conditions will a source host TCP process retransmit a segment?

The Four-Step Closing Host A has no more HTTP request messages to send, so it formally begins a close for the connection.

- It does so by sending a FIN segment (A7), which Host B acknowledges (B6).
- This means that Host A will not send new data. However, it will continue to send ACKs to segments sent by Host B. A FIN segment is a control segment consisting only of a header. The FIN bit is set in the header.
- In this case, Host B does have one more data segment to send, B7. When it sends this segment, Host A's transport process responds with an ACK (A8).
- After that exchange, Host B is finished sending data. It sends its own FIN segment (B8) and receives an acknowledgment (A9).
- The connection is closed.

Test Your Understanding

13. a) What are the four steps in the four-way close? b) When the side that initiates the close sends its FIN segment, does it stop transmitting more TCP segments? Explain.

MESSAGE SYNTAX IN STANDARDS

We have just looked at message ordering. Now we will turn to message **syntax**, which is how messages are organized. Messages are simply long strings of bits (1s and 0s). Logically, however, messages have several components, and the receiving process needs to know what these components are and where they are located in the bit stream. To give you a feeling for message syntax, we will look at the syntax of three important message types: IP packets, TCP segments, and UDP datagrams.

Syntax describes how messages are organized.

Syntax: General Message Organization

Before looking at the syntax of IP, TCP, and UDP messages, however, we need to look at syntax more generally. Figure 2-11 illustrates basic syntax elements.

Data Fields, Headers, and Trailers In Chapter 1, we saw several types of messages. In general, messages have three basic parts. The data field contains the information being delivered in the message. The definition of the header is simply everything that comes before the data field. Trailers? Everything that comes after the data field.

The header is everything that comes before the data field.

The trailer is everything that comes after the data field.

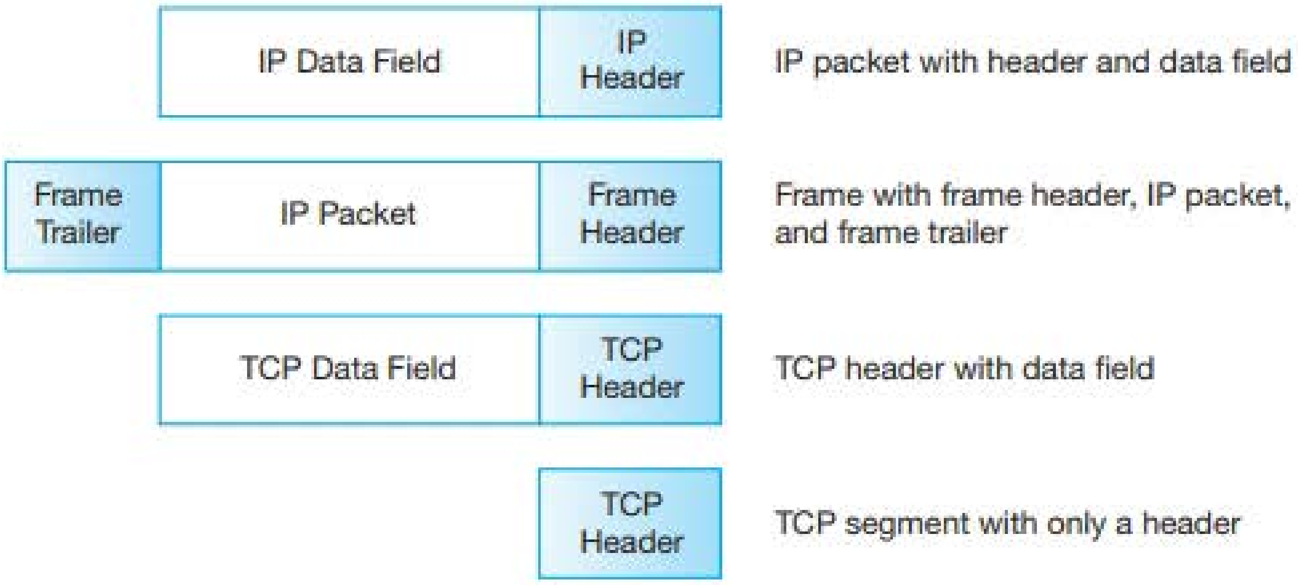
IP Data Field
IP Header
IP packet with header and data field
Frame Trailer
IP Packet
Frame Header
Frame with frame header, IP packet, and frame trailer
TCP Data Field
TCP Header
TCP header with data field
TCP Header
TCP segment with only a header

The header is defined as everything that comes before the data field.
The trailer is defined as everything that comes after the data field.
The header and trailer are divided into smaller parts called fields.

FIGURE 2-11 Headers, Data Fields, and Trailers

- We saw in Chapter 1 that an IP packet has a header and a data field.
- We saw that frames, in turn, often have a header, a data field, and a trailer. Frames are the only messages that typically have trailers, and not all of them have trailers.
- In Chapter 1, we also saw that TCP segments typically contain application message data in their data fields.
- We will see later in this section that some TCP segments are pure headers, with no data field. SYN, ACK, and FIN messages are examples of TCP segments without data fields. The supervisory information is contained entirely in the header.

Fields Headers and trailers are themselves divided into smaller parts called **fields**. In this section on message syntax, enumerating these fields and explaining some of them will be our focus in this section.

Headers and trailers are themselves divided into smaller parts called fields.

Test Your Understanding

14. a) What are the three general parts of messages? b) What does the data field contain? c) What is the definition of a header? d) Is there always a data field in a message? e) What is the definition of a trailer? f) Are trailers common? g) Distinguish between headers and header fields.

The Syntax of the Internet Protocol (IP) Packet

Having looked at message syntax in general, we now look at the syntax of a few individual standards. Figure 2-12 illustrates the syntax of an Internet Protocol (IP) version 4 (IPv4) packet.

32 Bits per Row An IP packet is a long string of bits (1s and 0s). Drawing the packet this way would require a page one line tall and several meters wide. Instead, Figure 2-12 shows that we usually depict an IP packet as a series of rows with 32 bits per row. This is the normal way to show syntax in TCP/IP standards, so you need to be familiar with it. In binary counting, the first bit is zero. Consequently, the first row shows bits 0 through 31. The next row shows bits 32 through 63.

Each row is subdivided into fields. For example, the first field is 4 bits long. This is the **Version Number Field**. In IPv4, it has the value 0100, which is the binary number for 4. As you might guess, value of this field in IPv6 is 0110, which is the binary number for 6.

Fields are distinct pieces of information in the bit stream of a message.

Source and Destination IP Address Fields We look at the IPv4 packet in more detail in Chapter 8, however, we note three fields in this chapter. Note that each IPv4 packet has a **Source IP Address Field** and **Destination IP Addresses Field** in the fourth and fifth rows. Each is 32 bits long, so each has a complete row. Routers use destination IP addresses to decide how to forward packets so that they will get closer to their destination.

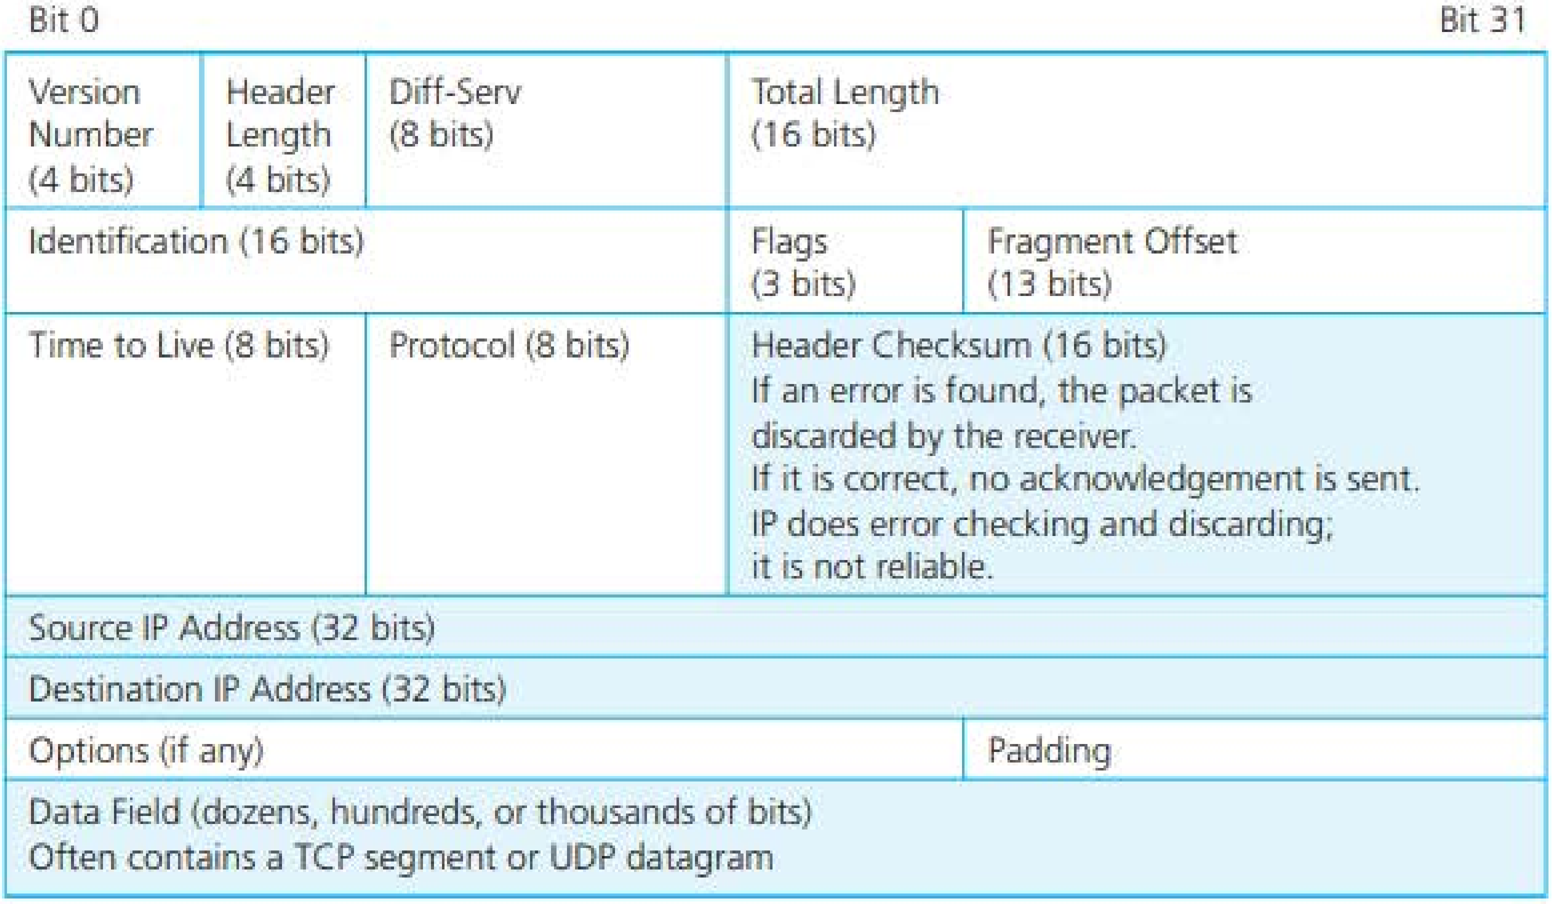

Bit 0 Bit 31

<table>
<tr><td>Version
Number
(4 bits)</td><td>Header
Length
(4 bits)</td><td>Diff-Serv
(8 bits)</td><td colspan="2">Total Length
(16 bits)</td></tr>
<tr><td colspan="3">Identification (16 bits)</td><td>Flags
(3 bits)</td><td>Fragment Offset
(13 bits)</td></tr>
<tr><td colspan="2">Time to Live (8 bits)</td><td>Protocol (8 bits)</td><td colspan="2">Header Checksum (16 bits)
If an error is found, the packet is
discarded by the receiver.
If it is correct, no acknowledgement is sent.
IP does error checking and discarding;
it is not reliable.</td></tr>
<tr><td colspan="5">Source IP Address (32 bits)</td></tr>
<tr><td colspan="5">Destination IP Address (32 bits)</td></tr>
<tr><td colspan="4">Options (if any)</td><td>Padding</td></tr>
<tr><td colspan="5">Data Field (dozens, hundreds, or thousands of bits)
Often contains a TCP segment or UDP datagram</td></tr>
</table>

FIGURE 2-12 The Internet Protocol (IP) Packet Syntax in IPv4

Unreliability The IPv4 **Header Checksum Field** is used for error detection. The sender computes a number based on all the other bits in the IP header. It places this value in the Header Checksum Field. The receiver redoes the calculation on the bits in the arriving IP packet header. If the numbers match, there have been no errors in transmission. The receiving internet process accepts the packet. If they do not match, then there has been an error. The receiver discards the packet.

Although the receiver checks for errors, it does not send an acknowledging packet if the packet is received correctly. The sending internet process has no way to know if the packet has been received correctly, so it cannot retransmit lost or damaged packets based on whether or not they have been received correctly. IP does error detection but not error correction. IP is an **unreliable protocol**.

IP does error detection and discards a packet containing an error. However, there is no retransmission of the lost message. IP is unreliable because reliability requires both error detection and error correction.

A Connectionless Protocol The Internet Protocol is a **connectionless protocol**. There is no need to formally agree to communicate or formally end the communication. It is like sending an e-mail. You just send it.

A connectionless protocol does not formally establish and then formally end communication sessions.

Test Your Understanding

15. a) List the first bit number on each IPv4 header row in Figure 2-12, not including options. (Remember that the first bit in Row 1 is Bit 0.) b) What is the bit number of the first bit in the Destination IP Address Field in IPv4? c) Describe how the internet process checks an arriving packet for errors. d) What does the receiving internet process do if it finds an error? e) What does it do if it does not find an error? f) Is IP reliable or unreliable? Explain. g) Is IP a connectionless or connection-oriented protocol?

Transmission Control Protocol (TCP) Segment Syntax

Earlier, we saw message ordering in the transmission of TCP segments. Now we will look at the syntax of TCP segments in a little more detail. We will see the rest of the TCP header syntax in Chapter 8. Most notably, this section describes how the TCP does what is necessary to be a reliable protocol.

Fields in TCP/IP Segments Figure 2-13 shows the organization of TCP segments. As in the case of IP packets, there are 32 bits on each line. This is the standard way in which the Internet Engineering Task Force shows syntax in its documents.

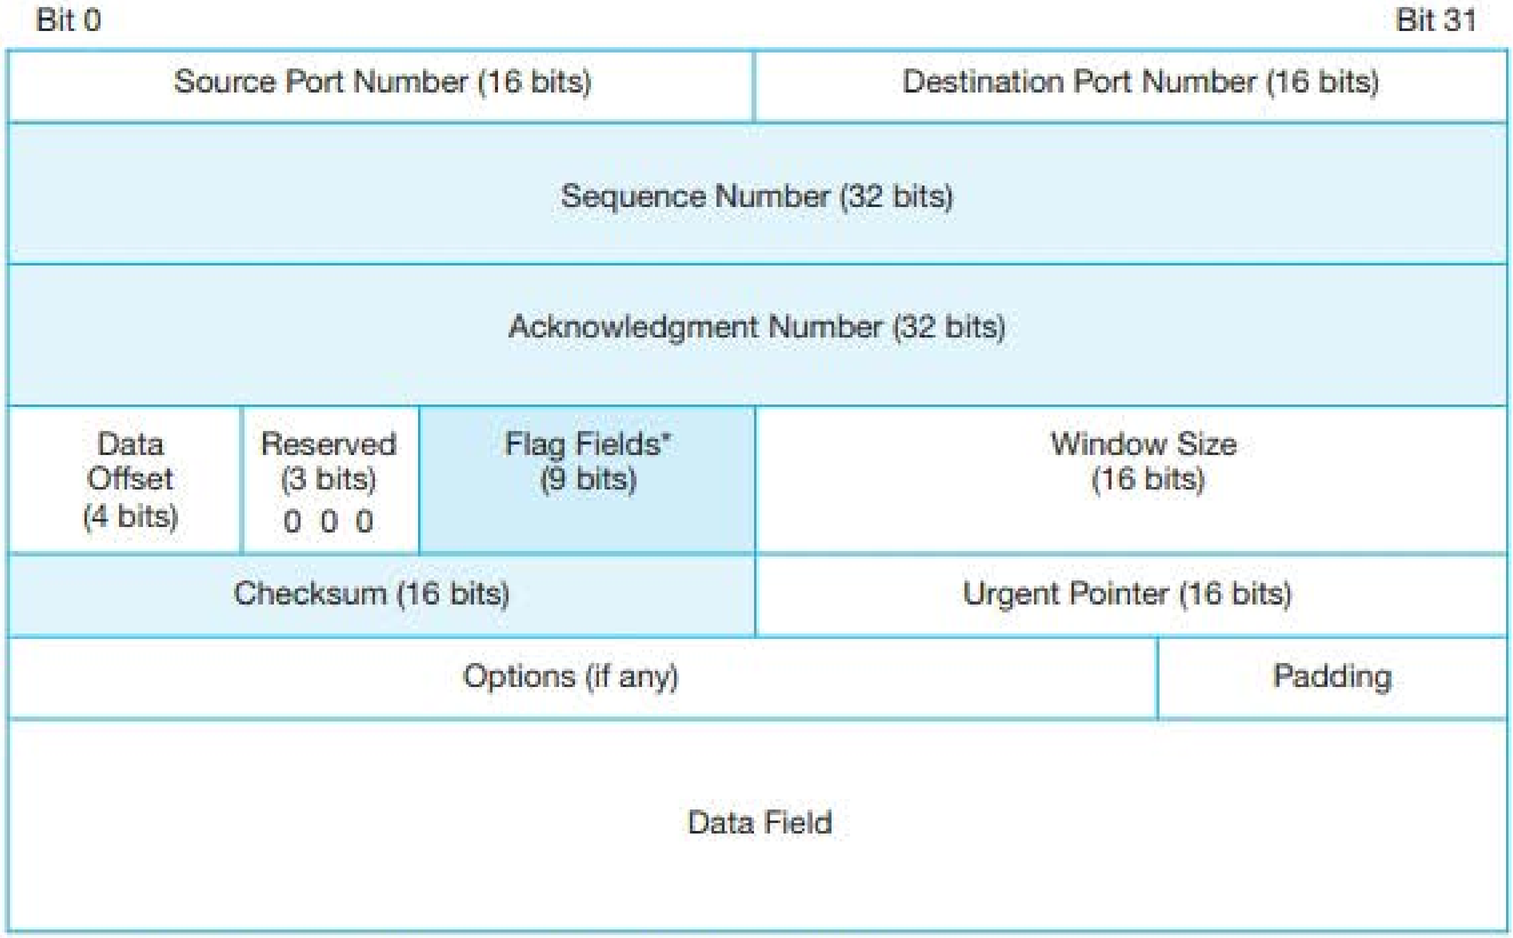

*Flag fields are 1-bit fields. They include SYN, ACK, and FIN bits.

FIGURE 2-13 TCP Segment

Flag Fields TCP has nine single-bit fields. Single-bit fields in general are called **flag fields**. If a flag field has the value 1, it is said to be *set*. If it has the value 0, it is said to be *not set*. In TCP, flag fields allow the receiving transport process to identify the kind of segment it is receiving. We will look at three of these flag bits:

- If the **ACK** (acknowledgment) bit is set (has the value 1), then the segment acknowledges another segment. When the ACK bit is set, the Acknowledgment Number Field also must be filled in to indicate which message is being acknowledged. If the ACK bit is not set, the TCP segment does not contain an acknowledgment. The receiver ignores the Acknowledgment Number Field.
- If the **SYN** (synchronization) bit is set, then the segment requests a connection opening.
- If the **FIN** (finish) bit is set, then the segment requests a normal connection closing.

Single-bit fields are called flag fields. If a flag field has the value 1, it is said to be set. (If it has the value 0, it is said to be not set.)

Earlier, we talked about TCP SYN segments, ACK segments, and FIN segments. These are simply segments in which the SYN, ACK, or FIN bits in the header are set, respectively. SYN and FIN segments have no data fields. ACK segments sometimes have no data fields.

An ACK segment is one in which the ACK bit is set (has the value 1).

Sequence Numbers Earlier, we mentioned the TCP **Sequence Number Field**. This field is 32 bits long.

Acknowledgment Numbers Earlier in this chapter, we noted that TCP uses acknowledgments (ACKs) to achieve reliability. The 32-bit **Acknowledgment Number Field** indicates which segment is being acknowledged.[13]

The acknowledgment number indicates which segment is being acknowledged.

Dual-Purpose Segments Note that TCP segments can contain both new information (usually an application message in the data field) and the acknowledgment of a received segment. This is done to minimize the number of TCP segments that are transmitted by the two internet processes.

Test Your Understanding

16. a) What are 1-bit fields called? b) If someone says that a flag field is set, what does this mean? c) If the ACK bit is set, what other field must have a value? d) Why are sequence numbers good? e) What is the purpose of the Acknowledgment Number Field? f) Do SYN segments have data fields? g) Can a single TCP segment both send information and provide an acknowledgment?

User Datagram Protocol (UDP) Datagram Syntax

Applications that cannot use the high functionality in TCP or that do not need this functionality can use the **User Datagram Protocol (UDP)** at the transport layer instead of TCP. UDP does not have openings, closings, or acknowledgments, and so it produces substantially less traffic than TCP.

UDP messages are called **datagrams**. Because of UDP's simple operation, the syntax of the **UDP datagram** shown in Figure 2-14 is very simple. Beside two port number fields, which we will see next in this chapter, there are only two header fields.

- There is a **UDP Length Field** so that the receiving transport process can know how long the datagram is. The packet in the datagram's data field has variable length, so the UDP datagram has variable length.
- There also is a **UDP Checksum Field** that allows the receiver to check for errors in this UDP datagram.[14] If an error is found, however, the UDP datagram is merely

[13] One might expect that if a segment has sequence number X, then the acknowledgment number in the segment that acknowledges it would have acknowledgment number X. The situation is actually more complex. The acknowledgment takes into account both the sequence number of the received TCP segment and its length. TCP does not have segment length information in its header.

[14] If the UDP Checksum Field has 16 zeroes, error checking is not to be done at all.

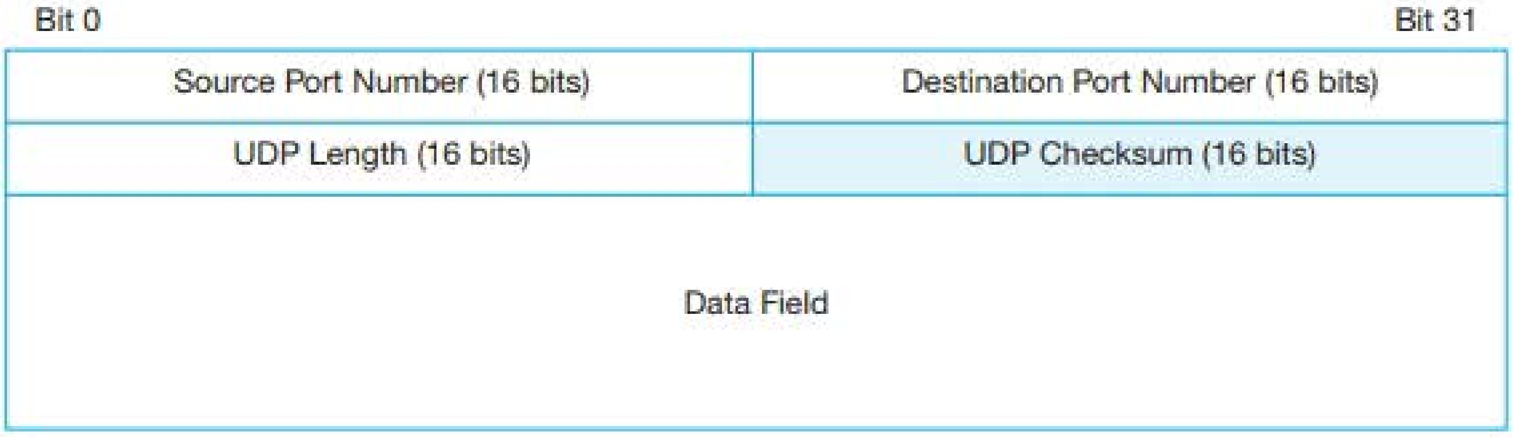

Bit 0	Bit 31
Source Port Number (16 bits)	Destination Port Number (16 bits)
UDP Length (16 bits)	UDP Checksum (16 bits)
Data Field	

FIGURE 2-14 UDP Datagram

discarded. In contrast to TCP but like IP, UDP has no mechanism for retransmission. Like IP, UDP is not reliable.

Test Your Understanding

17. a) What are the four fields in a UDP header? b) Describe the third. c) Describe the fourth. d) UDP does error detection and discarding but does not do the retransmission of damaged or lost datagrams. Is UDP reliable? Explain.

Port Numbers

Both TCP and UDP headers begin with two **port number** fields. The **Source Port Number Field** specifies the sender's port number, and the **Destination Port Number Field** gives the receiver's port number. Servers and clients use these port number fields differently.

Server Port Numbers Computers are multitasking machines, which means that they can run several application programs at the same time. Figure 2-15 shows a server running SMTP (the Simple Mail Transfer Protocol), HTTP, and FTP (the File Transfer Protocol) application programs.

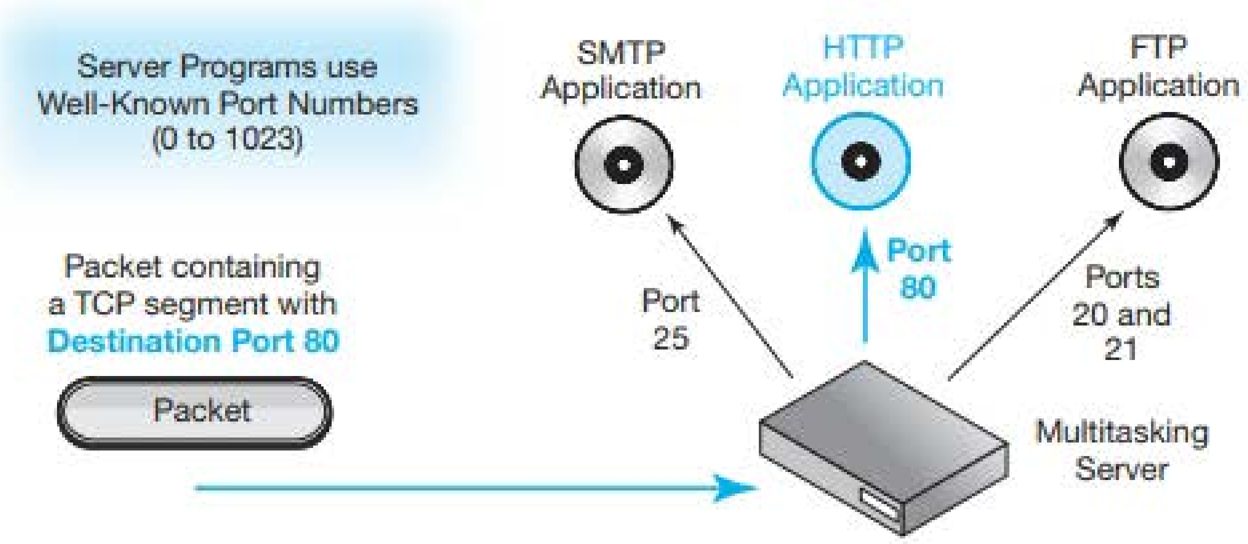

FIGURE 2-15 Server Port Numbers

If a packet arrives, how does the TCP or UDP process know which application program to give the message? This is where TCP and UDP port numbers come in. A server's port number specifies a particular application running on the server. Port 20 or 21 specifies the FTP (File Transfer Protocol) program, Port 25 specifies the SMTP (e-mail) program, and Port 80 specifies the HTTP (World Wide Web) application. These are **well-known port numbers**, which means that they are normally associated with particular application protocols.[15] Port 80 is normally used for HTTP, so if you see Port 80, you know that it probably is HTTP. The well-known port numbers have a port number range reserved for their use—0 through 1023. To send a TCP or UDP message to the application program on a server, the sender puts the appropriate port number in the Destination Port Number Field.

Well-known port numbers for server applications are normally associated with particular application protocols. The well-known port number of HTTP is 80.

Client Port Numbers Clients use port numbers differently. For every conversation a client initiates, it randomly generates an **ephemeral port number**. Ephemeral means that the port number is temporary. It is discarded when a conversation between the client and a particular webserver ends. If the client communicates with the same server program later, the client's transport process will generate a new ephemeral port number. On Windows computers, this is the range from Port 1024 to Port 4999.

Ephemeral port numbers on client computers are only used for a single set of interactions between the client and a server.

Figure 2-16 shows a client host (60.171.18.22) communicating with a blue server host (1.33.17.13). The server port number is Port 80, indicating that the client is communicating with the HTTP program on the server. The client has generated ephemeral Port 2707. When the client transmits to the server, the Source Port Number Field has the value 2707 and the Destination Port Number Field is 80. When the server replies, the source port number is 80 and the destination port number is 2707.

The client is simultaneously connected to an SMTP application on a server (123.30.17.120), which uses the well-known port number 25. For this conversation, the client randomly generates ephemeral Port 4400. When the client transmits, the source port number is 4400 and the destination port number is 25.

Sockets Figure 2-16 shows that a conversation always involves a source IP address and a source port number, plus a destination IP address and a destination port number. It is common to represent each IP address and port number as a **socket**, which

[15] An operating system does not have to use well-known port numbers for applications. For example, some systems administrators assign a different port number to webserver applications, believing that attackers will not be able to identify them as webservers. It doesn't work, and it tends to cause confusion for legitimate users and systems personnel.

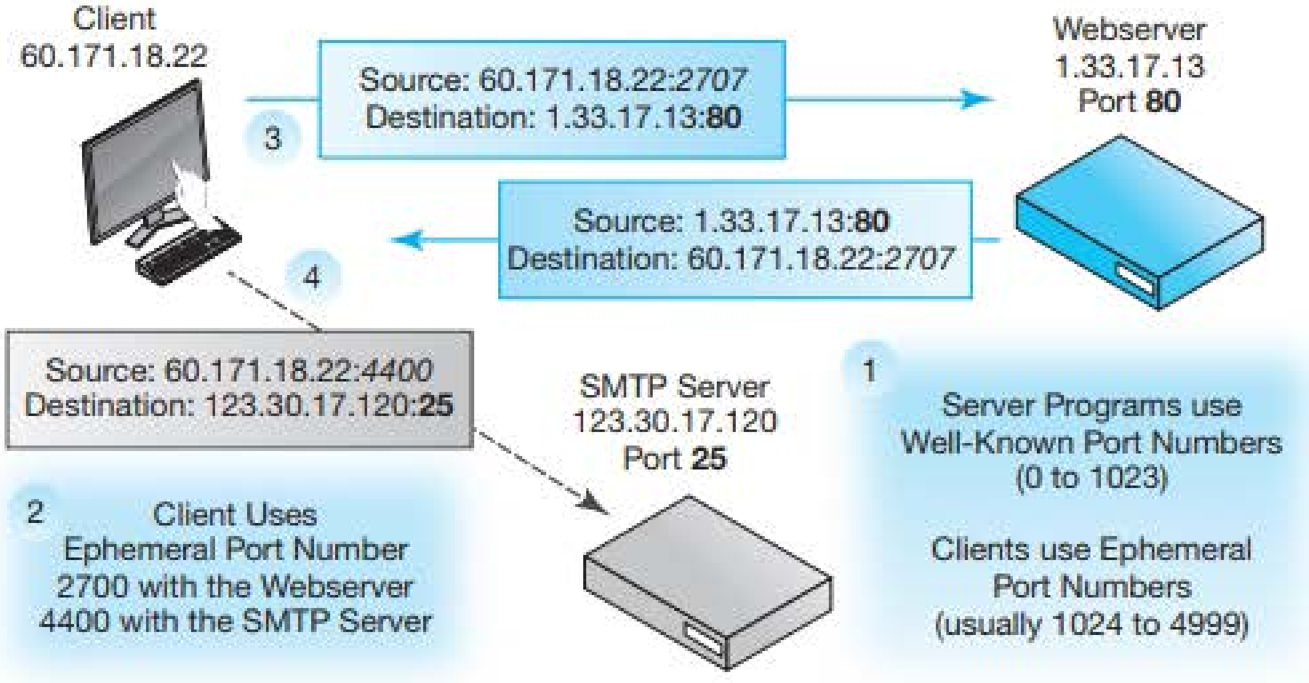

FIGURE 2-16 Client Port Numbers and Sockets

is simply the IP address, a colon, and the port number (Figure 2-17). When the client transmits to the webserver, the source socket is 60.171.18.22:2707 and the destination socket is 1.33.17.13:80. When the webserver replies, the source socket is 1.33.17.13:80 and the destination socket is 60.171.18.22:2707.[16]

Test Your Understanding

18. a) What type of port numbers do servers use for common server programs? b) What type of port numbers do clients use when they communicate with server programs? c) What is the range of port numbers for each type of port? d) How are ephemeral port numbers generated? e) Why are they called ephemeral?

19. a) What is the syntax of a socket? b) In Figure 2-16, when the client transmits to the mail server, what is the source socket? c) What is the destination socket? d) When the SMTP server transmits to the client host, what is the source socket? e) What is the destination socket?

IP Address	Port Number	Socket
60.171.18.22	2707	60.171.18.22:2707
123.30.17.120	80	123.30.17.120:80
60.171.18.22	4400	60.171.18.22:4400
1.33.17.13	25	1.33.17.13:25

FIGURE 2-17 A Socket Is an IP Address, a Colon (:), and a Port Number

[16] Note that the IP address and the port number are not even in the same header. The IP address is in the packet header, and the port number is in the TCP or UDP header.

Frame Syntax

At Layer 2, the data link layer, we have frames. Recall that a frame carries a packet through a single network. In Chapter 1, we looked briefly at the Point-to-Point Protocol and Ethernet. In this subsection, we look very briefly at the syntax of Ethernet frames. The main purpose is to illustrate how packets are encapsulated in frames.

Octets Field lengths are often measured in bits. Another common measure for field lengths in networking is the octet. An **octet** is a group of 8 bits. Isn't that a byte? Yes, exactly. Octet is just another name for byte. The term is widely used in networking, however, so you need to become familiar with it. Octet actually makes more sense than byte, because *oct* means "eight." We have octopuses, octagons, and octogenarians.[17] Octets are usually encountered in data link layer syntax.

An octet is a group of 8 bits.

Octets are usually encountered in data link layer syntax.

The Ethernet II Frame Figure 2-18 shows another frame, an **Ethernet II frame**. For Ethernet, the fields are not shown 32 bits on a line. Instead, the fields are shown in order, one after another. The source and destination EUI-48 addresses that we saw in Chapter 1 have 48-bit fields to hold them, as you would expect. Other field sizes are given in octets.

Note that this is the syntax of the Ethernet II frame. The IEEE 802.3 Working Group actually defines a different frame, which is generally called the 802.2 Ethernet

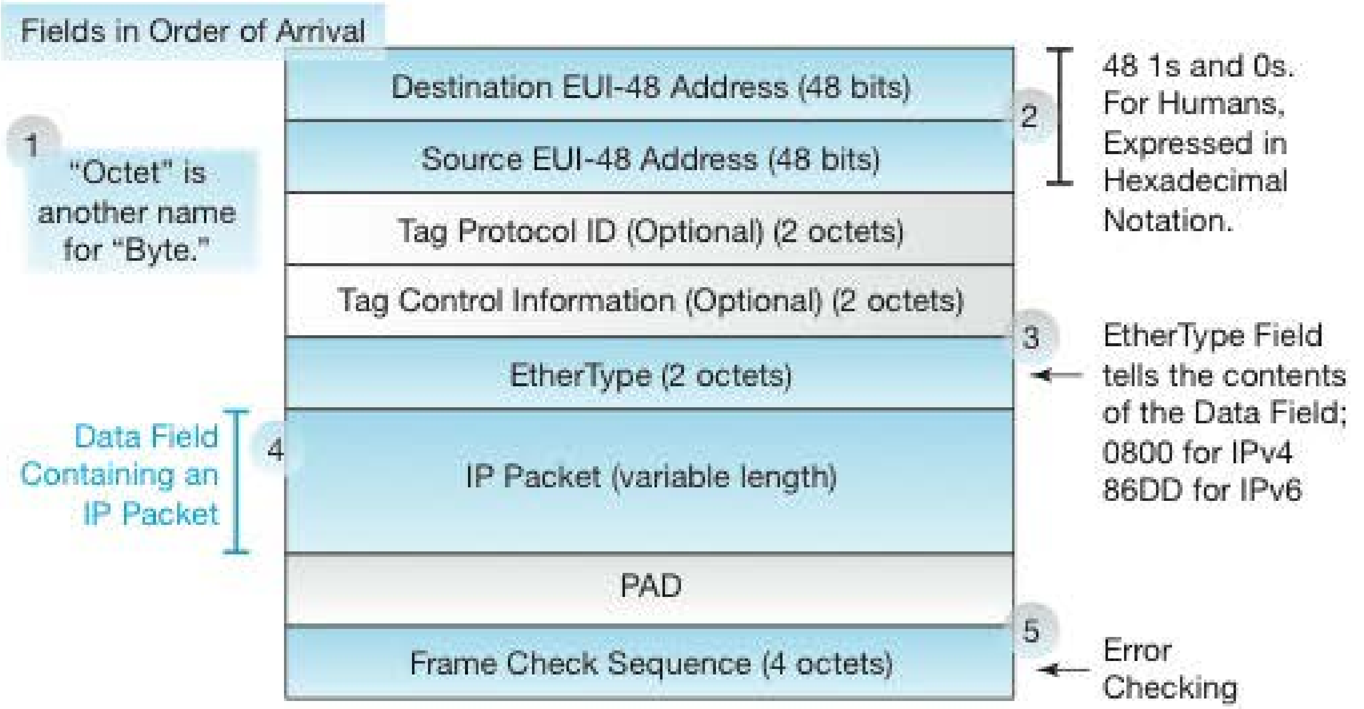
Fields in Order of Arrival
Destination EUI-48 Address (48 bits)
Source EUI-48 Address (48 bits)
Tag Protocol ID (Optional) (2 octets)
Tag Control Information (Optional) (2 octets)
EtherType (2 octets)
IP Packet (variable length)
PAD
Frame Check Sequence (4 octets)
1
"Octet" is another name for "Byte."
2
48 1s and 0s. For Humans, Expressed in Hexadecimal Notation.
3
EtherType Field tells the contents of the Data Field; 0800 for IPv4 86DD for IPv6
4
Data Field Containing an IP Packet
5
Error Checking

FIGURE 2-18 Ethernet II Frame

[17] What is the eighth month? (Careful. The Romans added months to honor Julius and Augustus Caesar.)

frame. The Ethernet II frame is actually the version that existed before the 802 LAN/MAN Standards Committee took over the standardization of Ethernet. In past editions of this text, we described the 802.2 frame syntax, which is arguably more "standard." However, the IETF has specified that IP packets should be encapsulated in Ethernet II frames, and this is usually done in practice. Given that IP packets are the dominant types of messages sent through corporate Ethernet networks, we focus on the Ethernet II frame in this edition of the book.

Ethernet frames can carry many types of information in their data fields. How does the receiver know if the data field contains an IPv4 packet, an IPv6 packet, or something else? That is the job of the **EtherType Field**. If there is an IPv4 packet in the data field, the EtherType Field has the value 0800 in the hexadecimal notation we will see in Chapter 5. In binary, this is 0000100000000000. For an IPv6 packet, the EtherType Field has the hexadecimal value 86DD, which is 1000011011011101 in binary. Figure 2-19 shows an IPv4 packet in the data field.

Both the PPP frame and the Ethernet frame have a **Frame Check Sequence Field**. These fields allow the receiver to check an arriving frame for errors. The sender does a calculation based on the bits in the frame and places the result in the Frame Check Sequence Field. When the frame arrives, the receiver repeats the calculation and compares its calculated result with the transmitted value in the arriving frame. If the two are different, an error has occurred.

In Ethernet, the receiver simply discards any frame with errors. This is error detection without retransmission, so Ethernet is an *unreliable protocol*. It is also connectionless.

Test Your Understanding

20. a) How is the syntax of Ethernet II frames depicted? b) In what field is the IP packet in carried Ethernet II frames? c) Why does this version of the book deal with Ethernet II frames? d) How does the receiving data link layer process know what is in the data field of an Ethernet II frame? e) Why is Ethernet unreliable despite having a Frame Check Sequence Field that is used to check for errors?

ENCODING APPLICATION MESSAGES INTO BINARY

Encoding

Application messages include letters, numbers, pictures, video streams, and other types of information. Lower-layer messages, as we saw earlier, consist of 1s and 0s. The application program itself must convert its various types of information into bit streams. This conversion of rich application data into binary is called **encoding**.

Encoding is the conversion of application messages into bits.

It is done by the application program.

Test Your Understanding

21. a) What is encoding? b) At what layer is the encoding of application messages done?

Encoding Text as ASCII

To encode **alphanumeric** information (text, numbers, and other keyboard characters), applications normally use the **ASCII code,**[18] whose individual symbols are each 7 bits long. Seven bits give 128 possibilities, as we will see later. This is enough for all keys on the keyboard plus some extra control codes.

Alphanumeric information consists of text, numbers, and other keyboard characters.

Figure 2-19 shows a few ASCII codes. Note that uppercase letters and lowercase letters have different ASCII codes. This is necessary because the destination application program must know whether to convert the encoded character into uppercase or lowercase. ASCII also encodes the digits from 0 through 9, as well as punctuation and other special characters. There are even ASCII control codes that tell the receiver what to do. For example, a carriage return is 0101110.

For transmission, the 7 bits of each ASCII character are placed in a byte. The 8th bit in the byte is not used today.[19]

Category	Example	7-Bit ASCII Code	8th bit in Transmitted Byte
Upper-Case Letters	A	1000001	Unused
Lower-Case Letters	a	1100001	Unused
Digits (0 through 9)	3	0110011	Unused
Punctuation	Period	0101110	Unused
Punctuation	Space	0100000	Unused
Control Codes	Carriage Return	0001101	Unused
Control Codes	Line Feed	0001010	Unused

FIGURE 2-19 Encoding Text as ASCII

[18] ASCII is not the only system for encoding alphanumeric data. The A in ASCII stands for "American." It does not represent diacritical marks, so there are variations for different languages. The better choice for international communication is UNICODE, which can represent any language–at the cost of more complexity.

[19] Early systems used the 8th bit in each byte as a "parity bit" to detect errors in transmission. The total number of bits in all bytes was made a whole odd (or even) number by selecting the parity bit. This could detect a change in a single bit in the byte. At today's high transmission speeds, however, transmission errors normally generate multibit errors rather than single-bit errors. Consequently, parity is useless and is ignored.

Test Your Understanding

22. a) What is alphanumeric information? b) Explain how many bytes it will take to transmit "Go team" without the quotation marks. (Answer: 7) c) Explain how many bytes it will take to transmit "Hello World!" without the quotation marks. d) Go to a search engine and find a converter to represent characters in ASCII. What are the 7-bit ASCII codes for "Hello world!" without the quotation marks? (Check: H is 1001000) Show this in a table with two columns. The first will show letters or other keyboard characters. The second will show the ASCII code for that character.

Converting Integers into Binary Numbers (1s and 0s)

Some application data consists of **integers**, which are whole numbers (0, 1, 2, 3, . . . 345, etc.). Humans write these as decimal numbers, in which each symbol is a digit from 0 through 9. The sending application program encodes integers as **binary numbers** (1s and 0s).

In decimal numbers, each symbol is a digit from 0 through 9.

In binary numbers, each symbol is a 1 or a 0.

Encoding Small Decimal Integers to Binary Using Your Brain Figure 2-20 shows how you can encode (convert) decimal integers to binary. Decimal is our normal number system. Integers are whole numbers.

- The decimal number to be converted to binary is 11.
- (1) Next, write the bit positions, from 0 through 6, writing them *right to left*. The first bit position on the right is 0, not 1.
- (2) Under each write the position value. This is 2 raised to the power of the position number. For bit position 5, this is 2^5 or 32.

0	Decimal number to be converted	11						
1	Bit position, b (begins with 0 from right)	6	5	4	3	2	1	0
2	Position Exponent	2^6	2^5	2^4	2^3	2^2	2^1	2^0
2	Position value, 2^b (Acts as the alternative in $a = 2^b$)	64	32	16	8	4	2	1
3	Available positions for the conversion of 11. (Less than or equal to 11)	-	-	-	Yes	Yes	Yes	Yes
4	Combination that produces 11 (Try combinations in your head)	-	-	-	1	0	1	1
5	The number in binary	1011						

FIGURE 2-20 Encoding (Converting) a Small Decimal Number to Binary

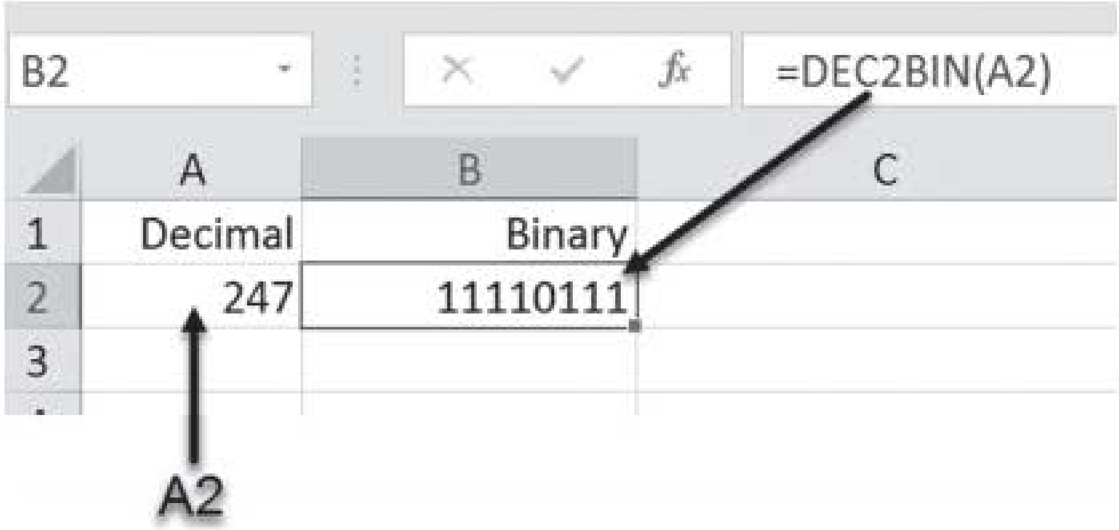

B2 =DEC2BIN(A2)

	A	B	C
1	Decimal	Binary	
2	247	11110111	
3			

FIGURE 2-21 Converting Decimal to Binary in Excel

- (3) Now note which positions are relevant for encoding 11. The largest value that will fit into 11 is 8, so only the first four bit positions are relevant.
- (4) Now, look at the values 8, 4, 2, and 1. Decide how to get 11 from them. The answer turns out to be 8 + 2 + 1. Put 1s in the 8, 2, and 1 positions. Another relevant position adds no value, so it gets a 0.
- The answer, then, is 1011.

Encoding Decimal Integers to Binary Using a Computer This works fine for small decimal numbers. For larger numbers like 247, Excel offers the dec2bin function (Figure 2-21). If you compute dec2bin(11), you will get 1011. If you compute dec2bin(247), you get 11110111. Most other spreadsheet programs have similar functionality. You can also use a search engine to find an online decimal to binary converter.

Converting Binary to Decimal Using Your Brain You should also know how to convert binary numbers that you come across back to decimal. Figure 2-22 shows how to do this for the binary number 1010. It is obviously the reverse of the encoding process. Excel offers bin2dec, and there are many binary to decimal converters on the Internet.

Binary number to be converted	1010				
Bit position, b (begins with 0 from right)	4	3	2	1	0
Position exponent	2^4	2^3	2^2	2^1	2^0
Position value, 2^b (Acts as the alternative in $a = 2^b$)	16	8	4	2	1
Binary number to be converted	0	1	0	1	0
Decimal Equivalents	0	8	0	2	0
Decimal Representation	1010 = 1*8 + 1*2 = 10				

FIGURE 2-22 Converting Binary Numbers to Decimal

Test Your Understanding

23. Answer the following questions without using a calculator or a computer. a) What is an integer? b) Is 4,307 an integer? c) Is 45.7 an integer? d) Is the first bit position on the right 0 or 1? e) Convert the decimal number 6 to binary without using a computer. (Answer: 110) f) Convert 0 to binary. g) Convert 15 to binary. h) Convert 62 to binary. i) This time using Excel or a decimal to binary converter, convert 128 to binary. (Answer:10000000). j) Also using Excel or a decimal to binary converter, convert 255 to binary. k) Convert the binary number 100 to decimal. (Answer: 4) l) Convert the binary number 1111 to decimal. m) Convert the binary number 10110 to decimal. n) Convert the binary number 100100 to decimal.

Encoding Alternatives

Some application data can be expressed as alternatives, such as North, South, East, or West. The application layer process will create a field in the application layer message and represent each alternative as a group of bits. For instance, the four cardinal compass points can be represented by a 2-bit field within the application message. North, South, East, and West can be represented as 00, 01, 10, and 11, respectively. (These are the binary numbers for 0, 1, 2, and 3.) There is no order to the alternatives, so any choice can be represented by any pair of bits.

We just saw that having four alternatives requires a 2-bit field. More generally, if a field has *b* bits, it can represent 2^b alternatives. This gives us the following equation:

Equation 1: $a = 2^b$, where *a* is the number of alternatives and *b* is the number of bits

We have just seen that a 2-bit field can represent 2^2 alternatives, or 4. Here, *b* is 2, so *a* is 4. What if you need to represent six alternatives? Two bits will not be enough, because 2^2 is only 4 and we need 6. A three-bit field will give us 2^3 alternatives, or 8. This gives us enough alternatives. Two alternatives will go unused.

If a field has b bits, it can represent 2^b alternatives.

Figure 2-23 illustrates how alternative encoding is done for fields that have 1, 2, 4, 8, 16, and 32 bits. It shows that with 1 bit you can encode yes or no, connection-oriented or connectionless, or any other dichotomy. Two bits, as we just saw, are good for the four cardinal compass points. With 4 bits, you can have up to 16 alternatives.

As noted in a previous example, not every set of categories will have exactly two-to-some-power items. Figure 2-23 shows that to represent the top 10 security threats, you need 4 bits, which can encode up to 16 alternatives. (Three bits will encode only eight alternatives.) Using 4 bits to represent 10 threats will "waste" six alternatives, but this is necessary.

Bits in Field (b)	Number of Alternatives (a) that can be Encoded (a =2^b)	Possible Bit Sequences	Examples
1	$2^1 = 2$	0, 1	Yes or No, Connection-oriented or connectionless, etc.
2	$2^2 = 4$	00, 01,10, 11	North, South, East, West; Red, Green, Blue, Black
4	$2^4 = 16$	0000, 0001, 0010, . . .	Top 10 security threats. 3 bits would give 8 alternatives. Not enough. 4 bits works. 6 values go unused
8	$2^8 = 256$	00000000, 00000001, . . .	One byte per color gives 256 possible colors levels.
16	$2^{16} = 65{,}536$	0000000000000000, 0000000000000001, '. . .	Two bytes per color gives 65,536 color levels.
32	$2^{32} = 4{,}294{,}967{,}296$	000000000000000 0000000000000000, etc.	Number of Internet Protocol Version 4 addresses

FIGURE 2-23 Binary Encoding to Represent a Certain Number of Alternatives

You should memorize the number of alternatives that can be represented by 4, 8, and 16 bits, because these are common field sizes. Each added bit doubles the number of possible alternatives, and each bit subtracted cuts the number of possible alternatives in half. So, if you remember that 8 bits can represent 256 alternatives, 7 bits (one less) can represent 128 alternatives (half as many), and 9 bits (one more) can represent 512 alternatives (twice as many). How many alternatives can 6 and 10 bits represent?

Test Your Understanding

24. a) How many alternatives can you represent with a 4-bit field? (Answer: 16) b) For each bit you add to an alternatives field, how many additional alternatives can you represent? c) How many alternatives can you represent with a 10-bit field? (With 8 bits, you can represent 256 alternatives.) d) If you need to represent 129 alternatives in a field, how many bits long must the field be? (Answer: 8) e) If you need to represent 18 alternatives in a field, how many bits long must the field be? f) Come up with three examples of things that should be encoded with 3 bits.

24. a) In TCP, port number fields are 16 bits long. How many possible port numbers are there? b) IPv6 addresses are 128 bits long. How many IPv6 addresses are there? Just represent the formula for calculating the value. c) The IP version number field is 4 bits long. How many possible versions of IP can there be? d) UDP length fields are 16 bits long. This field gives the number of bytes in the data field. How many bytes long may a UDP data field be? e) ASCII has a 7-bit code. How many keyboard characters can it represent?

PROTOCOLS IN THIS CHAPTER

Figure 2-24 lists information on several protocols we saw in this chapter, including layer number, whether the protocol is connectionless or connection-oriented, and whether the standard is reliable or unreliable. Note a very simple pattern. Only TCP among these major protocols is reliable and connection-oriented.[20] Although reliability appears to be a good thing, it is complex and resource-consuming. Connection orientation, in turn, is usually done to make reliability possible through message retransmission. Making all layers reliable would be extremely expensive.

Only TCP among these major protocols is reliable and connection-oriented.

Test Your Understanding

25. a) What protocols that we saw in this chapter are reliable? b) Why aren't all protocols reliable?

Layer	Protocol	Reliable or Unreliable?	Connection-Oriented or Connectionless
5	HTTP	Unreliable	Connectionless
4	TCP	Reliable	Connection-Oriented
4	UDP	Unreliable	Connectionless
3	IP	Unreliable	Connectionless
2	PPP	Unreliable	Connectionless
2	Ethernet	Unreliable	Connectionless

FIGURE 2-24 Protocols in this Chapter

[20] Why make TCP the reliable protocol? Recall that lower layers usually do error discarding. This means that if there has been an error at the layers below TCP, the TCP segment will not reach the transport process on the receiver. There will be no acknowledgment, so the source transport process will retransmit the TCP segment. TCP, then, automatically corrects errors at lower layers by retransmitting a discarded segment. TCP, furthermore, lies right below the application layer, so error correction at the transport layer gives the application program clean data. The application program should not have to deal with transmission errors.

END-OF-CHAPTER QUESTIONS

Thought Questions

2-1. How do you think TCP would handle the problem if an acknowledgment were lost, so that the sender retransmitted the unacknowledged TCP segment, therefore causing the receiving transport process to receive the same segment twice?

2-2. a) Compute the minimum number of TCP segments required to open a connection, send an HTTP request and response message, and close the connection. Justify this number by creating a table showing each message and its sequence number. b) Repeat the question, this time if the HTTP response message is damaged during transmission.

2-3. Compute the minimum number of TCP segments required to open a connection, send an HTTP request and response message, and close the connection if the HTTP response message must be fragmented across ten packets. Justify this number by creating a table showing each message and its sequence number.

2-4. a) In Figure 2-16, what will be the value in the destination port number field if a packet arrives for the e-mail application? b) When the HTTP program on a webserver sends an HTTP response message to a client PC, in what field of what message will it place the value 80?

2-5. Do the following without using a calculator or computer, but check your answers with a calculator or a computer. a) Convert 6 to binary. (Answer: 110) b) Convert 47 to binary. c) Convert 100 to binary. d) Convert 110100 to decimal. (Answer: 52) e) Convert 001100 to decimal.

2-6. Do the following without using a calculator or a computer, but check your answers with a calculator or a computer. You need to represent 1,026 different city names. How many bits will this take if you give each city a different binary number? Explain your answer.

2-7. a) The port number fields in TCP and UDP are 16 bits long. How many port numbers can they represent? b) In IP, the Time to Live Field is 8 bits in size. How many values can it represent? c) How many values can a flag field represent?

Internet Research: April 1 RFCs

2-8. Consult the Wikipedia Webpage April Fools' Day Request for Comments (https://en.wikipedia.org/wiki/April_Fools%27_Day_Request_for_Comments). Select one of the RFCs listed on the page and write a paragraph on its claimed purpose. (Don't just pick the first few.)

Perspective Questions

2-9. What was the most surprising thing you learned in this chapter?

2-10. What was the most difficult material for you in this chapter?

This page intentionally left blank

Chapter **3**

Network Management

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Discuss network quality of service (QoS) and specify service level agreement (SLA) guarantees.
- Design a network layout based on required traffic volumes between pairs of sites.
- Describe options for dealing with momentary traffic peaks.
- Describe the benefits and importance of centralized network management; discuss and compare three tools for centralizing network management: Ping, Traceroute, and the Simple Network Management Protocol (SNMP).
- Describe Software-Defined Networking (SDN), including why it is potentially revolutionary.

INTRODUCTION

Technology means nothing unless a company manages networks very well. In this chapter, we look at core issues and tools for network management. These concepts apply to everything networking professionals do at every level.

Today, we can build much larger networks than we can manage. Even a midsized bank is likely to have 500 Ethernet switches and a similar number of routers. Furthermore, network devices and their users are often scattered over large regions—sometimes internationally. Although network technology is exciting to talk about, it is chaos without good management.

A pervasive issue in network management is cost. In networking, you never say, "Cost doesn't matter." Network budgets are always stretched thin. Networking and security professionals always need to solve problems with minimum budgets. One way to do this is to automate as much network management work as possible.

NETWORK QUALITY OF SERVICE (QoS)

In the early days of the Internet, networked applications amazed new users. However, the next impression was, "Too bad it doesn't work better." Today, networks are mission-critical for corporations. If the network breaks down, much of the organization comes to an expensive halt. Today, networks must not only work, they must work *well*. Companies are increasingly concerned with network **quality-of-service (QoS) metrics**, that is, quantitative measures of network performance that define what "working well" means and measure how well the network is providing its services. Figure 3-1 shows that companies use several QoS metrics. Collectively, these metrics track the service quality that users receive.

Test Your Understanding

1. a) What are QoS metrics? (Do not just spell out the acronym.) b) Why are QoS metrics important?

Transmission Speed[1]

The first question people have about a newborn baby is, "Is it a boy or a girl?" For a network, the usual question is, "How fast is it?" The answer is important, but it is a little complicated.

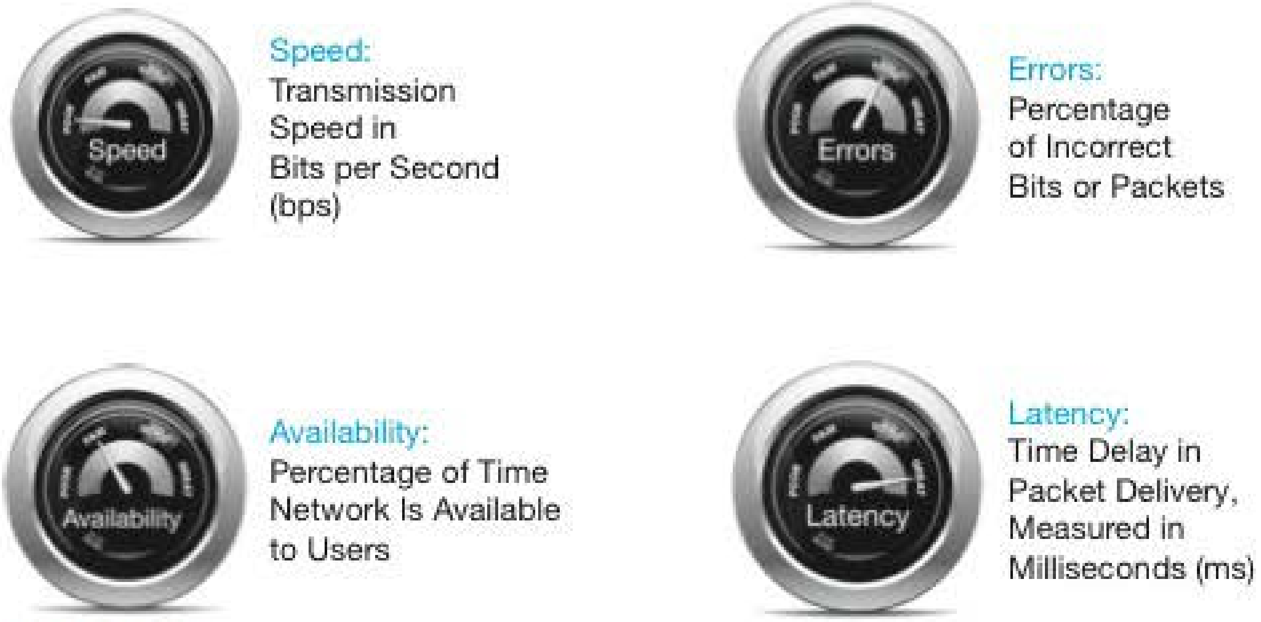
Speed
Speed:
Transmission Speed in Bits per Second (bps)
Errors
Errors:
Percentage of Incorrect Bits or Packets
Availability
Availability:
Percentage of Time Network Is Available to Users
Latency
Latency:
Time Delay in Packet Delivery, Measured in Milliseconds (ms)

FIGURE 3-1 Quality-of-Service (QoS) Metrics

[1] Strictly speaking, speed means velocity. But a host that transmits faster does not send its bits with higher velocity when it transmits faster. It merely transmits more bits each second. Speed is really about transmission rate, not transmission velocity. It is like talking faster, not running faster.

Designation	Abbreviation	Meaning	Example	Without a Metric Prefix
Kilobits per Second	kbps	1,000	4.5 kbps 93.047 kbps	4,500 bps 93,047 bps
Megabits per Second	Mbps	1,000,000	251.62 Mbps	251,620,000 bps
Gigabits per Second	Gbps	1,000,000,000	8 Gbps	8,000,000,000 bps
Terabits per Second	Tbps	1,000,000,000,000	12 Tbps	12,000,000,000,000 bps

FIGURE 3-2 Speeds in Metric Notation

Bits per second (bps) First, network speed is measured in **bits per second (bps)**. Note that this is *not bytes per second*. IS students and professionals tend to think in terms of bytes because of their file, database, and programming background. However, it is traditional to deal with bits in transmission systems. For things like file downloads, you occasionally do see speeds measured in bytes per second. In this case, the speed should be shown as Bps, not bps.

Speed is normally measured in bits per second (bps), not bytes per second (Bps).

High speeds are written in metric notation. As Figure 3-2 shows, transmission speeds are measured in **kilobits per second (kbps), Megabits per second (Mbps), Gigabits per second (Gbps)**, and **Terabits per second (Tbps)**. Notice an oddity in the metric system. Mega, Giga, and Tera have uppercase metric prefixes, but kilo is a lowercase k.[2]

Kilo is abbreviated with a lowercase k (kbps).

Application Requirements How much speed is necessary? That depends on the application that needs to be supported. Figure 3-3 shows download times for various applications at various download speeds. Note that for messaging and e-mail, any speed is fine. At the other extreme, high-definition video and full-disk backups need higher speeds than we generally get today.

Test Your Understanding

2. a) Is transmission speed usually measured in bits per second or bytes per second? b) For the HDTV program in Figure 3-3, which of the speeds shown will allow real-time streaming?

[2] In the metric system, uppercase K is the abbreviation for Kelvins, a measure of temperature.

Application	100 kbps	1 Mbps	10 Mbps	100 Mbps	1 Gbps
E-Mail Message (250 words)	0 sec	0 sec	0 sec	0 sec	0 sec
Photograph 5 MB	7 min	40 sec	4 sec	0 sec	0 sec
1-hr HDTV Program (7 Mbps)	3 da	7 hr	42 min	4 min	25 sec
Backup, 1 TB Hard Drive	31 mo	3 mo	7 da	22 hr	2 hr

FIGURE 3-3 Application Download Times at Various Transmission Speeds

Rated Speed and Throughput It is important to understand that there is a difference between rated speed and throughput. The **rated speed** is the speed the standard or the carrier specifies. **Throughput** is the speed you actually receive, which is lower, sometimes much lower. For working networking professionals, throughput is the only thing that is relevant.

The rated speed is the speed the standard or the carrier specifies.

Throughput is the speed you actually receive.

Shared Speed: Aggregate and Individual Throughput Things get even more confusing if the speed being delivered by the network is shared by several people. For example, if you are at a coffee shop that has an access router with a built-in access point, it will provide a certain **aggregate throughput** that is shared by everyone who is sending and receiving. Figure 3-5 shows this in an example. In the coffee shop, there are 5 users. Two are actively sending or receiving at this moment. The rated speed of the access point is 6 Gbps. Its aggregate throughput is 5 Gbps. This aggregate throughput is shared by the 2 users actively sending and receiving, not by all 10. Each has an **individual throughput** of 2.5 Gbps.

In a shared system, the aggregate throughput is the throughput available to all users.

The individual throughput is the aggregate throughput divided by the number of active users at the moment.

Rated Speed	The speed stated in the standard or the speed you are quoted by your provider.
Throughput	The speed you actually get. (Almost always lower, sometimes substantially.)

FIGURE 3-4 Rated Speed and Throughput (Study Figure)

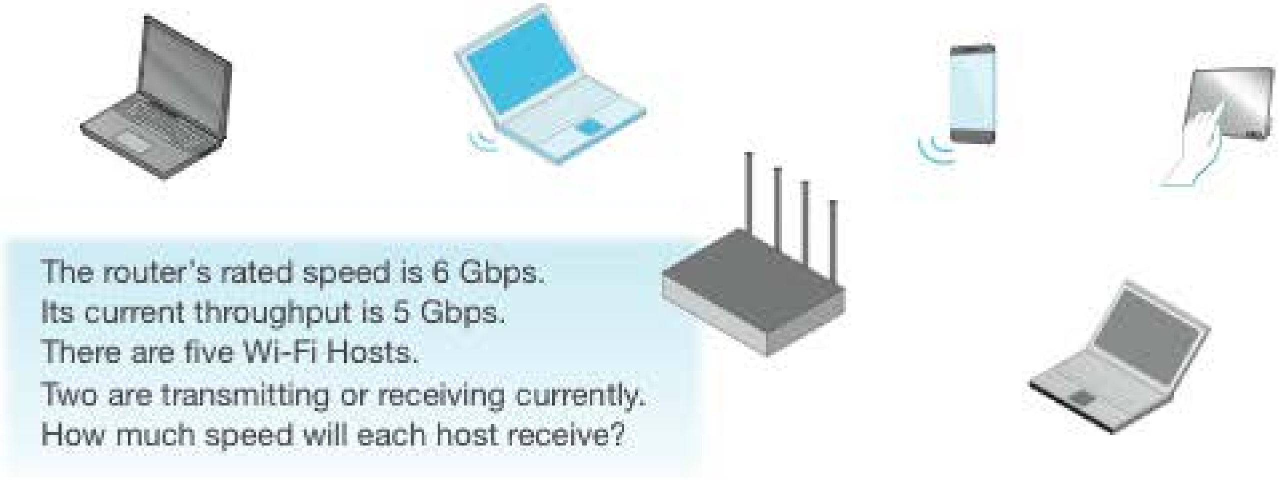

The router's rated speed is 6 Gbps.
Its current throughput is 5 Gbps.
There are five Wi-Fi Hosts.
Two are transmitting or receiving currently.
How much speed will each host receive?

FIGURE 3-5 Shared Throughput and Individual Throughput in a Coffee Shop

Test Your Understanding

3. a) Distinguish between rated speed and throughput. b) Distinguish between individual and aggregate throughput. c) You are working at an access point with 20 other people. Three are doing a download at the same time you are. The rest are looking at their screens or sipping coffee. The access point you share has a rated speed of 150 Mbps and provides a throughput of 100 Mbps. How much speed can you expect on average for a download? (Answer: 25 Mbps) d) In a coffee shop, there are 10 people sharing an access point with a rated speed of 2 Gbps. The throughput is half the rated speed. Several people are downloading. Each is getting an average of 100 Mbps. How many people are using the Internet at that moment?

Transmission Capacity on Multiplexed Transmission Links The transmission links that connect pairs of routers in the Internet may **multiplex** (combine) the traffic of thousands or millions of conversations. Multiplexing is shown in highly simplified form in Figure 3-6. In the figure, the multiplexed transmission link shares the traffic of only two connections. Hosts A and B generate 4 Gbps of traffic. For Hosts C and D, the traffic is 5 Gbps. Access links that connect each host to the network are not shared, so these **dedicated links** need to be able to carry the traffic of their individual hosts.

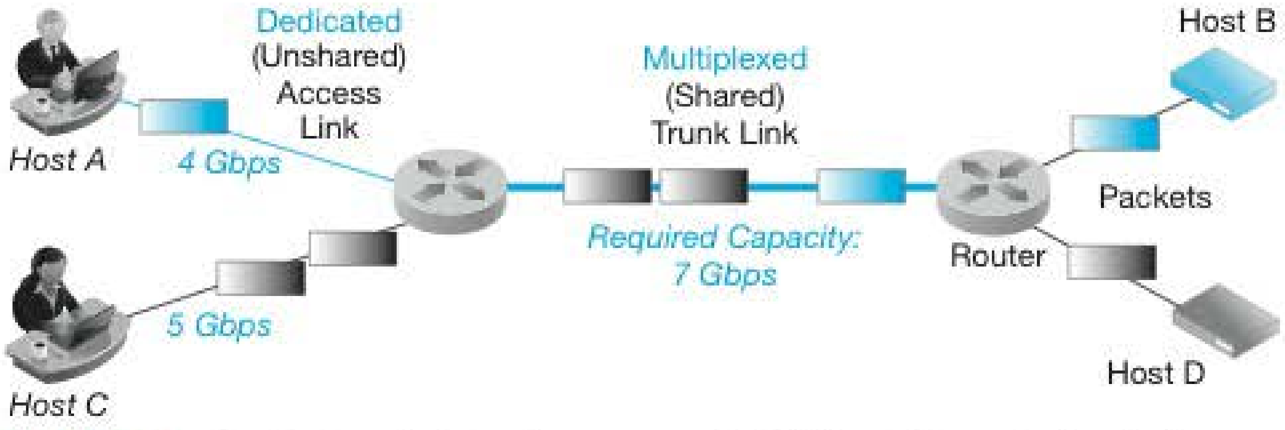

FIGURE 3-6 Traffic Capacity Requirement on a Multiplexed Transmission Link

Multiplexing transmits the traffic of multiple conversations over a shared trunk link, as opposed to unshared access links. This saves money.

Why Multiplexing? Why multiplex the traffic of many individual conversations on trunk links? The answer is that multiplexing reduces cost. It is cheaper to multiplex many conversations on a single line than to give each a line. There are economies of scale in transmission lines, so one big trunk link will be cheaper than having separate unshared links to give each pair its required capacity. In addition, hosts do not transmit constantly. They normally transmit in bursts separated by relatively long silences. Multiplexing packs frames or packets onto the line more efficiently, allowing a slower trunk line to be used. To give an analogy, although it might be nice for each car to have its own lane during rush hour, they need to share a few lanes to minimize cost.

Multiplexing reduces cost.

The multiplexed trunk link between the pairs of routers or switches must be sized to the *average* traffic of all conversations. Although both Host A and Host B generate a good deal of traffic, they do so by sending short bursts of packets separated by silences. So, although Host A and Host C collectively generate 9 Gbps of traffic, the multiplexed transmission line might only need to be, say, 7 Gbps to carry their combined traffic. Multiplexing, then, saves money by allowing multiple conversations to share multiplexed transmission links efficiently.

The Downside of Sharing Your mother probably told you that sharing is good. As you soon learned, she was wrong. Bad things happen when you share. Your individual throughput will vary with traffic, especially if the system is near capacity.

Test Your Understanding

4. a) Distinguish between dedicated and multiplexed transmission links. b) If 100 conversations averaging 50 Mbps are multiplexed on a transmission line, will the required transmission line capacity be less than 5 Gbps, equal to 5 Gbps, or more than 5 Gbps? c) What is the business benefit of multiplexing?

Other Quality-of-Service Metrics

Although network speed is important, it is only one element in networking quality of service. Figure 3-1 showed three other QoS categories. We look briefly at each.

Availability One is **availability**, which is the percentage of time that the network is available for use. Ideally, networks would be available 100% of the time, but that is impossible in reality. Now that networking is embedded in almost every aspect of business, a breakdown in availability quickly becomes intolerable.

Error Rates Ideally, all packets would arrive intact, but a few will not. The **error rate** is the percentage of bits or packets that are lost or damaged during delivery. (At the physical layer, it is common to measure bit error rates. At the internet layer, it is common to measure packet error rates.)

When the network is overloaded, error rates can soar because the network must drop the packets it cannot handle. Consequently, companies must measure error rates when traffic levels are high to have a good understanding of error rate risks.[3]

Latency When packets move through a network, they always encounter some delays. The amount of delay is called **latency.** Latency is measured in **milliseconds (ms).** A millisecond is a thousandth of a second. When latency reaches about 125 milliseconds, turn taking in telephone conversations becomes difficult. You think the other person has finished speaking, so you begin to speak—only to realize that the other party is still speaking.

The amount of network delay is called latency. Latency is measured in milliseconds (ms).

Jitter Figure 3-7 illustrates another time-related QoS concept, jitter. **Jitter** is the *average variability* in the latency between successive packets. Some packets will arrive farther apart in time, others closer in time. Jitter does not bother most applications, but voice over IP (VoIP) and streaming media are highly sensitive to jitter. If the sound is played back without adjustment, it will speed up and slow down. These variations often occur over millisecond times. As the name suggests, variable latency tends to make voices sound jittery.[4]

Jitter is the average variability in arrival times (latency).

Engineering for Latency and Jitter Most networks were engineered to carry traditional data such as e-mail and database transmissions. In traditional applications,

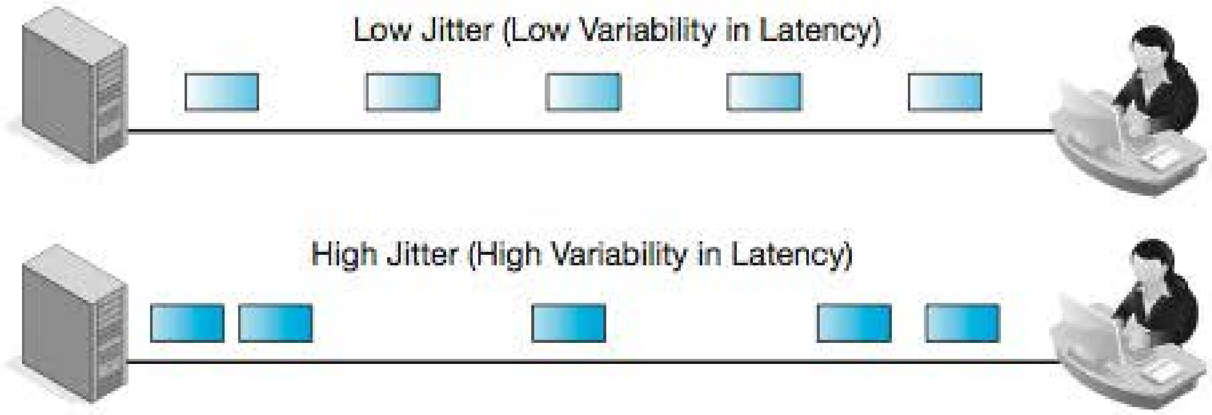
Low Jitter (Low Variability in Latency)
High Jitter (High Variability in Latency)

FIGURE 3-7 Jitter

[3] The impact of even small error rates can be surprisingly large. TCP tries to avoid network congestion by sending TCP segments slowly at the beginning of a connection. If these segments get through without errors, TCP sends the following segments more quickly. However, if there is a single error, the TCP process assumes that the network is overloaded. It falls back to its initial slow start rate for sending TCP segments and builds speed slowly. This can produce a major drop in throughput for applications.

[4] The technical term for jitter is "IP packet delay variation," but jitter is almost always used to designate the phenomenon. RFC 3393 describes how jitter can be measured. Do not attempt to read it unless you have strong headache medicine immediately available.

latency was only slightly important, and jitter was not important at all. However, as VoIP, video, and interactive applications have grown in importance, companies have begun to worry more about latency and jitter. They are finding that extensive network redesign may be needed to give good control over latency and jitter. This may include forklift upgrades for many of its switches and routers.

Test Your Understanding

5. a) What is availability? b) When should you measure error rates? Why? c) What is latency? d) In what units is latency measured? e) What is jitter? f) Why may adding applications that cannot tolerate latency and jitter be expensive?

Service Level Agreements (SLAs)

When you buy some products, you receive a guarantee that promises that they will work according to specifications and that lays out what the company must do if they do not. In networks, service providers may provide **service level agreements (SLAs)**, which are contracts that guarantee levels of performance for various metrics such as speed and availability. If a service does not meet its SLA guarantees, the service provider must pay a penalty to its customers.

Service level agreements (SLAs) are contracts that guarantee levels of performance for various metrics such as speed and availability.

Service Level Agreements (SLAs)

- Guarantees for performance
- Penalties if the network does not meet its service metrics guarantees

Guarantees Specify Worst Cases (No Worse than)

- Lowest speed (e.g., no worse than 100 Mbps)
- Maximum latency (e.g., no more than 125 ms)
- SLAs are like insurance policies—take effect when something bad happens

Often Written on a Percentage Basis

- E.g.: No worse than 100 Mbps 99% of the time
- As the percentage increases, cost of engineering increases in order to achieve it
- To specify 100% of the time would cost an infinite amount of money

Residential Services Are Rarely Sold with SLA Guarantees

- Engineering SLA-compliant networking would be too expensive

FIGURE 3-8 Service Level Agreements (SLA) (Study Figure)

Worst-Case Specification SLA guarantees are expressed as **worst cases**. For example, an SLA for speed would guarantee that speed will be *no lower* than a certain amount. If you are downloading webpages, you want at least a certain level of speed.

You certainly would not want a speed SLA to specify a *maximum* speed. More speed is good. Why would you want to impose penalties on the network provider for exceeding some maximum speed? That would give them a strong incentive not to increase speed! Making things better is not the SLA's job.

SLA guarantees are expressed as worst cases. Service will be no worse than a specific number.

For latency, in turn, an SLA specifying a worst case would require that latency will be *no higher* than a certain value. You might specify an SLA guarantee of a maximum of 65 ms (milliseconds). This means that you will not get worse (higher) latency.

Percentage-of-Time Elements Most SLAs have percentage-of-time elements. For instance, an SLA on speed might guarantee a speed of at least 480 Mbps 99.9% of the time. This means that the speed will nearly always be at least 480 Mbps but may fall below that 0.1% of the time without incurring penalties. A smaller exception percentage might be attractive to users, but it would probably require a substantially more expensive network. Nothing can be guaranteed to work properly 100% of the time, and beyond some point, cost grows very rapidly with increasing percentage guarantees. SLAs must always balance quality level and cost.

Corporations versus Individuals Companies that use commercial networks expect SLA guarantees in their contracts despite the fact that engineering networks to meet these guarantees will raise costs and prices. Businesses need these performance levels to do their work. Consumer services, however, rarely have SLAs because consumers are more price sensitive. For example, residential Internet access providers using a digital subscriber line (DSL), cable modem, or cellular network rarely offer SLAs. This keeps the price of residential services down, but there will be more instances of less-than-advertised performance.

Test Your Understanding

6. a) What are service level agreements? b) Does an SLA measure the best case or the worst case? c) Would an SLA specify a highest speed or a lowest speed? d) Would an SLA specify a highest availability or a lowest availability? e) Would an SLA specify highest latency or lowest latency? f) Would an SLA guarantee specify a highest jitter or a lowest jitter? g) What happens if a carrier does not meet its SLA guarantee? h) If carrier speed falls below its guaranteed speed in an SLA, under what circumstances will the carrier *not* have to pay a penalty to the customers? i) Does residential ISP service usually offer SLA guarantees? Why or why not? j) A business has an Internet access line with a maximum speed of 100 Mbps. What *two* things are wrong with this SLA?

NETWORK DESIGN

Network design is a core skill. The more you know about networking and your corporation's situation, the better your design will be. However, if there is something you do not know or think about, your design is likely to be a poor one. Network designers are governed by their worst moments.

Traffic Analysis

Network design always begins with traffic requirements. **Traffic analysis** asks how much traffic must flow over each of the network's many individual transmission links.

> *Traffic analysis asks how much traffic must flow over each of the network's many individual transmission links.*

Two-Site Analysis Figure 3-9 shows a trivial traffic analysis. A company only has two sites, A and B. They need to communicate at 1 Gbps. Obviously, the company needs a transmission link that can handle 1 Gbps.

Three-Site Analysis As soon as the number of sites grows beyond two, however, traffic analysis becomes challenging. Figure 3-10 shows a three-site traffic analysis. The figure shows that Site Q attaches to Site R, which attaches to Site S. There are two links: Link Q-R and Link R-S.

Site Q is west of Site R. Site S is east of Site R. Site Q needs to be able to communicate with Site R at 45 Mbps. Site R needs to be able to communicate with Site S at 2 Gbps. Site Q needs to be able to communicate with Site S at 300 Mbps.

Are you overwhelmed by the last paragraph? Anyone would be! In traffic analysis, it is critical to draw the picture. Figure 3-10 shows how the three sites are laid out and what links connect them.

After laying out the sites and links, it is straightforward to draw the three required traffic flows between each pair of sites. When you do that, you can see that some traffic flows are limited to a single transmission link, whereas traffic between Q and S must travel over both links.

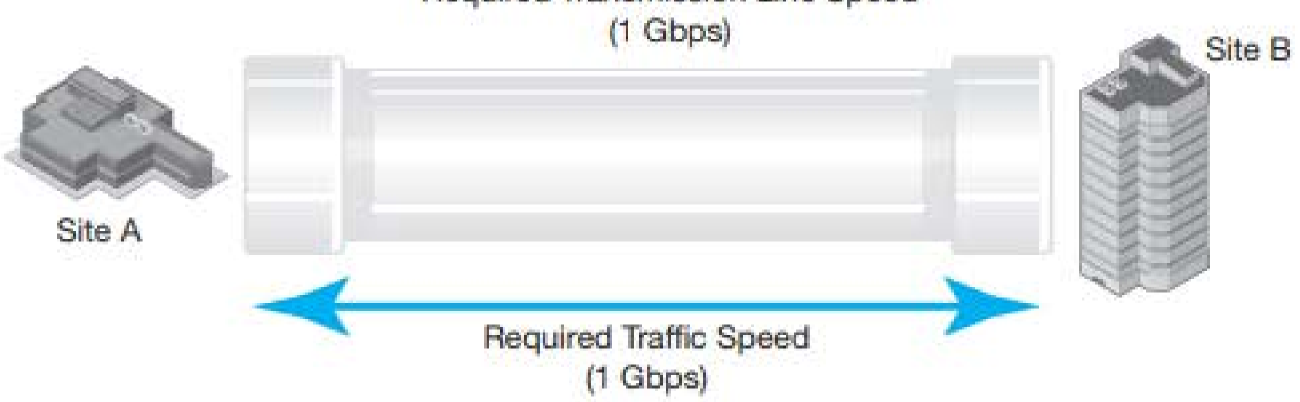

FIGURE 3-9 Two-Site Traffic Analysis

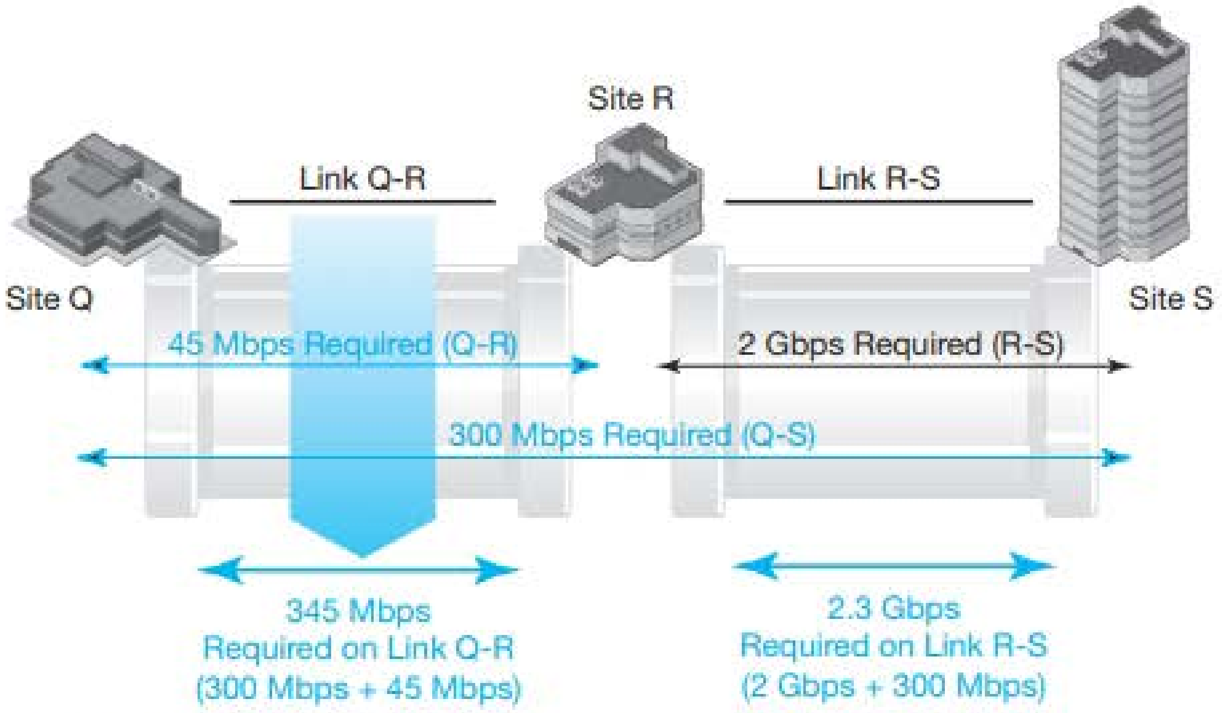

FIGURE 3-10 Three-Site Traffic Analysis

Now, you add up all the traffic flowing over each link.

- The link between Q and R must handle both Q–R traffic (45 Mbps) and Q–S traffic (300 Mbps). It does not handle any of the traffic between R and S, however. Consequently, Link Q-R must be able to handle 345 Mbps.
- Similarly, Link R-S must be able to handle R–S traffic (2 Gbps) and Q–S traffic (300 Mbps). This means that the transmission link between R and S must be able to handle 2.3 Gbps.

This problem can be handled well with the figure alone. However, Figure 3-11 shows a more general tabular way to do the analysis, a traffic table.

- The first column shows all possible combinations of pairs of sites that must communicate. The general rule is that if there are N sites, you will have N*(N-1)/2 site pairs. In this case, we have three sites, and 3*(3-1)/2 is 3. These three site pairs are Q-R, R-S, and Q-S. Each row shows traffic flowing between each possible site pair.
- The rows also show which links this traffic flows over. This requires you to see how traffic between each site pair will travel over the network. Again, you need the figure to let you understand the situation. Traffic between Q and R only flows over Link Q-R. Similarly, traffic flowing between R and S only flows over Link

Site Pairs / Link Pairs	Link Q-R	Link R-S	Remarks
Q-R Traffic	45 Mbps		Traffic goes over a single link.
R-S Traffic		2 Gbps	Traffic goes over a single link.
Q-S Traffic	300 Mbps	300 Mbps	Traffic goes over both links.
Total	345 Mbps	2.3 Gbps	Required speed for link.

FIGURE 3-11 Traffic Table for Figure 3-10

R-S. However, traffic flowing between Q and S needs to travel over two links: Q-R and R-S.

- Once you set up the table and enter the site-to-site traffic requirements for each row, you can simply total down each column to compute total traffic flowing over the link.

The traffic table looks like more work than just examining the figure and solving the totals visually. For more complex situations, however, the traffic table is the only approach that leaves you reasonably sane after the calculations. One reason the calculations are easy to do visually with the picture is that the sites are all laid out in a single line. As you might suspect, that doesn't often happen in the real world.

Four-Site Analysis Here is another, slightly more complex example. (Master the previous example before doing this one.) A company has offices in Honolulu, Seattle, Ogden, and Dublin, Ireland. There are three transmission links: Honolulu and Seattle, Seattle and Ogden, and Ogden and Dublin.

Seattle needs to communicate at 1 Gbps with each other site. Honolulu and Dublin only need to communicate with each other at 1 Mbps. Ogden and Dublin need to communicate at 2 Gbps. Honolulu and Ogden need to communicate at 10 Gbps. How much traffic will each transmission link have to carry? The analysis in Figure 3-12 shows how to calculate this.

- The first step, again, is to draw a picture showing the sites and transmission lines. Figure 3-12 shows this information at the top.
- Second, draw traffic requirements for each link between sites. The figure has also done that.
- Third, using the picture and no traffic table, find the total traffic that must flow over each link. Figure 3-12 does this for the link between Honolulu and Seattle. Note that the traffic flowing over the Honolulu-Seattle Link includes 1 Gbps flowing between Honolulu and Seattle, 1 Mbps (0.001 Gbps) between Honolulu and Dublin, and 10 Gbps between Honolulu and Ogden.
- Therefore, the total traffic flowing over the Honolulu-Seattle Link is 11.001 Gbps.

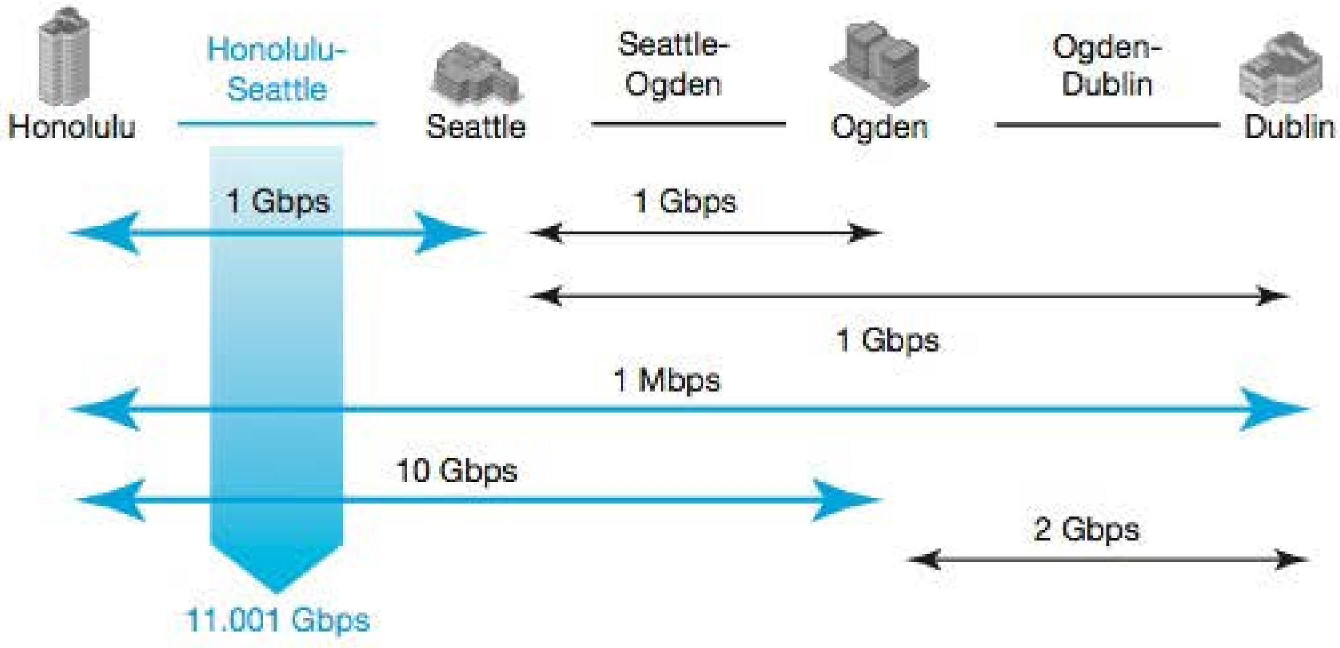

FIGURE 3-12 A Four-Site Traffic Analysis

Traffic Requirement for Each Pair of Sites	Transmission Link		
	Honolulu–Seattle	Seattle-Ogden	Ogden–Dublin
Honolulu-Seattle	1 Gbps		
Honolulu-Ogden	10 Gbps	10 Gbps	
Honolulu-Dublin			
Seattle-Ogden			
Seattle-Dublin			
Ogden–Dublin			

FIGURE 3-13 Traffic Table for Figure 3-12

The corresponding traffic table is shown in Figure 3-13. To make a traffic table, first note that there are three transmission links between sites to analyze. This means that there will be three columns of data. In addition, there are now four sites, so there are 4*(4-1)/2 possible combinations. This is six site pairs for traffic analysis. In Figure 3-13, the first two rows have been filled in. Your job is to fill in the rest of the table.

Test Your Understanding

7. a) Complete the traffic table in Figure 3-13. b) In Figure 3-12, add 392 Mbps of traffic for Seattle-Ogden communication. Using a picture like the one in the figure, show your work. c) Do it again with a traffic table. d) In Figure 3-10, remove the link between Q and R but add a link between Q and S. Using a picture, calculate requirements, showing your work. (Do not add the 392 Mbps in Part a.) e) Now use a traffic table to do the calculations. f) If you have 10 sites connected by seven transmission links, how many rows of traffic data will you have in your traffic table? g) How many columns?

Reliability Through Redundancy

Transmission lines sometimes fail, of course. The failure of even a single transmission line can wreak havoc within a network. Figure 3-14 repeats the four-site analysis in Figure 3-12. Actually, it repeats it twice.

Failed Transmission Line The top of Figure 3-14 shows what happens when the transmission line between Seattle and Ogden fails. Honolulu can still talk to Seattle, and Ogden can still talk to Dublin. However, Honolulu and Seattle cannot talk to Ogden or Dublin.

Adding Redundancy The lower half of the figure repeats the situation. This time, however, there is an extra transmission line. This line connects Honolulu and

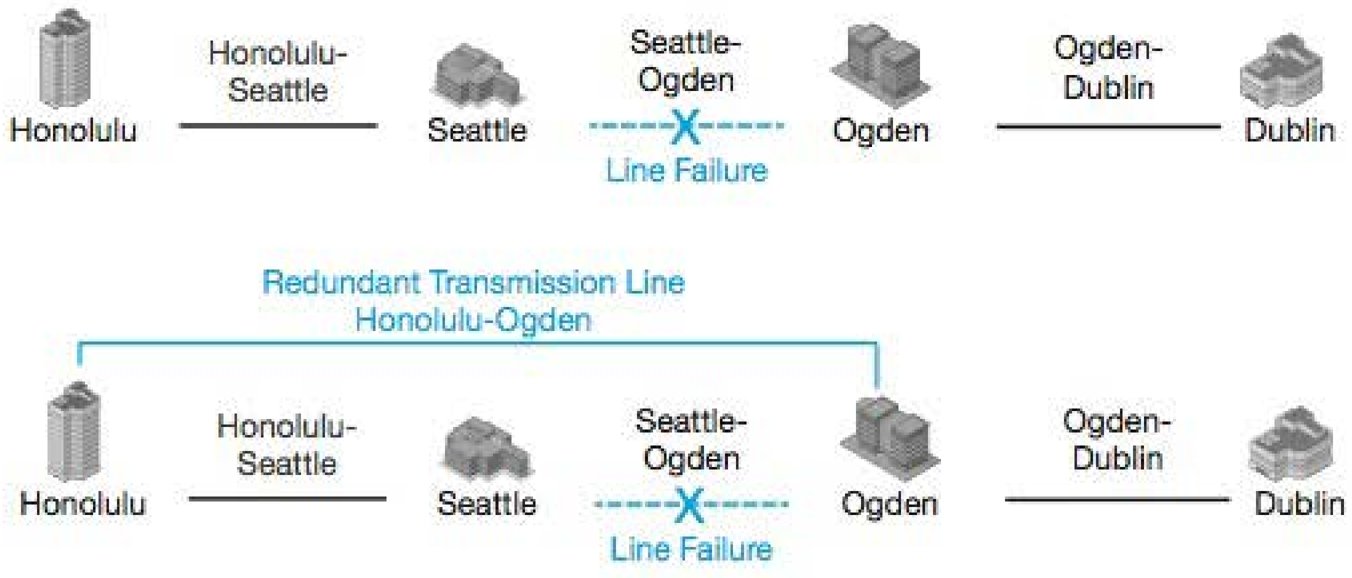

FIGURE 3-14 Reliability Through Redundancy

Ogden. Now, the failure in the Seattle-Ogden link is not a problem. Honolulu can still talk to Ogden directly, and through Ogden, Honolulu can still talk to Dublin. The added transmission link gives **redundancy**; if a link fails, there can still be transmission among at least some sites that would not be able to communicate after the failure.

Redundant transmission links ensure that if a link fails there can still be transmission among at least some sites that would not be able to communicate after the failure.

The situation for Seattle is a little more complicated. If Seattle wants to talk to Ogden, it must do so through Honolulu. Seattle first transmits to Honolulu, which lies in the opposite direction to Ogden. The traffic then passes from Honolulu through Ogden over the redundant transmission line.

With multiple ways for traffic to get between sites, computing required transmission line capacities by hand becomes impractical, growing impossible if there are many transmission lines and redundant transmission lines. Fortunately, software is available to do the calculations and make tradeoffs while still leaving spare capacity to achieve target levels of service if certain transmission lines failed.

Test Your Understanding

8. a) If the backup line in Figure 3-14 were to connect Seattle and Dublin instead of Honolulu and Ogden, list the path that traffic would take between Honolulu and Dublin if the line between Seattle and Ogden failed. b) Repeat for traffic between Seattle and Ogden. c) In the network in the lower half of Figure 3-14, what sites cannot communicate if the link between Ogden and Dublin fails?

Traffic Requirements versus Leased Lines

If the sites are miles apart, they will be connected by leased lines from a telephone carrier. Leased lines are point-to-point links between pairs of sites. They are "always on," so they are always available. Figure 3-15 shows the most common leased line speeds in

Line	Transmission Speed
T1	1.544 Mbps
T3	44.7 Mbps
OC-3	155.58 Mbps
OC-12	622.08 Mbps
OC-48	2,488 Mbps
OC-192	9,953 Mbps

FIGURE 3-15 Leased Line Speeds in the United States

the United States.[5] If you need a transmission speed of 30 Mbps between two sites, you cannot lease a 30 Mbps line. You would need a T3 line, with a speed of 44.7 Mbps.

Here is another example. In Figure 3-10, the link between Q and R needs to have a capacity of 345 Mbps. An OC-3 leased line of 155.58 Mbps would be too slow, so the company would need an OC-12 line running at 622.08 Mbps. This is a lot more capacity than the link needs, but there is nothing between OC-3 and OC-12.

In turn, Link R–S in Figure 3-10 requires a speed of 2.3 Gbps. In this case, it is easy to see that an OC-48 line at 2.488 Gbps will do the job with little wasted capacity. (However, when things are so close, it would be wise to ask if 188 Mbps is enough room for growth. Line requirements are based on growth forecasts for several years. Forecasts are never perfect.)

A very similar situation exists within local area networks that use Ethernet. We will see in Chapter 5 that Ethernet transmission also comes at a limited number of standard speeds.

Test Your Understanding

9. a) What leased line do you need if you have a capacity requirement for 2 Gbps? b) For 500 Mbps? c) For the situation in Figure 3-9, what leased line would the link require? d) Repeat for Figure 3-12. Do this for *all links* regardless of whether their capacity is shown in the figures.

Momentary Traffic Peaks

Traffic volume varies randomly (Figure 3-16). Consequently, there will *inevitably* be occasional **momentary traffic peaks** that exceed capacity. These only persist for milliseconds or a second or two, but they can be disruptive. Traffic will be delayed, creating latency. Some traffic may even be discarded because switches and routers have only

[5] To give more flexibility at the low end of speeds, one can purchase fractional T1 and T3 lines that give a fraction of the speed for a fraction of the cost. (The cost fraction is always bigger than the speed fraction, of course.) Fractional offerings vary by carrier. For example, the lowest fractional T1 line might be 128 kbps for one carrier, whereas for another it might be 256 kbps. Some might not offer fractional T1 services at all because demand for services below 1.544 Mbps is small today.

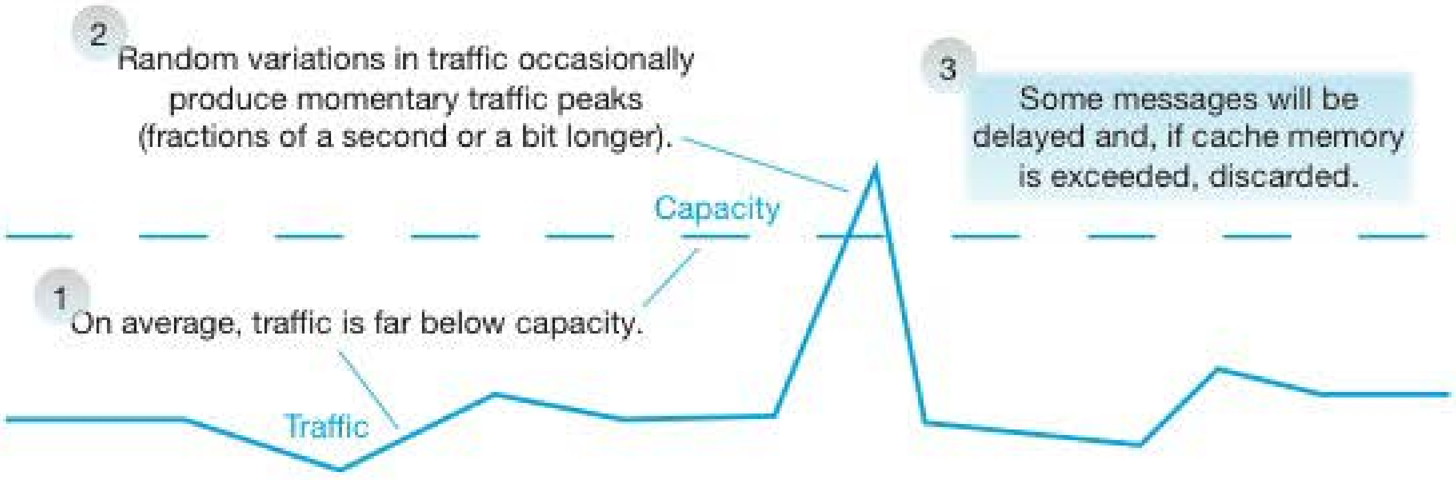

FIGURE 3-16 Momentary Traffic Peaks

a small amount of memory to store delayed messages. Once this "cache memory" is exceeded, frames or packets need to be dropped.

Adding More Capacity Figure 3-17 shows three techniques for addressing momentary traffic peaks. The first is to add more capacity. Ideally, one would add enough more capacity to eliminate momentary traffic peaks entirely. Given the nature of randomness, however, momentary traffic peaks will still occur, but they will be rarer and far shorter in duration. Adding more capacity is expensive in terms of transmission

Lack of Capacity	Amelioration	Description	Considerations
Momentary (milliseconds to a few seconds)	Add more capacity	Momentary traffic peaks will become extremely rare and brief	Expensive in terms of transmission cost. But requires no ongoing management labor.
	Prioritize traffic	Send higher-priority traffic through first.	Delay-intolerant high-priority traffic such as voice gets through immediately. Delay tolerant traffic such as e-mail will get lower priority, so if it is delayed briefly, harm is minimal. Requires ongoing management labor.
	Give QoS guarantees	QoS guaranteed capacity for certain traffic	Like reserved seating in a sports stadium. Traffic with a QoS guarantee will absolutely get through, up to the amount of capacity reserved. Other traffic only gets what is left over—even if the guaranteed traffic is not using its capacity.

FIGURE 3-17 Addressing Momentary Traffic Peaks

facilities, but it adds no ongoing management labor. Given the cost of labor, this is often a good tradeoff.

Priority A second approach to dealing with momentary traffic peaks is to assign a **priority level** to frames or packets, based on their tolerance for latency and loss.

- VoIP is extremely latency intolerant. Any noticeable delay will compromise the user experience substantially. It should be given very high priority.
- On the other hand, e-mail can easily tolerate a delay of several seconds. Consequently, e-mail gets low priority because a delay of a few seconds is not a problem in e-mail.

All commercial switches and routers in corporations come with the ability to use priority, so priority does not increase capital expense. Priority makes momentary traffic peaks tolerable to all types of uses, unless the peak is quite long. On the negative side, assigning priority to different applications and managing priority on switches and routers requires considerable ongoing management labor, which is expensive.

Momentary traffic peaks can be addressed by assigning a priority level to frames or packets, based on their tolerance for latency and loss.

Quality-of-Service Guarantees An extreme approach is to give **QoS guarantees** to certain traffic flows such as VoIP. Regardless of momentary traffic peaks, this traffic will always get through. It is like having season ticket seats for a sports team. To provide QoS guarantees, the company must allocate **reserved capacity** on each switch, router, and transmission line. This is great for traffic flows with QoS guarantees. However, it means that all other traffic only gets what is left over, even if the reserved capacity is not being used.

QoS guarantees reserve capacity for certain traffic flows such as VoIP. Regardless of momentary traffic peaks, this traffic will always get through.

Traffic Shaping The three coping mechanisms in Figure 3-17 only deal with ways to handle traffic *after it has entered the network*. A more fundamental way to deal with congestion is to limit what traffic *enters the network in the first place*. As Figure 3-18 shows, this is called **traffic shaping**. Traffic enters the network through an **edge router**, that is, a router at the edge of the network. This edge router has an **access control list (ACL)** that specifies what to do with different kinds of traffic. To implement this ACL, the edge router has the ability to recognize the types of applications that have generated the traffic.

Some applications are approved. These might include e-mail, database, web browsing, and other normal business applications. Approved applications are permitted to enter the network.

Other applications are *forbidden* by the access control list. These are simply blocked from entering the network. Still other applications may have some utility. An example might be YouTube. However, they cannot be permitted to take up much of the network's

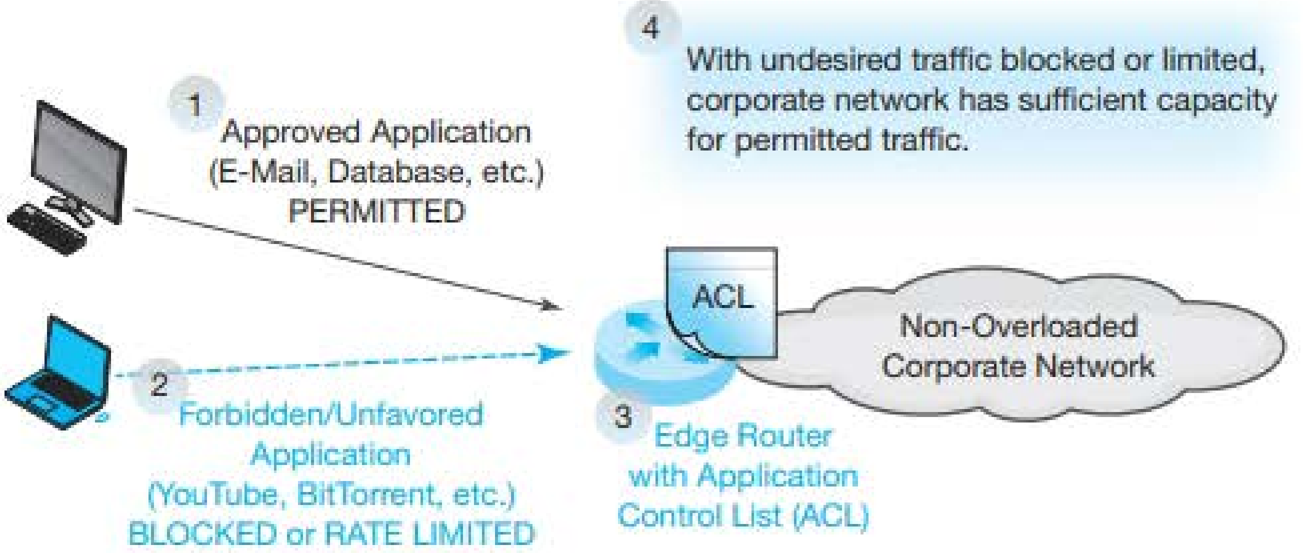

FIGURE 3-18 Traffic Shaping

capacity. These *unfavored* applications are **rate-limited**, meaning that they are limited to a certain small percentage of the network's traffic.

Test Your Understanding

10. a) Distinguish between chronic lack of capacity and momentary traffic peaks. b) How long do momentary traffic peaks last? c) What two problems do they create? d) What three choices do you have for reducing the impact of delays on latency-intollerant traffic? e) What is the advantage of each compared to the others? f) Compared to e-mail and VoIP, what priority would you give to network control messages sent to switches and routers? (The answer is not in the text.) Explain your reasoning. g) Is traffic shaping done before or after the traffic enters the network? h) What two choices does traffic shaping present for forbidden or undesirable traffic submitted to the network? i) If there is a chronic lack of capacity, which of the mechanisms described in these sections can help? j) What must be done if this is not possible or not sufficient?

CENTRALIZED NETWORK MANAGEMENT

Given the complexity of networks, network managers use **network management programs** to reduce their work. These programs allow managers working in centralized **network operation centers (NOCs)** to comprehend (and change) what is going on throughout their networks. Figure 3-19 shows the basic functions of an NOC.

Network Visibility The most important goal of network management is **network visibility**—the ability of the network manager to see what is going on throughout the network. This makes problem diagnosis possible even when the network has thousands of devices spread nationally or even worldwide. It also provides an understanding of network traffic trends and errors needed for planning. To achieve network visibility, every device must frequently send information about its configuration, traffic data, and error data.

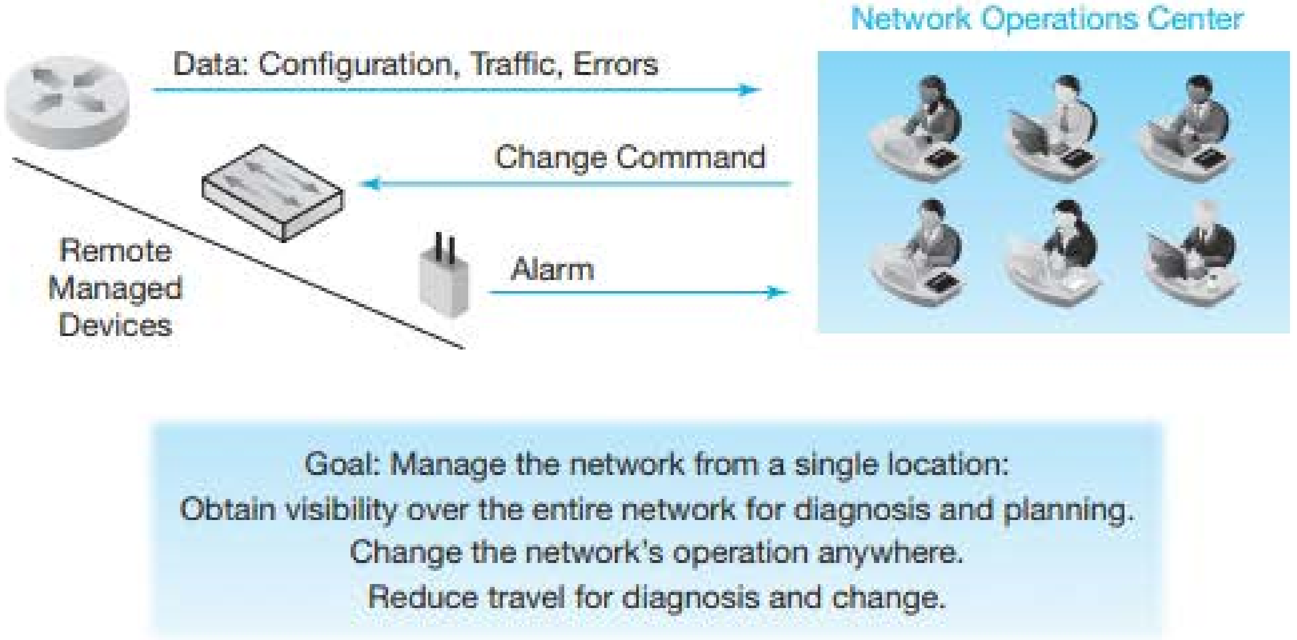
Network Operations Center
Data: Configuration, Traffic, Errors
Change Command
Alarm
Remote
Managed
Devices

Goal: Manage the network from a single location:
Obtain visibility over the entire network for diagnosis and planning.
Change the network's operation anywhere.
Reduce travel for diagnosis and change.

FIGURE 3-19 Centralized Network Management at a Network Operations Center (NOC)

Network visibility is the ability of the network manager to see what is going on throughout the network.

Network management tools are critical in large and distributed networks. If a network administrator had to travel to each device and transmission line to collect operating data for diagnosis when problems occurred or because he or she wished to optimize the network, the cost would be prohibitive.

Network management tools in general are costly to purchase and require considerable labor to operate, but they reduce labor expenses more than they cost. Beyond that, these tools enable network administrators to fix problems far more rapidly and to quickly make network changes that would be prohibitively difficult otherwise.

Sending Commands Of equal importance, network management programs allow network administrators to send commands to individual devices to change the way they operate. For example, the network administrator can tell a device to test the operation of a specific port and report its results. The administrator can also tell a device to turn off a port to change network traffic patterns or to shut down a malfunctioning port. This allows the network administrator to route traffic around congestion, to turn off expensive transmission lines when they are not needed, and to do many other things that affect overall network operation.

Alarms Individual devices sometimes take the initiative in communication. If they detect something wrong or at least suspicious, they can send messages called alarms to the network management software. Alarms get the network management software's attention and provide as much information on the situation as possible.

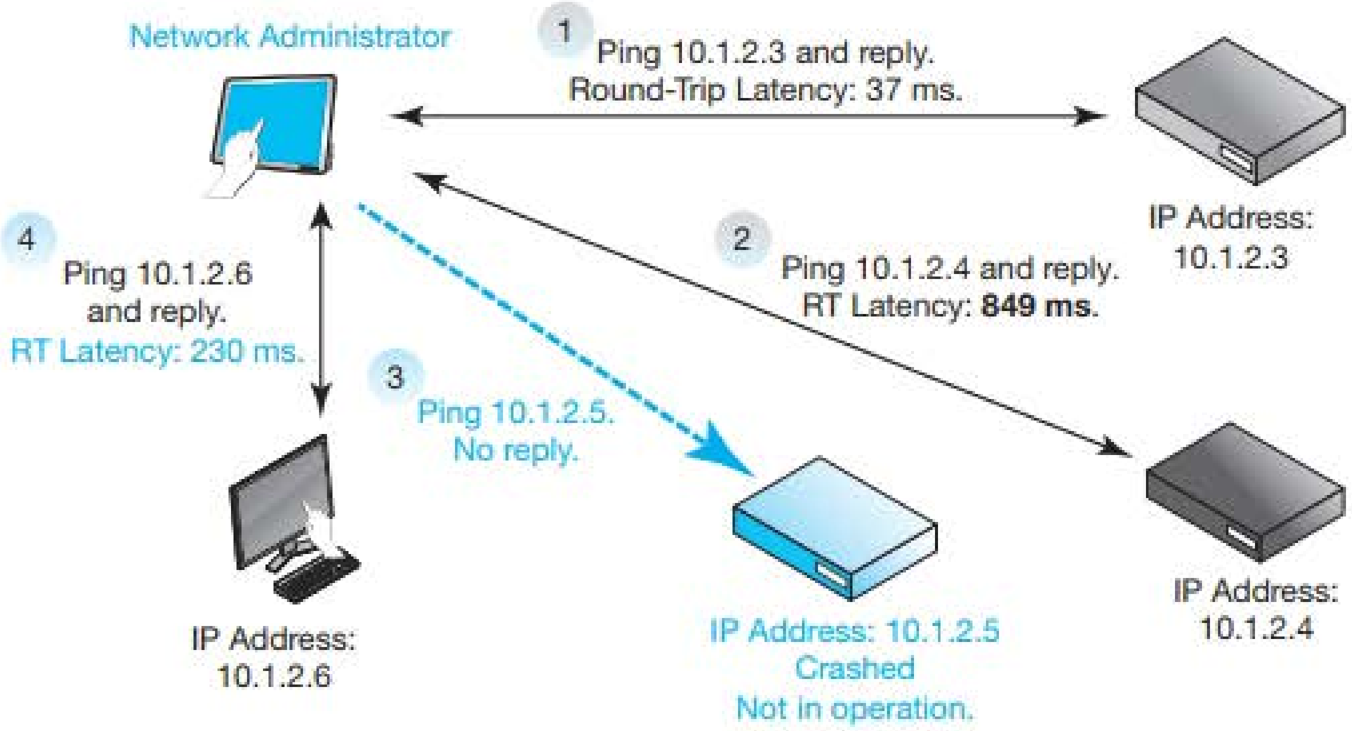

FIGURE 3-20 Ping

Test Your Understanding

11. a) What do network visibility tools allow a manager to do? b) Do they cost more money than they save? Explain.

Ping

The oldest network visibility tool is the **Ping** command available in all operating systems. If a network is having problems, a network administrator can simply Ping a wide range of IP addresses in the company. When a host receives a Ping, it should send back a reply. If it replies, it is reachable. If it does not, there is a problem. In Figure 3-20, host 10.1.2.5 does not respond to the Ping, signaling that it is either down or unreachable due to another problem in the network.

Reachability By analyzing which hosts and routers respond or do not respond, then drawing the unreachable devices on a map, the administrator is likely to be able to see a pattern that indicates the root cause of the problem. Of course, manually Pinging a wide range of IP addresses could take a prohibitive amount of time. Fortunately, there are many programs that Ping a range of IP addresses and portray the results.

Latency Problems Even if a host responds, there may still be a problem. Ping also reports the **round-trip latency** between the transmission of the Ping and the reception of the response. If the round-trip latency is substantial, there may be communication problems that need to be solved. This appears to be the case with Host 10.1.2.4, which has a two-way latency of 849 ms.

Round-trip latency is the time between sending a message and getting a response.

Traceroute

A related network visibility tool, **Traceroute**, gives you more granularity by reporting the round-trip latency for each hop between routers along the route. This can help you determine where a latency problem lies. Figure 3-21 shows how.

The first column shows the sequence number of each router along the route. There are 17 routers along the way, followed by the destination host. (Traceroute will actually show you the names of the routers and the host instead of just a number.) This will often help you identify who owns a particular router that seems to be causing problems.

The second column gives the round-trip latency to each particular device. For example, 18 has the value 670. This means that the round-trip latency to the destination host is 670 ms. This latency is about two-thirds of a second.

The third column shows how much latency each data link adds to the transmission time. In this figure, the jump to Router 12 adds the most to latency. The latency to Router 11 is 38 ms. The latency to Router 12 is 560 ms. Therefore, the data link between

Router	Round-Trip Latency (ms)	Difference
1	1	N.A.
2	7	6
3	7	0
4	7	0
5	8	1
6	10	2
7	12	2
8	13	1
9	29	16
10	34	5
11	38	4
12	560	522
13	563	3
14	567	4
15	590	23
16	603	13
17	620	17
18	670	50
Total	**670**	**N.A.**

FIGURE 3-21 Traceroute

Router 11 and Router 12 adds 522 seconds of round-trip latency—far more than any other jump between routers. This may indicate a problem.

Test Your Understanding

12. a) If you Ping a host and it does not respond, what can you conclude? b) What *two* things does Ping tell you about a host that replies? c) What types of latency do Ping and Traceroute give you? d) If a router causes problems, how can you diagnose this with Ping? e) Distinguish between Ping and Traceroute. f) In Figure 3-21, what jump causes the *second* most latency?

The Simple Network Management Protocol (SNMP)

Ping and Traceroute can tell you if a host is reachable and, if so, the latency in reaching that host. This is useful information, but it is extremely limited. For example, they do not let you query Router 12 in Figure 3-21 to look for indications of a problem. Full network management programs do.

Network management programs are built by many different vendors. So are routers, switches, access points, firewalls, and other network devices. A standard to govern their communication and what data they collect is necessary. This standard exists. It is the **Simple Network Management Protocol (SNMP)**, which Figure 3-22 illustrates. In the network operations center, a computer runs a program called the **SNMP manager**. The manager communicates with a large number of **managed devices**, such as switches, routers, access points, firewalls, servers, and PCs.

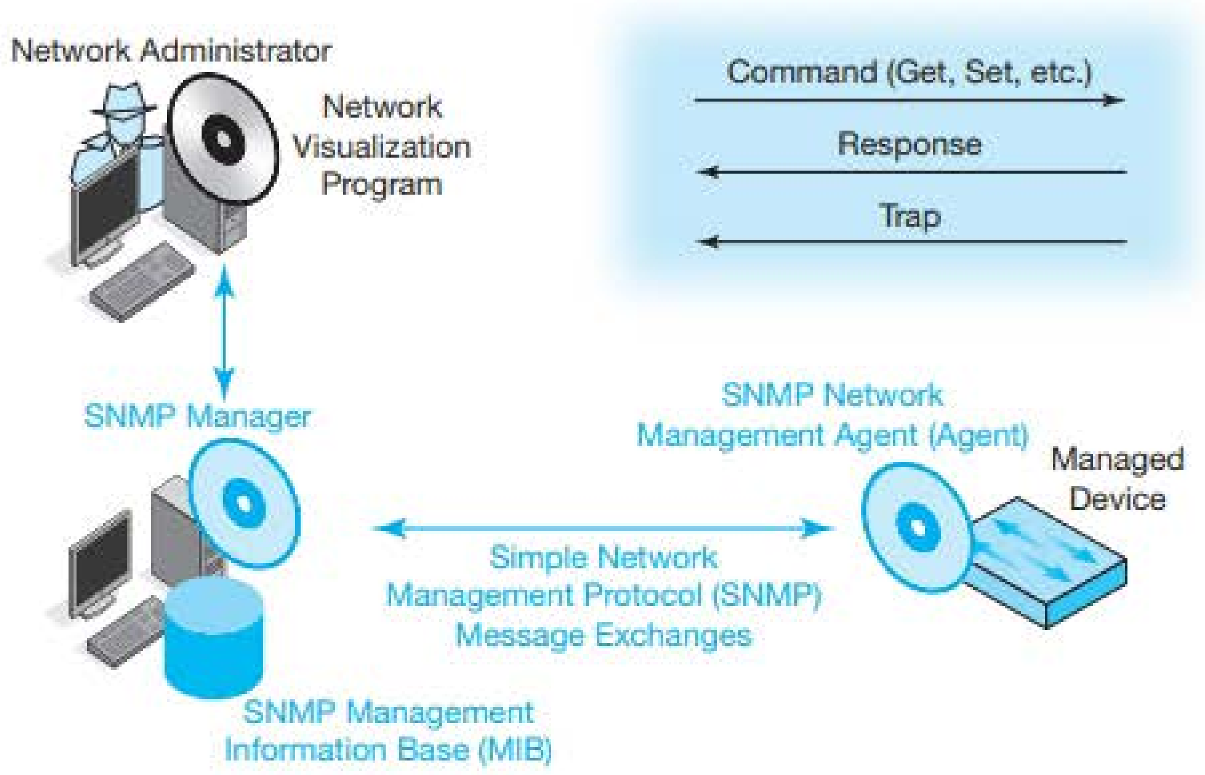
Network Administrator
Network Visualization Program
Command (Get, Set, etc.)
Response
Trap
SNMP Manager
SNMP Network Management Agent (Agent)
Managed Device
Simple Network Management Protocol (SNMP) Message Exchanges
SNMP Management Information Base (MIB)

FIGURE 3-22 Simple Network Management Protocol (SNMP)

SNMP Agents The manager does not talk directly with the managed devices. Rather, each managed device has an **SNMP agent**, which is hardware, software, or both. The manager talks only to the agent. To give an analogy, recording stars have agents who negotiate contracts with studios and performance events.

SNMP Get Commands The network operations center constantly collects data from the managed devices using **SNMP Get** commands. The Get command specifies what data is to be provided. A response message delivers the data.

SNMP Management Information Base (MIB) When the data from a device arrives, the manager stores it in its SNMP **management information base (MIB)**. Data in the MIB allows network administrators to understand the traffic flowing through the network. It is the basis for network visibility.

Databases have schema, which specify the particular types of data they can store. There are different **SNMP schemas** for different types of devices, such as Ethernet switches and routers. The MIB stores them separately but can integrate some of their information. The schemas for each type of device are extremely rich and specific. For example, if a router is failing and then rebooting frequently, this can make it difficult to diagnose. Ping commands, for instance, may not indicate a problem if they arrive when it is up. However, one element of a router's SNMP schema is the time since last reboot. A short time since last reboot will indicate an intermittent failure. The MIB also can contain data for various types of errors to help pinpoint a problem.

SNMP Set Commands and Security In addition, the manager can send **SNMP Set** commands to managed devices. The agents of these devices send response messages to confirm that they have made the changes.

Many companies do not use Set because an attacker can do infinite mayhem within the network unless Set messages are highly secured. Forgoing Set is safe, but it is extremely costly because it requires a great deal of more costly network management labor. Companies that have good security, however, can use Set safely. This is an example of how good security can be a money-saving enabler, not simply a cost.

SNMP Trap Earlier, we saw that managed devices can send alarms if they detect an issue. SNMP calls these alarms **SNMP traps**.

Network Visualization Program There is one more program in the figure—the **network visualization program**. This program takes results from the MIB and interprets the data to display results in maps, find root causes for problems, and do other tasks. Note that this functionality is *not* included in the Simple Network Management Protocol. SNMP simply collects the data in a way that network visualization programs can use. This lack of specification allows network visualization program vendors to innovate without being constrained by standards. The network visualization program also can issue commands to the SNMP manager to query a managed device for data or to change the way a device operates.

Test Your Understanding

13. a) List the main elements in SNMP. b) Does the SNMP manager communicate directly with the managed device? Explain. c) Distinguish between SNMP Get and SNMP Set commands. d) Where does the SNMP manager store the information it receives from Get commands? e) Why can good security save money in network management? f) What kind of message can agents initiate? g) What is the relationship between the network visualization program and the SNMP manager?

Automation

Many other network management chores can be automated to reduce the amount of work that network administrators need to spend on minutia. For example, many routers are given a standard corporate configuration when they are installed. This greatly reduces the time needed to configure each router and reduces configuration errors. However, it is possible to create a standard configuration, store it, and simply download it onto new routers. These router configurations must then be adjusted for their devices' particular roles in the network.

Test Your Understanding

14. a) Why is the automation of network management tasks important? b) Why are standard configurations desirable?

SOFTWARE-DEFINED NETWORKING (SDN)

We close with a new trend that may redefine network management. **Software-Defined Networking (SDN)** is a radically new way to configure switches, routers, access points, and other devices. Even a medium-sized bank has hundreds of switches and hundreds of routers. As just noted, many companies have "standard configurations" that are downloaded to new switches and routers. Afterward, however, network engineers may have to modify this configuration when conditions change, and these changes have become increasingly frequent, making the traditional way we have configured devices a limiting factor in **control agility**—the ability to rapidly change how the network operates when conditions change.

Control agility is the ability to rapidly change how the network operates when conditions change.

Traditional Configuration and Its Discontents

Figure 3-23 shows that network administrators have traditionally modified each device's configuration file individually and manually. Often, the administrator must travel to the switch or router and work with it physically. Sometimes, it can be reached from the network operations center via the network, but even this only saves some of the required time to reconfigure individual routers, switches, access points, firewalls, and other network devices.

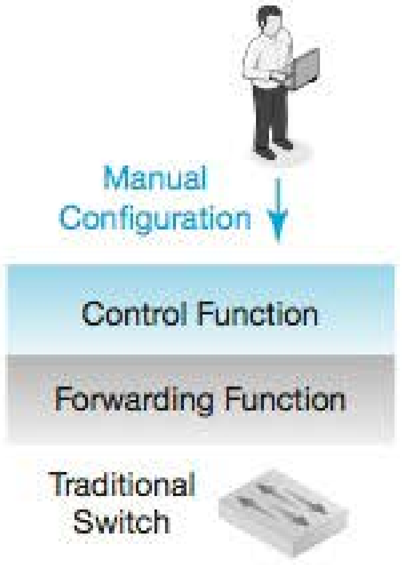

Traditionally, each device had to be configured individually and manually.

Configuration is done on the control function, which is the policy-based rule set that determines how forwarding should be done.

The forwarding function executes forwarding decisions for individual frames and packets. It forwards based on the control function's rule set.

Traditionally, the rule set rarely changed, so changing devices manually was inconvenient but unfrequent.

Today, control function rule set changes are frequent. A new way to configure many devices is needed.

FIGURE 3-23 Traditional Individual Manual Device Configuration and Reconfiguration

The Forwarding Function Figure 3-23 shows that the managed device normally has two functions. The obvious one is the **forwarding function**, which consists of switching arriving frames back out or routing arriving packets back out. This forwarding function consumes nearly all of a switch's or router's resources.

The forwarding function consists of switching arriving frames back out or routing arriving packets back out.

The Control Function There is also the control function, which is not as obvious. The **control function** is the policy-based reconfiguring of network devices. For example, switches forward frames based on information in a switching table, which tells the switch which port to use to send the arriving frame back out. Part of the control function on a switch is creating these switching tables. More broadly, the control function embraces configuration in general. One element in a router's configuration is whether the routing parameter is set to on or off. If it is off, the router will not route packets. One of the first steps in diagnosing a problem is determining whether the router is routing. In general, the configuration of switches, routers, access points, firewalls, and other devices is a complex process. Forwarding consumes most of a device's time, but control consumes most of the administrator's time in dealing with the devices.

The control function is the policy-based reconfiguring of network devices.

Forwarding consumes most of a device's time, but control consumes most of the administrator's time in dealing with the devices.

Policy-Based Configuration Note that the control function uses **policy-based configuration**. Policies are broad mandates for network management.

- For a company that manages servers for several different customers, a policy might be that the servers of one customer must never be able to communicate with the servers of other customers.
- Another might be that expensive transmission links between the company's sites should be shut down at night so that less expensive transmission links are used to handle the lower traffic at night.

Policies typically must be applied to multiple devices and the management of multiple transmission links. This means that multiple devices must be configured when policies change.

Initially, it made sense to do configuration on individual switches and routers. This way, a firm with only one or two switches or routers did not have to use a sophisticated set of tools to manage the routers. As the number of network devices grew, however, the cost of making control function changes grew proportionally. So did the cost of devices, which depend considerably on control system requirements. Figure 3-24 emphasizes the burdensome nature of doing configuration on each device.

Control Agility Traditionally, the control function did not change much or frequently. Having to change it manually and perhaps locally was not too much of a burden. That is no longer the case. For example, consider what happens in a cloud computer server farm with thousands of servers managed by a cloud service provider such as Amazon.com. Every time a new server is added, the cloud service provider's configuration policy may mandate changes so that other corporate users of the server farm cannot get access to the new customer server. This will require reconfiguring each router's Access Control List immediately to specify changes in which servers which customers can reach from their servers. Figuring out how to change the ACLs of hundreds or thousands of routers is time consuming. Actually making the changes adds to the required time. As noted earlier, companies today need control agility—the ability to take policy-based control actions rapidly when conditions require change.

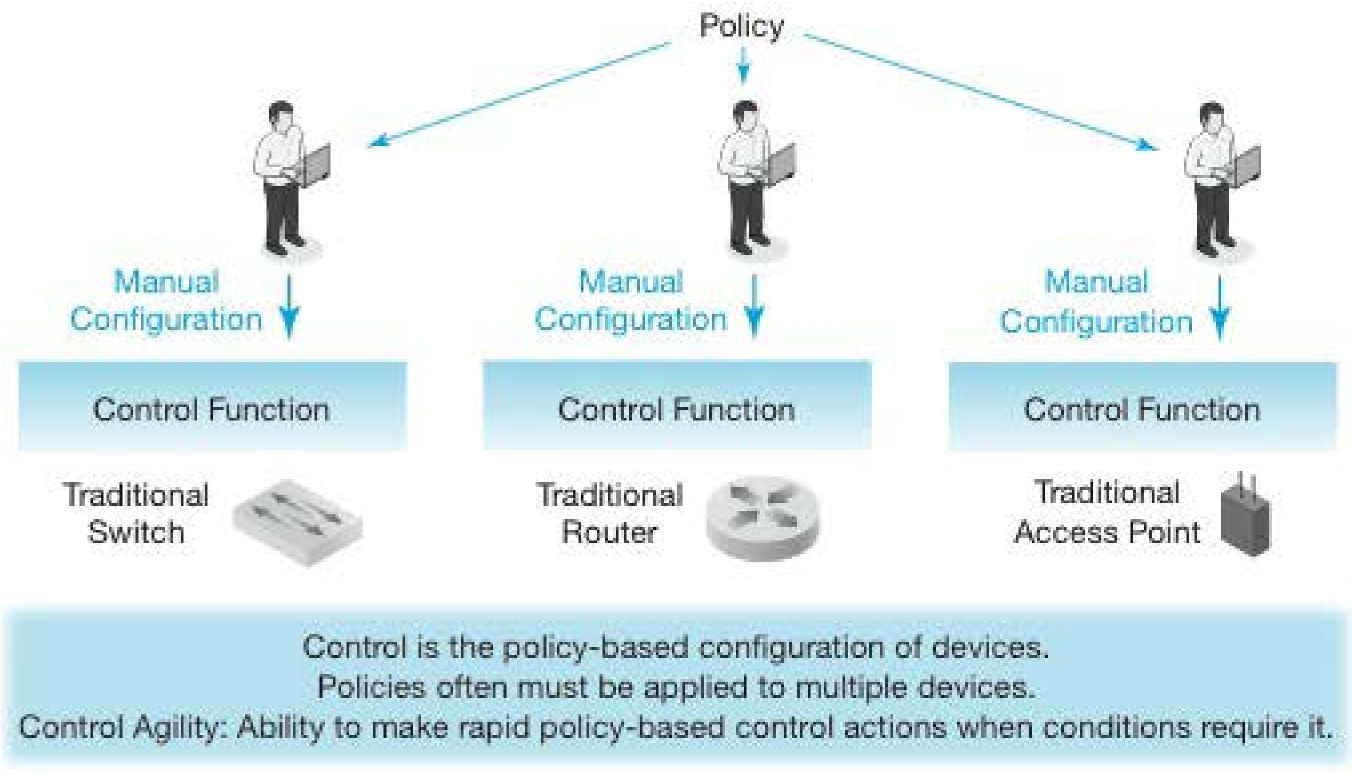

Control is the policy-based configuration of devices.
Policies often must be applied to multiple devices.
Control Agility: Ability to make rapid policy-based control actions when conditions require it.

FIGURE 3-24 Traditional Configuration for Multiple Devices

How often does reconfiguration happen today? To give an example, the customers of Amazon Web Services frequently add servers as their needs increase and release servers when their needs decrease. For example, Netflix uses Amazon Web Services for its recommendation software to suggest television shows and movies to individuals based on what other users with similar viewing profiles have watched. The CPU cycles and storage needed by the software spike during prime time viewing hours in the evening and plunge at night. Netflix changes the number of servers it uses for this function several times a day. Each change requires reconfiguration. If Netflix adds 300 servers that were used by other customers minutes or hours earlier, many switches and routers need to be reconfigured to change which servers each customer can reach.

Software-Defined Networking Operation

Figure 3-25 shows how Software-Defined Networking changes the picture. Most obviously, each device is stripped of its control function. This considerably reduces the cost of a switch, router, or other device. The figure shows that the control function is instead centralized in an **SDN controller**. When changes are made on the controller, new forwarding rule sets are sent to the affected devices.

With centralization, the administrator simply gives a high-level command. This might be to add a new customer's server. This broad command is converted into ACL rules appropriate for individual devices, and these control rules are sent to the affected devices.

Figure 3-26 illustrates that this requires **SDN application programs** that run on the SDN controller. These applications allow the administrator to do complex tasks with SDN doing remarkably little work compared to manual configuration. For example, one of the applications in the figure is a traffic segmentation program that manages changes in access when new servers are added or released by a customer.

Figure 3-26 shows **APIs**, which are **application program interfaces**. APIs are standardized interfaces between programs. For example, the SDN controller has a set of APIs that application programs use to talk to it in a standardized way. This means that

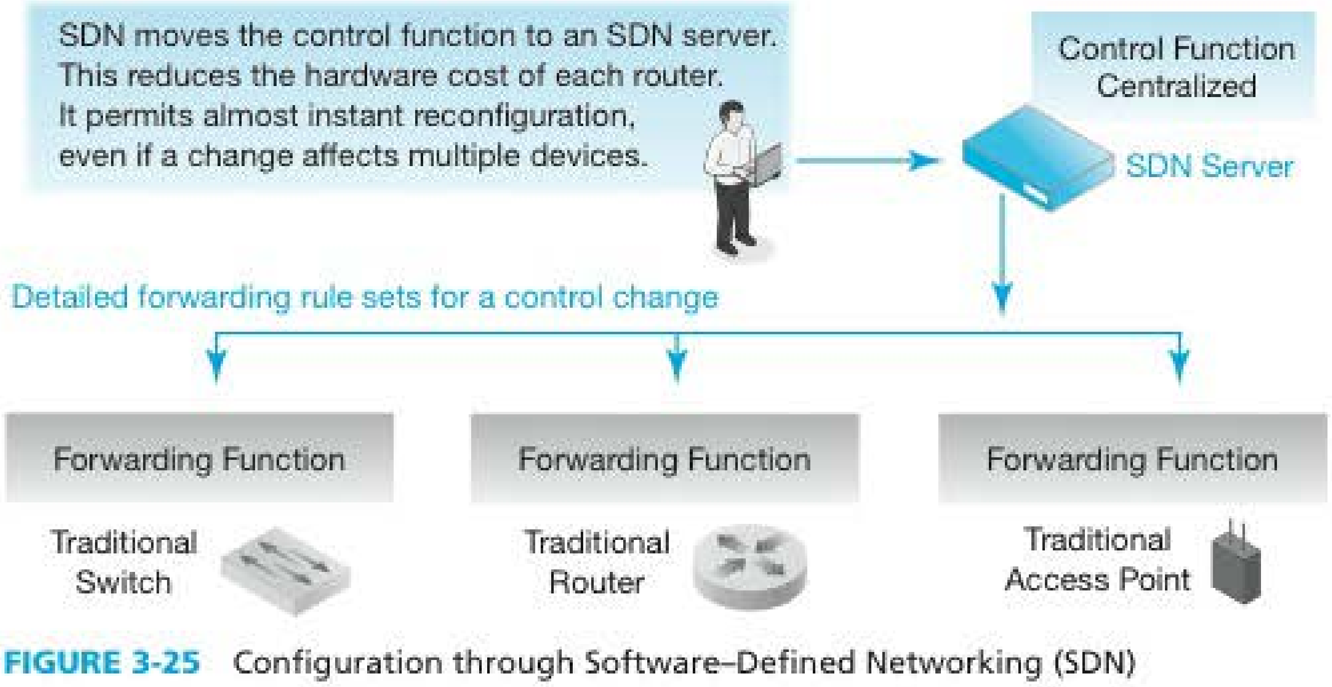
SDN moves the control function to an SDN server. This reduces the hardware cost of each router. It permits almost instant reconfiguration, even if a change affects multiple devices.

FIGURE 3-25 Configuration through Software–Defined Networking (SDN)

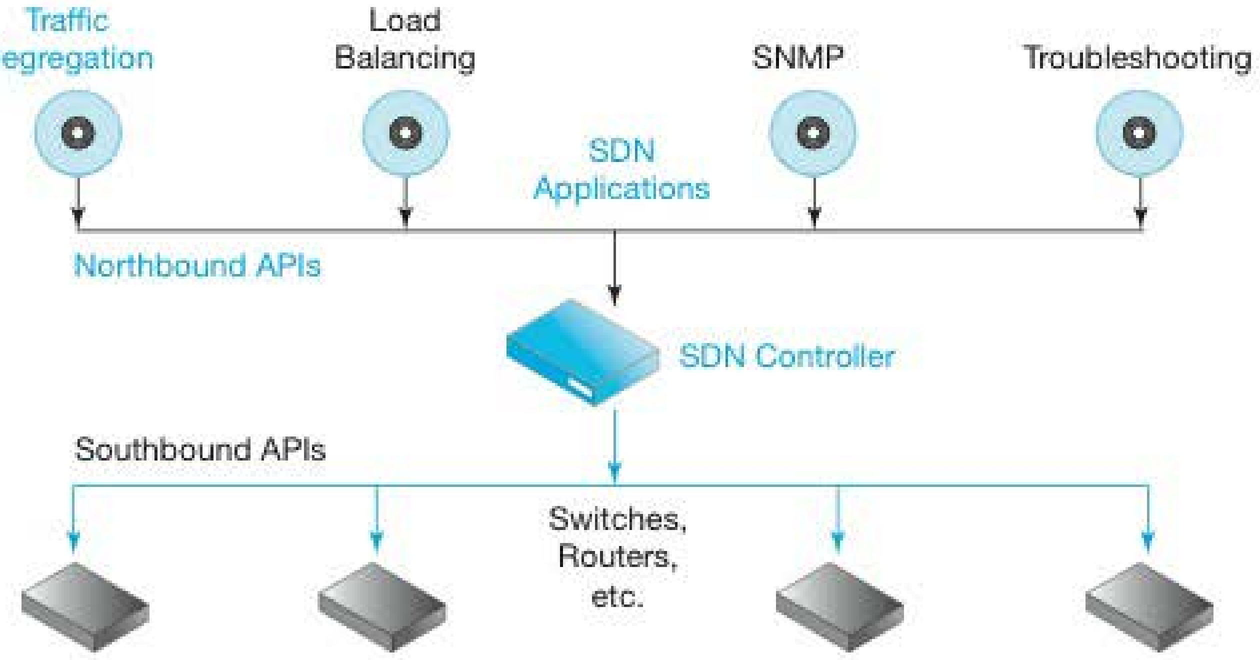

FIGURE 3-26 SDN Applications and Application Program Interfaces (APIs)

any software company can write application programs to work with the SDN controller. It also means that these applications can run on SDN controllers from other vendors if these vendors follow the same APIs. This is a classic example of how standards enable technology competition.

Figure 3-26 shows that SDN controllers have *two* sets of APIs. **Northbound APIs** govern communication between application programs and the controller. **Southbound APIs** are different. They standardize communication between the SDN controller and the individual switches, routers, and other devices they configure.

Unfortunately, today there are several SDN API families. Complicating matters greatly, Cisco Systems, which dominates sales in routers and many switch categories, has its own approach that is designed to keep its routers and switches expensive by not stripping the control function out of them but offering many of the benefits of SDN.

Although this unsettled market environment is making most firms wary of SDN, some firms have implemented it on a large scale and have reaped extensive benefits. One is Amazon Web Services, which hosts servers for other companies. AWS has even created and used its own SDN routers, which are far cheaper than commercial routers because their control functions have been removed. Software-Defined Networking has done more than reduce costs, however. It has brought extreme control agility to Amazon's vast network of hosted servers. This agility has allowed AWS to implement network changes that could not have been imagined before Software-Defined Networking.

Test Your Understanding

15. a) What are the benefits of Software-Defined Networking? b) Distinguish between the control function and the forwarding function. c) Where was the control function placed traditionally? d) Where is it placed in SDN? e) What do northbound APIs connect? f) What do southbound APIs connect? g) Which type of API must router and switch designers support? h) Why are applications necessary for SDN to be successful? (The answer is not in the text.)

END-OF-CHAPTER QUESTIONS

Thought Questions

3-1. Your home is connected to the Internet. You get to create SLAs that the ISP must follow. Being reasonable, write SLAs you would like to have for the following things: a) Write an SLA for speed. b) Write an SLA for availability. c) Write an SLA for latency. **Do not just say what each SLA should include. Actually write the SLAs as the ISP would write them in the form of specific guarantees. Failure to do this will result in a substantial grading penalty.**

3-2. Redo the analysis in Figure 3-12. Remove the link between Ogden and Seattle but add a link between Seattle and Dublin. On each link, what traffic capacity will be needed, and what leased line would you select for it? Use a traffic table to do the analysis.

3-3. Figure 3-27 shows four sites communicating. Each site needs to communicate with each other site at 2 Mbps, except for Paris. Paris needs to communicate with each other site at 5 Gbps. Create a traffic table and solve it. (Partial Answer: For London–Munich, the total traffic is 5.004 Gbps.)

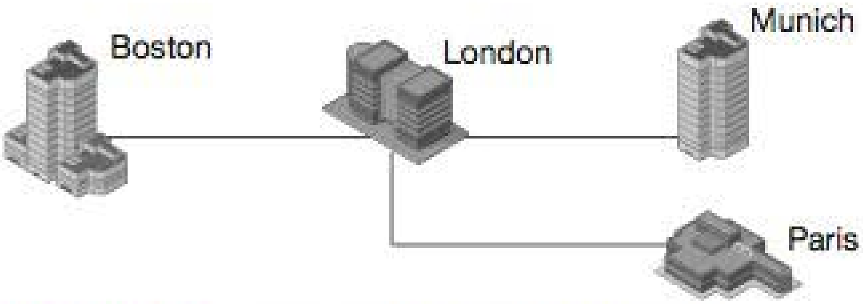

FIGURE 3-27 Layout for Thought Question 3-3

3-4. Figure 3-28 has data from a Traceroute analysis. a) Add a third column showing the change in latency between the router in that row and the router in the preceding row. (Check figure: The change for Router 2 from Router 3 is zero.) Shade the row for any large latency problem or problems you find. b) For the first problem you find, state whether the problem might be in the router in the row, in the previous router, or something else.

Router	Latency (ms)
1	1
2	7
3	7
4	7
5	8
6	10
7	12
8	13
9	29
10	34
11	38
12	52
13	75
14	567
15	590
16	603
17	1002
18	1017

FIGURE 3-28 Data for Thought Question 3-4

3-5. a) Why *must* the forwarding function remain on the network device? (The answer is not in the text.) b) What might be holding back SDN in many firms?

Perspective Questions

3-6. What was the most surprising thing you learned in this chapter?

3-7. What was the most difficult part of this chapter for you?

Chapter **3a**

Hands-On: Microsoft Office Visio

LEARNING OBJECTIVE

By the end of this chapter, you should be able to:

- Create a simple Visio diagram.

WHAT IS VISIO?

Microsoft Office Visio is a drawing program. The professional version has special symbols for drawing network diagrams. Visio is widely used by network professionals to visualize networks they are designing.

USING VISIO

Visio is part of the Microsoft Office family. Installing Visio is like installing any other Office product.

Figure 3a-1 shows how to start a Visio drawing. Of course, this begins by selecting File and then New. In the figure, Network has been selected for the type of drawing. Detailed Network Diagram has been selected.

As Figure 3a-2 shows, this brings up a window with a canvas on which you can drag shapes. In the figure, the shape of a generic server has been dragged onto the screen. As you can see, many other network diagramming shapes can be dragged onto the screen.

After you have added the devices you need, it is time to begin showing how they are connected. As Figure 3a-3 shows, there is a connector icon at the top of the screen.

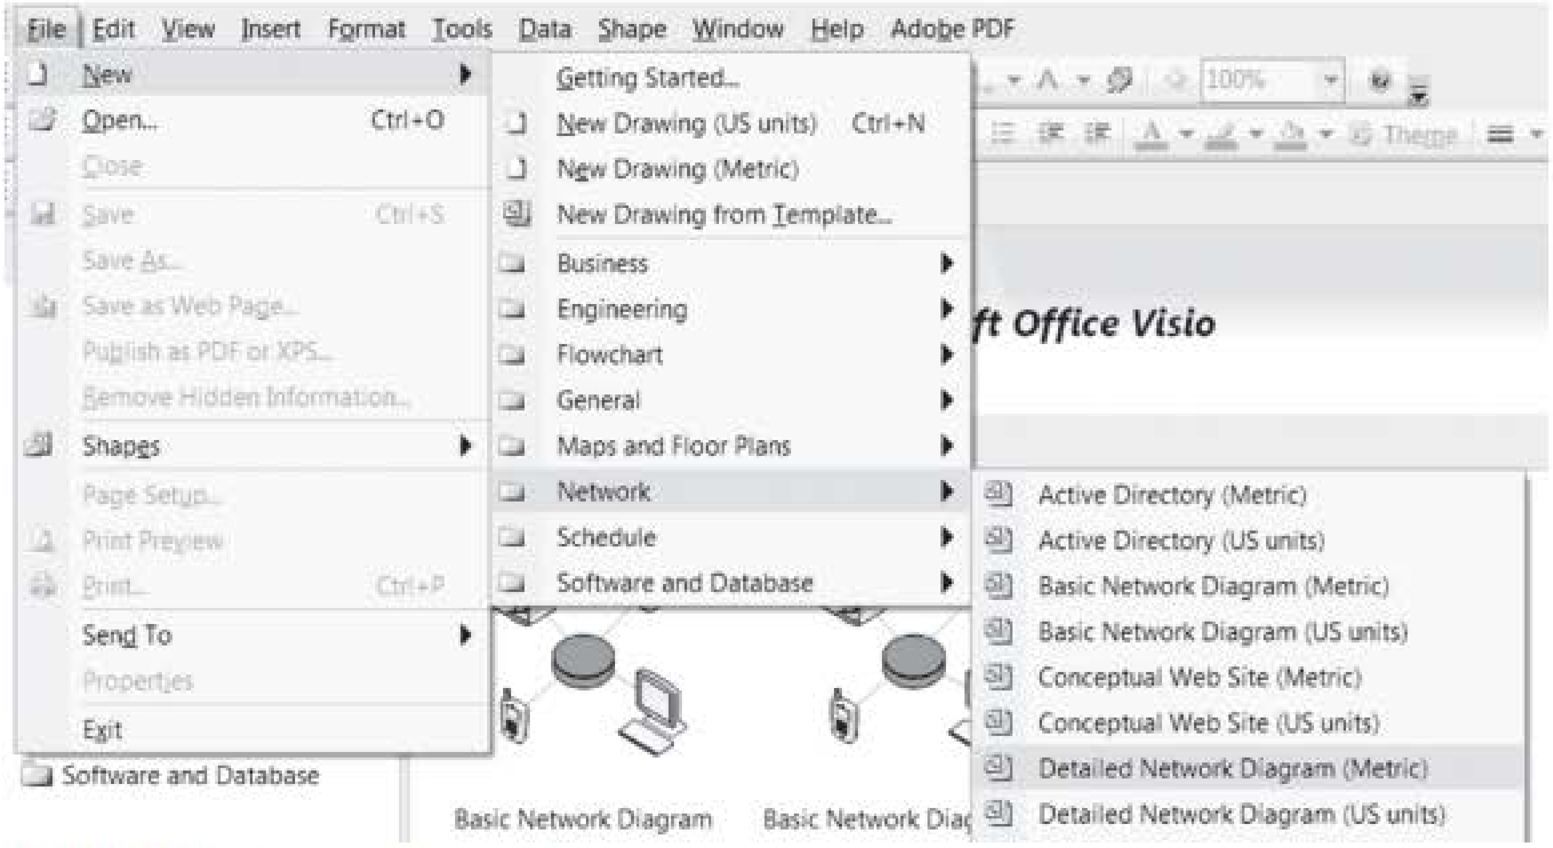

FIGURE 3a-1 Starting a Visio Drawing

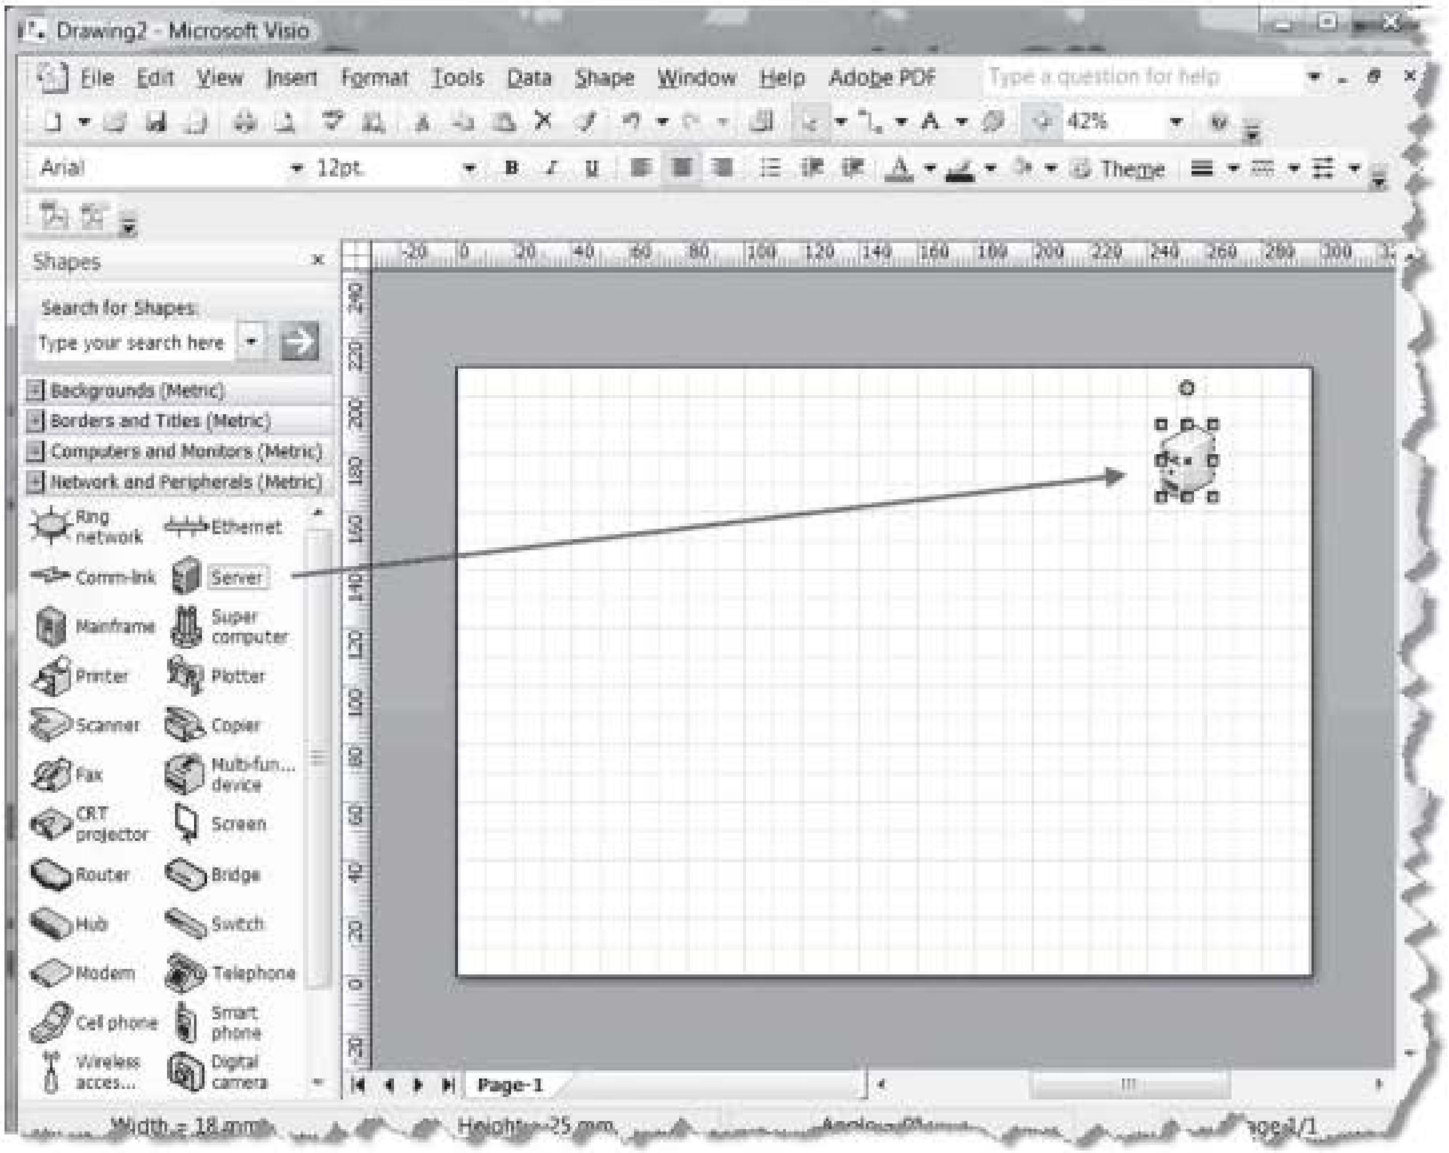

FIGURE 3a-2 Drawing Canvas with Icon Being Dragged

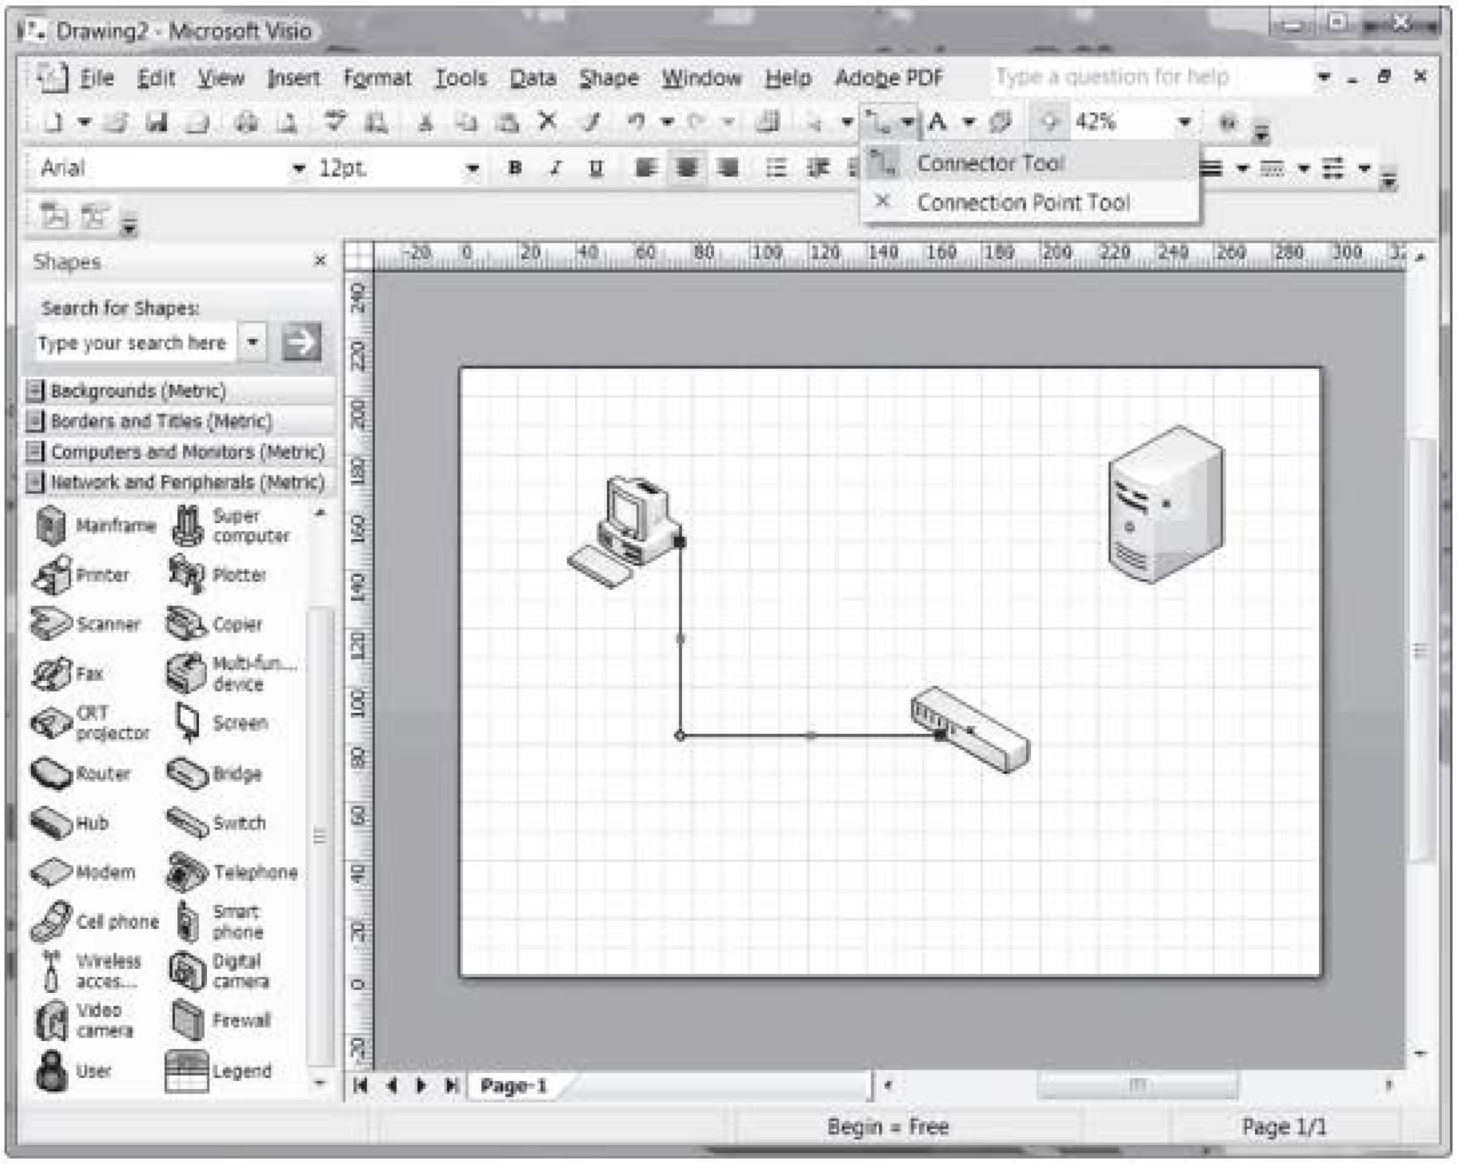

FIGURE 3a-3 Adding Connections

Select the connector tool. Then drag between the two icons to connect them. After you have connected them, try dragging one of the connected devices. You will see that the connectors move with them.

Not shown in the figure, you can double-click on an icon. This adds text below the icon. Visio is not fussy about preventing lines from overlapping text. Overall, Visio diagrams are easy to create but not extremely pretty.

HANDS-ON EXERCISES

In Microsoft Office Visio, create something like the drawing in Figure 3a-4.

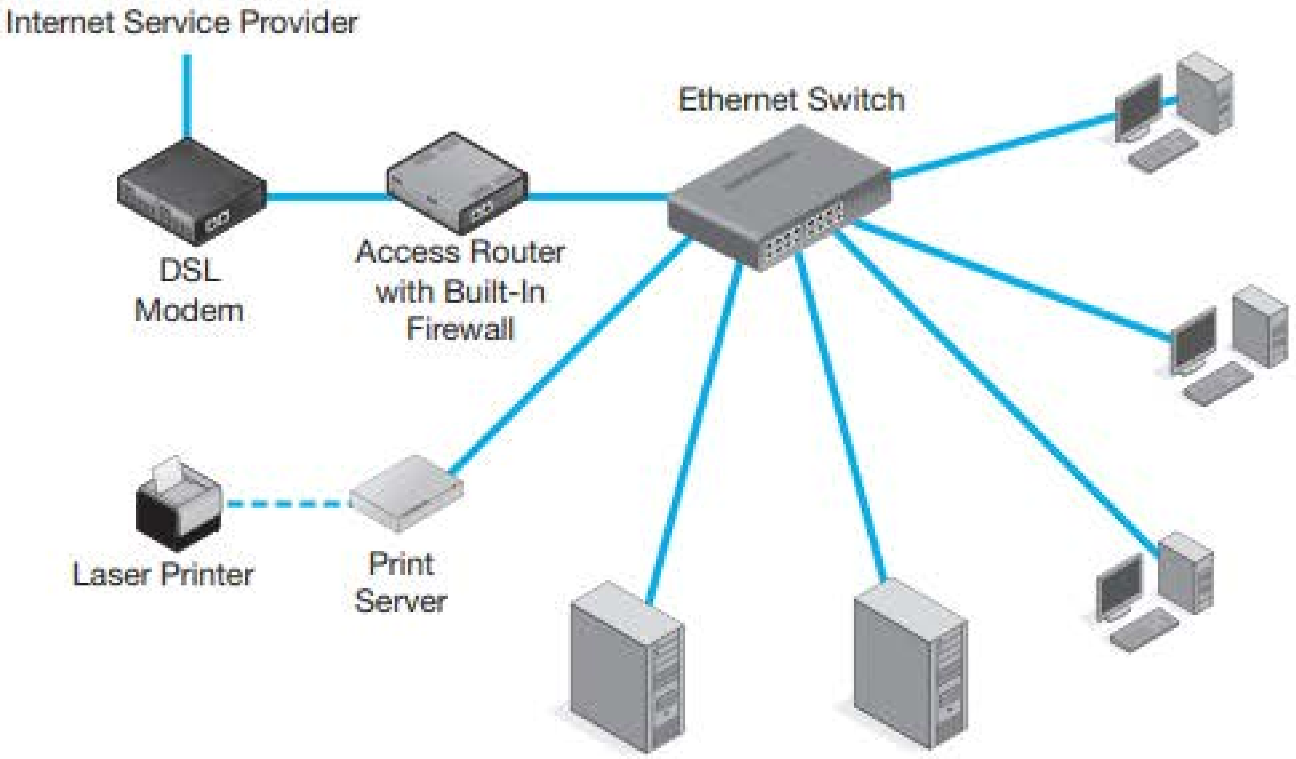

FIGURE 3a-4 Sample Drawing

This page intentionally left blank

Chapter **4**

Network Security

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Describe the threat environment, including types of attacks and types of attackers.
- Explain how to protect dialogues by cryptography, including encryption for confidentiality, electronic signatures, and host-to-host virtual private networks (VPNs).
- Evaluate alternative authentication mechanisms, including passwords, smart cards, biometrics, digital certificate authentication, and two-factor authentication.
- Describe firewall protection, including stateful packet inspection, next-generation firewalls, and related intrusion prevention systems.
- Describe the role of antivirus protection.

THE TARGET BREACH

Near the end of the 2013 holiday season, Target announced that thieves had stolen data from 40 million credit cards scanned at Target stores in preceding weeks.[1] The attackers had done this by downloading malware to nearly all point-of-sale (POS) systems in American Target stores. It captured magnetic stripe information and sent it to data thieves.[2] Target initially did not reveal the fact that thieves were already committing fraud with the stolen card data. A month later, Target announced that a separate but

[1] Alastair Jamieson and Erin McClam, "Millions of Target Customers' Credit, Debit Card Accounts May Be Hit by Data Breach," NBC News, December 19, 2013. http://www.nbcnews.com/business/consumer/millions-target-customers-credit-debit-card-accounts-may-be-hit-f2D11775203.

[2] Jaikumar Vijayan, "Security Firm IDs Malware Used in Target Attack," Computerworld.com. http://www.computerworld.com/s/article/9245491/Security_firm_IDs_malware_used_in_Target_attack.

related theft had occurred during roughly the same period. Attackers had stolen personal information on roughly 70 million Target customers.[3] Consumers were shocked and worried by these thefts. Many canceled their charge cards and demanded new cards from their banks. Within weeks, a barrage of lawsuits began.

The POS Attack

Target released little information about either compromise, but analysts gradually constructed a likely picture of how the credit card number theft had occurred. News reports naturally focused on the POS systems, but the theft involved a complex series of steps inside and outside Target. Figure 4-1 shows the most important of these steps.

The theft did not begin with a direct attack on Target. Rather, it began with an attack on Fazio Mechanical Services, which provided services to Target in the mid-Atlantic region.[4] Fazio had credentials on a vendor server that handled electronic billing and other matters. The attackers probably sent an employee a spear phishing e-mail that tricked the employee into loading malware onto his or her machine. The malware captured the Fazio credentials on the vendor server and sent it back to the attackers. The attackers then used these credentials to get access to the vendor server. From this initial foothold, they were able to move more deeply into the Target network.

Now inside the Target network, thieves installed POS malware, which they had purchased from an online crimeware shop, to a malware download server within

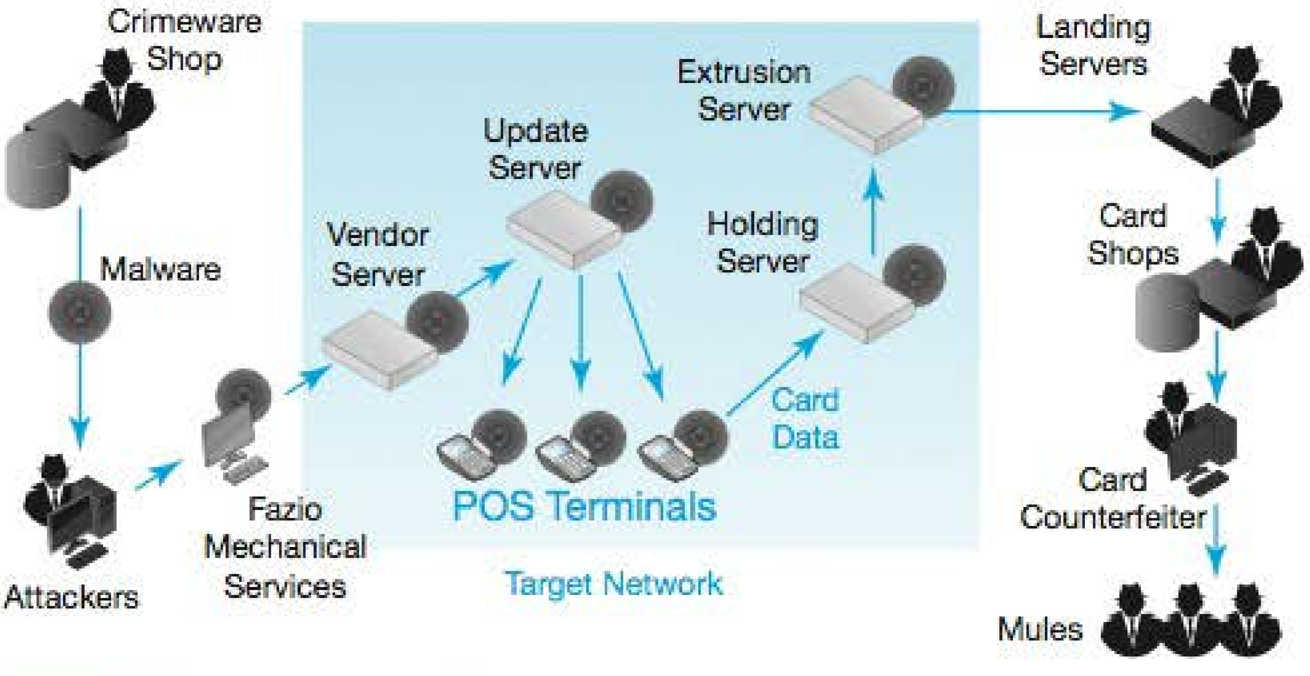

FIGURE 4-1 The Target Breach

[3] Target, "Target Provides Update on Data Breach and Financial Performance," January 10, 2014. http://pressroom.target.com/news/target-provides-update-on-data-breach-and-financial-performance.

[4] Fazio Mechanical Services, "Statement on Target Data Breach," accessed April 26, 2014. http://faziomechanical.com/Target-Breach-Statement.pdf; Brian Krebs, "Target Hackers Broke in Via HVAC Company," KrebsOnSecurity.com, February 5, 2014. http://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company.

Target. There is suspicion that the thieves actually took over Target's internal server that downloaded updates to the POS systems.[5] In any case, the malware was downloaded to a few POS systems initially and then to nearly all Target POS systems in the United States.[6]

The malware was a variant of the BlackPOS malware that had been in existence for almost a year.[7] It was readily available at online crimeware shops for about $2,000.[8] The attackers probably modified the software to attack Target's specific POS terminals.[9] They probably also modified it so that existing antivirus programs would not detect it.[10] It is common for hackers to maintain small server farms to test malware against popular antivirus products.

The malware collected magnetic stripe data from every card swiped at the terminal. This occurred before the information was encrypted and sent over the Target network. Most sources called the malware a *RAM scraper*, indicating that it sent everything in the POS terminal's memory to the attackers.[11] Actually, it was more selective, stealing only data on the magnetic stripes of swiped cards.[12] This included the primary account number, the expiration date, the name of the card owner, and optional information. Stolen data *did not* include the card security code, which is a 3-digit or 4-digit number printed on a credit card. Companies ask you for this number when you cannot present your card physically. For credit cards, there was sufficient information on the magnetic stripe to create counterfeit credit cards. For debit cards, the theft included encrypted personal identification numbers (PINs), but there is no indication that these PINs were decoded.[13]

Data collected at the POS terminal went, as usual, to legitimate Target servers. However, the malware also sent the data to a compromised holding server where the data from all of the POS terminals was stored temporarily.[14] For data extrusion, the attackers compromised another server that would deliver the data to the

[5] Brian Krebs, "These Guys Battled BlackPOS at a Retailer," KrebsonSecurity.com, February 14, 2014. http://krebsonsecurity.com/2014/02/these-guys-battled-blackpos-at-a-retailer/.

[6] Krebs, "Target Hackers Broke in Via HVAC Company."

[7] Vijayan, "Security Firm IDs Malware Used in Target Attack."

[8] Ibid.

[9] Ibid.

[10] Ibid.

[11] Target, "Target Provides Update on Data Breach and Financial Performance," January 10, 2014. http://pressroom.target.com/news/target-provides-update-on-data-breach-and-financial-performance.

[12] Krebs, "These Guys Battled BlackPOS at a Retailer."

[13] Adam Greenberg, "Hackers Seek to Decrypt PIN Codes Likely Stolen in Target Breach," *SC Magazine*, January 8, 2014. http://www.scmagazine.com/hackers-seek-to-decrypt-pin-codes-likely-stolen-in-target-breach/article/328529/.

[14] Keith Jarvis and Jason Milletary, "Inside a Targeted Point-of-Sale Data Breach," Dell SecureWorks, January 24, 2014. http://krebsonsecurity.com/wp-content/uploads/2014/01/Inside-a-Targeted-Point-of-Sale-Data-Breach.pdf.

attackers outside the Target network.[15] This extrusion server pulled batches of card data sets from the holding server and transmitted them to landing servers in Russia, Brazil, Miami, and other locations.[16] The thieves could not conceal the Internet protocol (IP) addresses of the landing servers, so they probably moved the data quickly to other servers.

Now the attackers monetized their stolen data. They wholesaled batches of data to online *card shops* that then sold the data to counterfeiters. These card shops held stripe information in a searchable database. This allowed counterfeiters to purchase card stripe information selectively. For example, thieves know that using a credit card in a city that is not the owner's may result in a credit freeze. Consequently, card shops allowed customers to search by zip code. Counterfeiters also refined their purchases in other ways, based on such factors as whether the card had a high debt limit. Based on the characteristics of each card, counterfeiters paid from $20 to more than $100 per card. The first customers received a money-back guarantee that 100% of the card data was useable.[17] Over time, the guaranteed percentage fell, and prices declined.

The counterfeiters used the card data to create fake credit cards that looked legitimate down to the graphics used by individual banks. They then copied data from a single legitimate card onto the magnetic stripe of each counterfeit card. This allowed them to purchase high-end merchandise and then sell the merchandise to traditional fences. However, the counterfeiters did not make the purchases themselves. Instead, they hired a small corps of "mules" to make the actual purchases or take cash out of ATMs.

One thing is missing from the figure. The attacks needed to transmit control messages frequently into the Target network in order to compromise servers and take actions to direct actions on these servers during the attack. All of these messages had to go through Target's firewalls. Showing this information in Figure 4-1 would create an unintelligible spiderweb of arrows. However, it was critical for the attackers to maintain a hole in the victim's firewalls during the entire attack process.

Test Your Understanding

1. a) How did the attackers gain access to Target's network? b) List the internal Target servers the attackers compromised. c) How did the attackers exfiltrate the card data? d) List the criminal groups, besides the main attackers, who were involved in the overall process. e) What benefit did the attackers seek to obtain from their actions? f) Critique (positively or negatively) the fact that Target knew that fraud was already occurring with the stolen card data but did not reveal this when it announced the breach.

[15] Ibid.

[16] Brian Krebs, "Non-US Cards Used at Target Fetch Premium," KrebsonSecurity.com, December 13, 2014. http://krebsonsecurity.com/2013/12/non-us-cards-used-at-target-fetch-premium/.

[17] Brian Krebs, "Cards Stolen in Target Breach Flood Underground Markets," KrebsonSecurity.com, December 20, 2014. http://krebsonsecurity.com/2013/12/cards-stolen-in-target-breach-flood-underground-markets/.

Damages

It may take years to fully understand the damage from the Target breach. However, it is easy to identify victims. One was Target itself. In the period from the breach revelation to February 2014, Target sales fell 5.3% from the previous year, and profits fell 46%.[18] This profit decline was roughly $500 million. In addition, Target has probably paid out several hundred million dollars due to lawsuits brought by commercial and government organizations. The company's chief technical officer resigned fairly soon after the breach,[19] and the company's CEO resigned in May 2014.[20]

Consumers are protected against fraudulent credit card purchases—but only if they notify their credit card company quickly of fraudulent charges on their bills. Credit card companies will drop these transactions from bills. However, this process is time-consuming and frustrating. It sometimes even involves disagreements about whether charges are truly fraudulent. There is even more time lost if the consumer cancels the credit card and gets a new card to get peace of mind. Finally, the prospects of credit card fraud and identity fraud created psychological costs for many cardholders.

Surprisingly, banks and credit card processors usually do not lose money in the case of reported fraudulent purchases. Just as the customer does not pay them, banks and credit card processors do not pay the retail stores in which the fraudulent purchases were made. Beyond this, financial services companies face substantial costs in the replacement of compromised cards. However, they are likely to recover these costs successfully in lawsuits.

Fraud hits retailers the hardest. They rarely recover merchandise purchased fraudulently. However, there is one thing that physical retailers can do to reduce these losses. Counterfeiters normally only create a single card master from which all counterfeit cards in a batch are made. All counterfeit cards in the batch have the same printed name, credit card number, expiration date, and other information. The magnetic stripe data, however, will be specific to a single compromised credit card. This is why store clerks look at the last four digits of the card number on the physical credit card. If this is different from information on the magnetic stripe, the card is fraudulent.

Test Your Understanding

2. a) How was Target damaged by the breach? b) Were banks and credit card bureaus damaged by the breach? c) How were consumers damaged by the breach? d) How were retailers damaged by the breach? e) What can retailers do to defend themselves against counterfeit credit cards? f) What individual victim or group of individual victims suffered the most harm?

[18] "Target Profits Plunge 46% after Holiday Security Breach," BBC.com, February 26, 2014. http://www.bbc.com/news/business-26358556.

[19] Anne D'Innocenzio, "Target's Chief Information Officer Resigns," Associated Press, March 5, 2014. http://www.nytimes.com/2014/03/06/business/targets-chief-information-officer-resigns.html?_r=0.

[20] Clare O'Connor, "Target CEO Gregg Steinhafel Resigns in Data Breach Fallout," Forbes, May 5, 2014. http://www.forbes.com/sites/clareoconnor/2014/05/05/target-ceo-gregg-steinhafel-resigns-in-wake-of-data-breach-fallout/.

Perspective

The Target breach was not an isolated incident. Surveys have found that most firms suffer at least one compromise each year. Successful attacks are becoming ever more frequent, sophisticated, and damaging. In 2012, the director of the Federal Bureau of Investigation Robert Mueller made the following statement: "Terrorism remains the FBI's top priority. But in the not too distant future, we anticipate that the cyber threat will pose the number one threat to our country."[21] In 2014, the Center for Strategic and International Studies estimated global damage from cybercrime.[22] It concluded that cybercrime reduced the entire world's gross domestic product by almost 1%. Cybercrime is not a small or distant threat, and it is growing explosively. In 2015, British insurer Lloyds estimated that cybercrime was costing businesses $400 million a year.

> *"Terrorism remains the FBI's top priority. But in the not too distant future, we anticipate that the cyber threat will pose the number one threat to our country."*
>
> *Robert Mueller, Director of the FBI*

INTRODUCTION

Networks give us access to almost anything, anytime, anywhere. Unfortunately, they give the same access to criminals, national governments, terrorists, and just plain jerks. Wherever there has been opportunity, there has been crime and vandalism. Networks are no exception. Security is the snake in the network garden.

Network thinking focuses on software bugs and mechanical breakdowns. In contrast, security thinking must anticipate the actions of intelligent adversaries who will try many things to succeed and adapt to the defenses you put in place.

> *Network thinking focuses on software bugs and mechanical breakdowns. In contrast, security thinking must anticipate the actions of intelligent adversaries who will try many things to succeed and adapt to the defenses you put in place.*

Giving you even a broad view of security is too much for one chapter. The appendix looks more broadly at how to manage security as part of overall network management. As security expert Bruce Schneier has said in many of his writings, "Security is a process, not a product."

Test Your Understanding

3. How does security thinking differ from network thinking?

[21] Federal Bureau of Investigation, Speech by Robert S. Mueller III, Director, Federal Bureau of Investigation (Press release), RSA Cyber Security Conference, San Francisco, California, March 1, 2012.

[22] Center for Strategic and International Studies, "Net Losses: Estimating the Global Cost of Crime," June 2014. http://www.mcafee.com/us/resources/reports/rp-economic-impact-cybercrime2-summary.pdf.

TYPES OF ATTACKS

We begin by looking at the threat environment that corporations face. The **threat environment** consists of the types of attacks that companies face and the types of attackers who engage in these attacks. We begin by looking at *types of attacks*.

> *The threat environment consists of the types of attacks that companies face and the types of attackers who engage in these attacks.*

Malware Attacks

Malware is a generic name for evil software. It includes viruses, worms, Trojan horses, and other dangerous attack software. Malware attacks are the most frequent problems that companies face. Nearly every firm has one or more significant malware compromise each year.

> *Malware is any evil software.*

Test Your Understanding

4. a) What is malware? b) What are the most frequent types of attacks on companies?

Malware

- A general name for evil software

Vulnerabilities and Patches

- Vulnerabilities are security flaws in specific programs
- Vulnerabilities enable specific attacks against these programs to succeed
- Software vendors release patches to close vulnerabilities
- However, users do not always install patches promptly or at all, so continue to be vulnerable
- Also, zero-day attacks occur before the patch is released for the vulnerability

Social Engineering

- For when there is no vulnerability
- Trick the user into doing something that will compromise security, such as opening an e-mail attachment
- Phishing involves e-mail messages that appear to be legitimate to a group of people (e.g., customers of a particular bank)
- Spear phishing is aimed more selectively at individuals or a few individuals (more effective because it is personal)
- Going to a website and being tricked into downloading malware

FIGURE 4-2 Malware and Vulnerabilities

Vulnerabilities and Patches

Most types of malware can only succeed if a program under attack has a security vulnerability. A **vulnerability** is a flaw in a program that permits a specific attack or set of attacks to succeed against the program. Vulnerabilities are found frequently in popular application programs.[23]

A vulnerability is a flaw in a program that permits a specific attack or set of attacks against this program to succeed.

When a software vendor discovers a vulnerability, the company issues a **patch**, which is a small program designed to fix the security vulnerability. After patch installation, the program is safe from attacks based on that particular vulnerability. Too often, however, users fail to install patches, so their programs continue to be vulnerable. Even if they do install patches, they may delay, giving the attacker a long window of opportunity.

Of course, if attacks begin before the program vendor creates a patch (or even learns about the attack), then all attacks against vulnerable computers will succeed. A vulnerability-specific attack that occurs before a patch is available is called a **zero-day attack**. In such cases, there would be no signature to check for yet. On the security black market, well-funded adversaries can often purchase information that allows them to create zero-day attacks.

A vulnerability-specific attack that occurs before a patch is available is called a zero-day attack.

Test Your Understanding

5. a) What is a vulnerability? b) How can users eliminate vulnerabilities in their programs? c) What name do we give to attacks that occur before a patch is available?

Social Engineering: No Vulnerability Necessary

Even if the software being attacked has no vulnerabilities, attackers can succeed if they can get the user to take an action that compromises security. This is called **social engineering**. A prime example of social engineering is an e-mail phishing attack. A **phishing** attack pretends to be from a company the user does business with or from another seemingly trustworthy source. The text of the e-mail message is also convincing. Using HTML, it may look exactly like e-mail messages the source usually sends.

[23] A 2014 study by Cenzic found that 96% of all applications tested had at least one vulnerability. The median number of flaws per application was 14. Andy Patrizio, "Nearly All Apps Are Vulnerable in Some Way," NetworkWorld, March 3, 2014. http://www.networkworld.com/article/2226448/microsoft-subnet/nearly-all-apps-are-vulnerable-in-some-way–report-says.html.

Social engineering consists of tricking the user into taking an action that compromises security.

An e-mail phishing attack involves sending a message that pretends to be from a company the user does business with or from another seemingly trustworthy source. However, it is really from an attacker.

Spear phishing is even more specific. The attacker personalizes the e-mail message to a particular person, such as the chief executive officer of the company. Spear phishing e-mails are even more convincing because they typically appear to come from a specific trusted person and contain information that only that person is likely to know. For example, it may mention specific projects or locations while traveling.

In some cases, a social engineering attack entices the user to click on a link that will take the victim to a site that asks the person to download a program to view a particular attachment. This downloaded program will actually be malware. In other cases, the e-mail may contain the malware directly, in the form of an attachment.

Test Your Understanding

6. a) What kind of attack may succeed against a system with no technological vulnerabilities? b) What is the goal of social engineering? c) Distinguish between phishing and spear phishing attacks.

Types of Malware

There are many types of malware. Figure 4-3 lists three common types.

Viruses The first common type of malware is the virus. A **virus** attaches itself to a legitimate program, just as a human virus attaches itself to a person's cells. If the victim runs the program, the virus will spread to other programs on the computer.

Typically, the virus will then try to propagate to other computers. It cannot do this directly. Instead, it propagates through e-mail attachments, peer-to-peer file transfer networks, social networks, and websites that ask the visitor to download a special program to experience their contents. They also propagate through USB RAM sticks. In Afghanistan, the Taliban left infected USB RAM sticks in public places. When U.S. forces found these drives and inserted them into their USB ports, they spread the infection throughout their networks.

Worms Worms are similar to viruses. However, instead of attaching themselves to other programs, **worms** are full programs. Normally, they propagate between computers with the same mechanisms that viruses use.

However, some worms are able to jump directly between computers without human intervention. This means that no social engineering is necessary. However, **directly propagating worms** have a major limitation. They must be written to exploit a particular vulnerability on the receiving host, and that host must have the vulnerability. Humans are often gullible, but propagation through social engineering takes time. Worms that propagate directly can do so in milliseconds, and each worm that succeeds will launch itself against many more victims. In 2003, the Slammer worm infected 90% of all vulnerable computers attached to the Internet within 10 minutes.

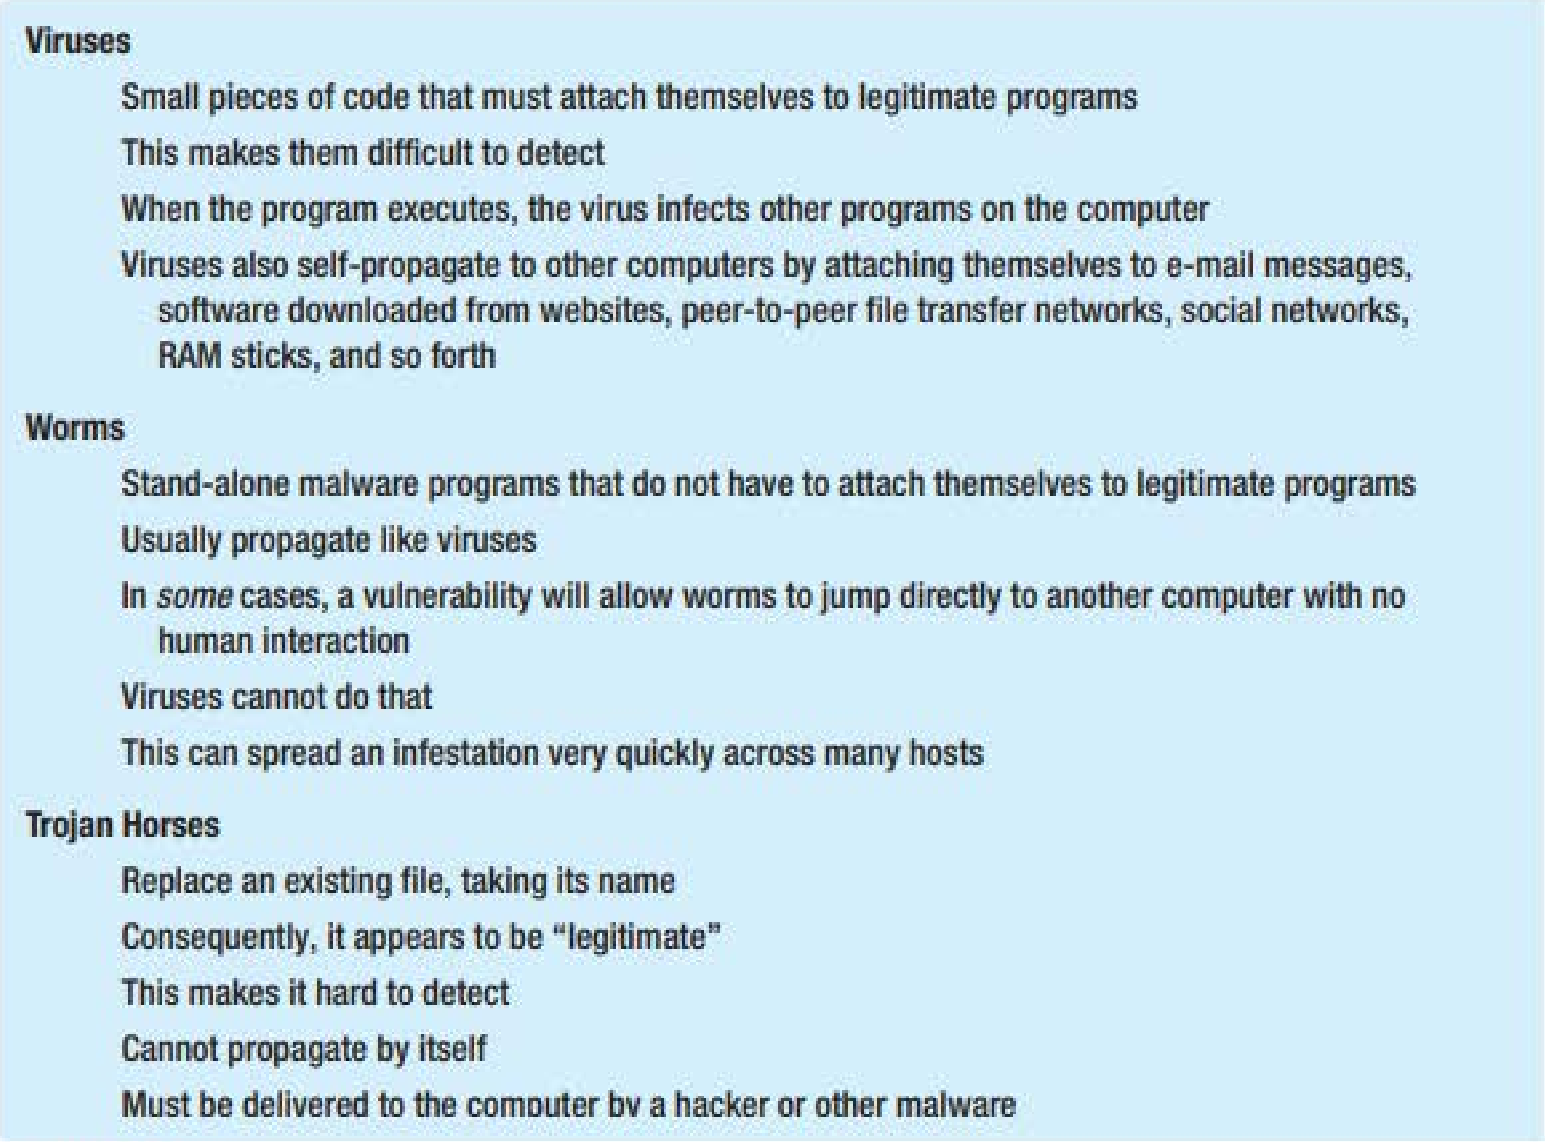

Viruses

Small pieces of code that must attach themselves to legitimate programs

This makes them difficult to detect

When the program executes, the virus infects other programs on the computer

Viruses also self-propagate to other computers by attaching themselves to e-mail messages, software downloaded from websites, peer-to-peer file transfer networks, social networks, RAM sticks, and so forth

Worms

Stand-alone malware programs that do not have to attach themselves to legitimate programs

Usually propagate like viruses

In *some* cases, a vulnerability will allow worms to jump directly to another computer with no human interaction

Viruses cannot do that

This can spread an infestation very quickly across many hosts

Trojan Horses

Replace an existing file, taking its name

Consequently, it appears to be "legitimate"

This makes it hard to detect

Cannot propagate by itself

Must be delivered to the comouter bv a hacker or other malware

FIGURE 4-3 Common Types of Malware

Trojan Horses In *The Iliad*, the Trojan horse was supposed to be a gift offering. It was really a trap. The Greeks left it at the gate and let the Trojans bring it inside. In malware, a **Trojan horse** is similar.

- First, it disguises itself as a legitimate file. This makes it difficult to detect.
- Second, in contrast to viruses, worms, and mobile code, a Trojan horse cannot propagate to another computer on its own initiative. It must be placed there by another piece of malware, by a human hacker, or by a user downloading the program voluntarily.

A Trojan horse cannot spread from one computer to another by itself.

Test Your Understanding

7. a) How do viruses and worms differ? b) How do viruses and worms propagate using social engineering? c) Do all worms spread by direct propagation? d) Why is direct propagation especially dangerous? e) What are Trojan horses? f) How do Trojan horses propagate to computers?

Payloads

In war, when a bomber aircraft reaches its target, it releases its payload of bombs. Similarly, after they spread, viruses, worms, and other types of malware may execute pieces of code called **payloads**. Malicious payloads can do extensive damage. Figure 4-4 gives some example of this.

Erasing or Encrypting Your Hard Drive Most people do not back up their files regularly or effectively. Some malware maliciously erases a hard drive, creating a devastating loss of critical data. More recently, ransomware has encrypted everything on a hard drive and has then told the user to pay a ransom to get the data unencrypted. This ransom typically must be paid in Bitcoins to a particular server. Typically, the thieves do provide the decryption key, but this is not always the case. In the massive WanaCry ransomware attack that took place in May 2017, the thieves had a poor payment system and generally did not decrypt their victims' files. This was extremely damaging because this massive attack encrypted the data on about 300,000 computers around the world.[24] Ironically, it is thought that the thieves reaped less than $100,000 in ransom during the attack. Most ransom attacks are smaller but bring in more money.

Turn Your Computer into a Spam or Pornography Server Nobody likes getting spam, which is unsolicited commercial e-mail, often of a fraudulent nature. Where does it come from? Actually, it may be coming from your own computer.

After Propagation, Malware May Execute Payloads
- Code that does damage

Malicious Payloads Intend to Do Damage
- Can erase your hard drive
- Ransomware encrypts your files, forcing you to pay ransom to be able to read them
- Can make your computer into a spam source or pornography distribution site
- Spyware can steal information from your computer and send it to attackers
  - Keystroke loggers capture what you type
  - Data miners search your storage for Social Security numbers, bank account numbers, etc.
- Credit card number theft
  - Steal credit card numbers, make unauthorized purchases
  - Credit card companies will reimburse, but the process can be painful
- Identity theft
  - Steal enough information to impersonate the victim in large financial transactions
  - No reimbursement for stolen funds
  - Repairing credit can be difficult

FIGURE 4-4 Payloads

[24] Dustin Volz, "Cyber Attack Eases, Hacking Group Threatens to Sell Code," *Reuters*, May 17, 2017. http://www.reuters.com/article/uscyberattackidUSKCN18B0AC.

Spammers often install spam-generating software on compromised computers. (Why pay for their own computers to send spam?) More seriously, some attackers will turn a compromised computer into a pornography server, even a child pornography server. This will, of course, litter the computer with pornography.

Spyware One concern on the list is **spyware**, which can steal information from your computer and send it to attackers. **Keystroke logger** spyware captures what you type and analyzes it for login credentials and other things you type. It then sends these keystrokes back to the spymaster. At a more sophisticated level, **data miners** actively search your storage for Social Security numbers, bank account numbers, and other sensitive information. Data miners can extract a great deal of sensitive data in a very short period of time.

Credit Card Number Theft Two other payloads are very common. One is malware to do **credit card number theft**. The thief can use this information to make unauthorized purchases. Credit card firms will refund money spent on purchases by the thief, but getting this refund can be a painful process.

Identity Theft In some cases, thieves collect enough data about a victim (name, address, Social Security number, driver's license number, date of birth, etc.) to impersonate the victim in complex financial transactions. This impersonation is called **identity theft**. Thieves commit identity theft in order to purchase expensive goods, take out major loans using the victim's assets as collateral, obtain prescription drugs, get a job, enter the country illegally, and do many other things. Identity theft is more damaging than credit card theft because it can involve large monetary losses that are not reimbursed by anyone. In addition, correcting the victim's credit rating can take months. Some victims have even been arrested for crimes committed by the identity thief.

Test Your Understanding

8. a) What are payloads? b) What is ransomware? c) What is spyware? d) What is the difference between the two types of spyware mentioned in the text? e) Distinguish between credit card number theft and identity theft. g) Which is more harmful to the victim? Why?

Human Break-Ins (Hacking)

A virus or worm typically has a single attack method. If that method fails, the attack fails. However, human adversaries can attack a company with a variety of different approaches until one succeeds. This flexibility makes human break-ins much more likely to succeed than malware break-ins.

What Is Hacking? Breaking into a computer is called hacking. Legally, **hacking** is defined as *intentionally* using a computer resource *without authorization* or *in excess of authorization*. The key issue is authorization.[25] If you see a password written on a note

[25] Note that the unauthorized access must be intentional. Proving intentionality is almost always necessary in criminal prosecution, and hacking is no exception. However, damage does not have to be intentional for a break-in to be hacking.

Humans Can Use Many Attack Methods

- This makes them more dangerous than malware, which usually has only one or two attack methods

Hacking

- Intentionally using a computer resource
- without authorization or
- in excess of authorization

If an Action Fits the Definition, It Is Hacking

- For example, if you find username and password on a piece of paper negligently left around, you are still not authorized to use the account, so to use it would be hacking

Irrelevant Considerations

- Not well-protected: does not excuse hacking
- Just testing the resource's security: does not excuse hacking

Penalties Depend on the Amount of Damage Done

- Easy to do damage accidentally

FIGURE 4-5 Human Break-Ins (Hacking) (Study Figure)

attached to a computer screen, this does not mean that you have authorization to use it. Also, note that it is hacking even if a person has legitimate access to an account but uses the account for *unauthorized* purposes.

Hacking is intentionally using a computer resource without authorization or in excess of authorization.

All hacking is illegal. Penalties differ by the type of asset that is hacked and by the amount of damage done, but it is very easy to do enough harm accidentally to merit a jail term, and "intentionally" only applies to intending to use the asset, not intending to do damage.

Test Your Understanding

9. a) What is the definition of hacking? b) If you see a username and password on a sticky note on a monitor, is it hacking if you use this information to log in? Explain in terms of the definition. (Answer: No, you did not receive authorization to use it.) c) You discover that you can get into other e-mail accounts after you have logged in under your account. You spend just a few minutes looking at another user's mail. Is that hacking? Explain in terms of the definition. d) If you click on a link expecting to go to a legitimate website but are directed to a website that contains information you are not authorized to see, is that hacking? Explain in terms of the definition.

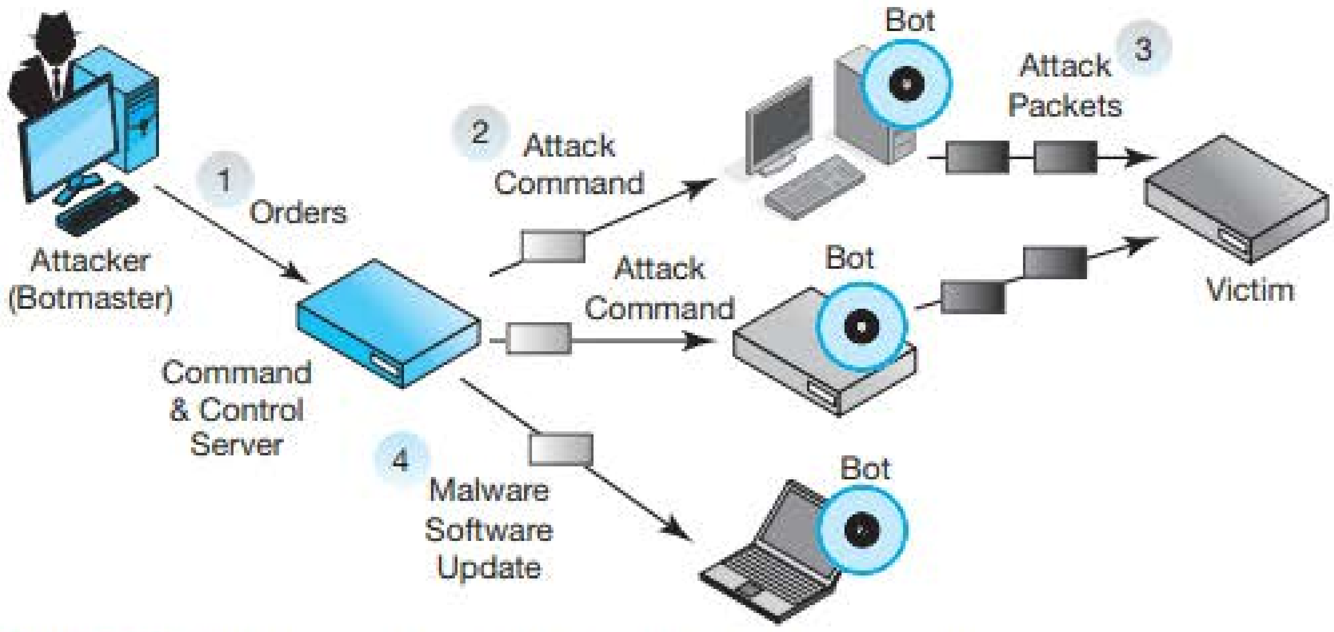

FIGURE 4-6 Distributed Denial of Service (DDoS) Attack Using Bots

Denial-of-Service (DoS) Attacks

The goal of **denial-of-service (DoS) attacks** is to make a computer or an entire network unavailable to its legitimate users. We saw the massive DDoS attack against KrebsOnSecurity.com in Chapter 1. Let's look at these attacks in a bit more detail. As Figure 4-6 shows, most DOS attacks involve flooding the victim computer with attack packets. The victim computer becomes so busy processing this flood of attack packets that it cannot process legitimate packets. The overloaded host may even fail. Or transmission lines may be so clogged with distributed denial-of-service (DDoS) traffic that the host may remain active but be unreachable.

The goal of denial-of-service (DoS) attacks is to make a computer or an entire network unavailable to its legitimate users.

The attacker begins by compromising computers and installing malware programs called **bots** on hundreds or thousands of PCs or servers. This collection of compromised computers is called a **botnet**, and the attacker is called the **botmaster**. When the user sends these bots an attack command, they all begin to flood the victim with packets.

Typically, the adversary does not communicate with bots directly. Rather, he or she sends orders to a **command and control server**, which then sends attack commands to the bots. In effect, the attacker is two levels removed from the attack, making the botmaster difficult to identify.

In many cases, the bot malware will not function properly when it is first fielded. With bots, however, the botmaster can send updates, as Figure 4-6 illustrates. Thus fixed, the bots can be effective in subsequent attacks.

More radically, the bot malware can be changed from one type to another. Many botnets are created initially to generate spam. As these attacks become less effective, the botmaster can turn the botnet into a DDoS machine and later something else.

Test Your Understanding

10. a) What is the purpose of a denial-of-service attack? b) Which programs directly attack the victim in a distributed denial-of-service attack? c) What is a collection of compromised computers called? d) What is the person who controls them called? e) To what computer does the attacker send messages directly? f) What are the implications of the fact that bots can be updated?

Advanced Persistent Threats (APTs)

In the past, criminal attacks were brief and limited—the electronic equivalent of smash-and-grab thefts in jewelry stores. Increasingly, however, we are experiencing **advanced persistent threats (APTs)** in which the adversaries have multiple objectives that they continue to pursue for days, weeks, months, or even years. These are true nightmares for corporations.

The adversary must first break into the firm. In a large majority of cases, he or she does this through an extremely well-crafted spear phishing attack that gives the attacker access to critical authentication credentials. (This was probably the case in the Target breach case at the beginning of this chapter.) "The adversary uses the initial foothold to explore and break into other parts of the firm's IT infrastructure. The attacker may also install Trojan horses and other exploitation programs. In each of these steps, the attacker uses advanced penetration and exploitation methods. This is the origin of "advanced" in the name.

APTs are expensive to mount. Consequently, they were only done by national governments in the past. However, if there are good prospects for a large theft, criminal groups now may launch them. This was the case with the Target attack.

Test Your Understanding

11. a) Explain "advanced" in the term *advanced persistent threat*. b) Explain "persistent" in the context of APTs. c) How do adversaries often enter the system and then expand to other parts of it? d) Who mounts APTs today?

Prolonged Attack

- Days, weeks, months, sometimes years
- Initial foothold, then move to other systems
- Plenty of time to learn systems and do damage

Advanced

- Uses attack techniques well beyond typical hacks
- Although often begins with a relatively simple spear phishing attack

Difficult and Expensive

- Only worth it for major objectives
- Once done only by nation-states, now done by criminal hackers (e.g., Target)

FIGURE 4-7 Advanced Persistent Threats (Study Figure)

TYPES OF ATTACKERS

The threat environment consists of types of attacks and types of attackers. As Figure 4-8 shows, there are several different types of attackers facing organizations today.

Cybercriminals

When most people think of attackers, they normally have two pictures in their mind. The first is the old-school hacker driven by curiosity, the thrill of the break-in, and the desire to increase one's reputation among other old-school hackers. They were seen as annoying but not too damaging.

This view is completely out of date. Hackers today are overwhelmingly **cybercriminals** who attack to make money. This has been true since the beginning of this century. Cybercriminals often work in loosely structured gangs. Funded by their

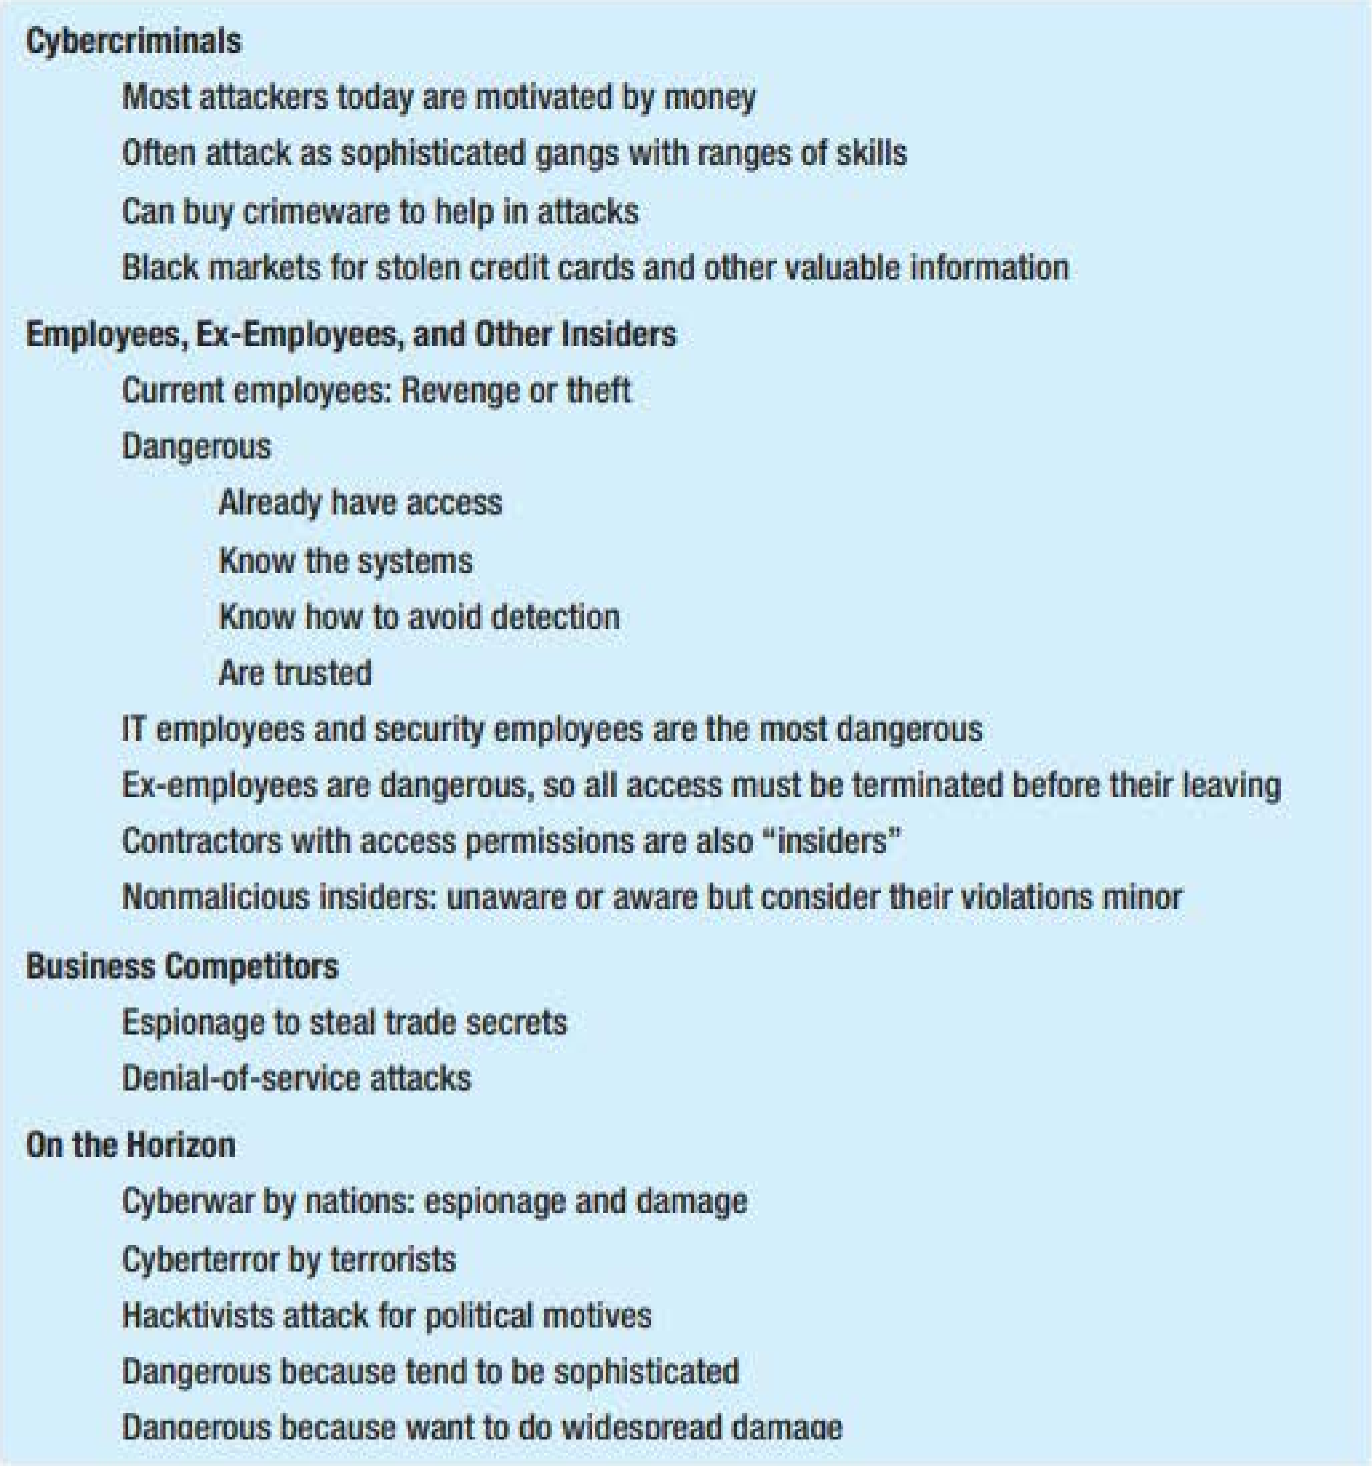

Cybercriminals

- Most attackers today are motivated by money
- Often attack as sophisticated gangs with ranges of skills
- Can buy crimeware to help in attacks
- Black markets for stolen credit cards and other valuable information

Employees, Ex-Employees, and Other Insiders

- Current employees: Revenge or theft
- Dangerous
  - Already have access
  - Know the systems
  - Know how to avoid detection
  - Are trusted
- IT employees and security employees are the most dangerous
- Ex-employees are dangerous, so all access must be terminated before their leaving
- Contractors with access permissions are also "insiders"
- Nonmalicious insiders: unaware or aware but consider their violations minor

Business Competitors

- Espionage to steal trade secrets
- Denial-of-service attacks

On the Horizon

- Cyberwar by nations: espionage and damage
- Cyberterror by terrorists
- Hacktivists attack for political motives
- Dangerous because tend to be sophisticated
- Dangerous because want to do widespread damage

FIGURE 4-8 Types of Attackers

crimes, many criminals can afford to hire the best hackers and to enhance their own skills. Consequently, criminal attacks are not just growing in numbers; they also are growing very rapidly in technical sophistication.

Today, most hackers are cybercriminals.

Criminal attackers have access to a vast online cybercrime community that gives them access to **crimeware** programs with slick user interfaces and prepaid annual updates. There are e-commerce black markets for them to buy and sell credit card numbers and identity information. Many elements of this black market are in countries where law enforcement is minimal at best.

Test Your Understanding

12. a) What type of adversary are most hackers today? b) Why is this type of attacker extremely dangerous? c) What resources can they purchase and sell over the Internet?

Employees, Ex-Employees, and Other Insiders

A large number of attacks are undertaken not by outsiders but by employees. Often, they are disgruntled employees who attack for revenge. However, they also may be employees who simply want to steal. Employees are especially dangerous for four reasons:

- They are knowledgeable about corporate systems,
- They typically have access to key systems,
- They have knowledge about how to avoid detection, and
- They tend to be trusted.

The most dangerous employees are IT staff members and especially IT security staff members. An ancient Roman question, "Quis custodiet ipsos custodes?" means "Who guards the guardians?" It is a serious question in security.

For ex-employees, revenge is a common motive. Another is stealing trade secrets that the employee worked on and believes are "his" or "hers." It is important to terminate all ex-employees' access to internal resources after they leave. In fact, even *before* employees leave, it is important to monitor their access for signs that they are infiltrating company intellectual property.

Often, contractors and service providers are given access credentials. This makes them **insiders**, and they must be considered dangerous as a consequence. When Edward Snowden stole files from the National Security Agency in early 2013, he was an employee of contractor Booz Allen Hamilton in Hawaii. In the Target breach, the account that thieves used to break into Target's computers was that of an air conditioning service company performing services for Target.

Malicious insiders have garnered the most attention in the past. However, companies also need to be concerned with **nonmalicious insiders** who commit security violations through ignorance or because they consider that the violation will do little or no damage. The Target breach and many other breaches begin with unsafe acts by nonmalicious insiders.

Test Your Understanding

13. a) Why may employees attack? b) For what four reasons are employees especially dangerous? c) Who are the most dangerous employees? d) Why may ex-employees attack? e) What should be done before an employee leaves the firm? f) Why are contractor firms more dangerous than other outside firms?

Business Competitors

Your business competitors may also attack you. All businesses have trade secrets such as customer lists, production schedules, product formulations, lists of projects, and lists of employees. Especially in some industries, it is common for business competitors to attempt to find these trade secrets. This may range from legal activities such as looking at your public website to hacking attacks to find well-protected information. In some cases, competitors will actually attack you using denial-of-service attacks and other disruptive assaults. They also may attack your firm's reputation via social media.

Test Your Understanding

14. What three types of attacks may come from your firm's business competitors?

Cyberterrorists and National Governments

On the horizon is the danger of far more massive **cyberterror** attacks by terrorists and even worse **cyberwar** attacks by national governments. These could produce unprecedented damages in the hundreds of billions of dollars.

The United States has acknowledged that it has long had cyberwar capabilities, and it established a consolidated Cyberwar Command in 2009. It is clear that several other countries have these capabilities as well (especially China). Countries could use IT to do espionage to gather intelligence, conduct attacks on opponents' financial and power infrastructures, or destroy enemy command and control facilities during physical attacks.[26]

Russia is especially focused on disrupting political processes in other countries. Russia's hacking and selective release of information during the United States presidential election in 2016 has received the most attention, but they have been active during the elections of several other countries as well.

Cyberterror attacks by terrorists are also likely. During physical attacks, terrorists might disable communication systems to thwart first responders and to spread confusion and terror among the population. Cyberterrorists could also conduct purely

[26] A 2009 article in the *New York Times* reported that before the 2003 invasion of Iraq, the United States considered an attack that would shut down Iraq's entire financial infrastructure (John Markoff and Thom Shanker, "'03 Plan Displays Cyberwar Risk," *New York Times*, August 1, 2009. www.msnbc.msn.com/id/3032619/%2328368424). This attack was not approved, but not because it was infeasible. It was held back because its impact might have spread beyond Iraq and might even have damaged the U.S. financial system. More recently, attacks by the United States and Israel used the Stuxnet worm to damage a specific group of nuclear centrifuges in a specific factory in Iran. The researchers who discovered Stuxnet were amazed by its complexity and by the scope of the operation that produced and tested it. It even involved forged digital certificates for important firms.

IT-based attacks. Nation-states are concerned about the side effects of cyberwar attacks, but terrorists have no such qualms.

Cyberwar and cyberterror are particularly dangerous for three reasons. First, funding allows them to be extremely sophisticated. Second, they focus on doing damage instead of committing thefts. Third, they are dangerous because they are likely to be directed against many targets simultaneously for massive damage.

Espionage has more limited objectives than destructive attacks. In spying, the goal is to learn an enemy's secrets. Several countries are doing this on a massive scale. In many cases, they also are targeting commercial enterprises to steal trade secrets useable by firms in their countries. The Chinese have been very effective in penetrating classified U.S. defense resources in recent years.

Hacktivists attack for political motives. They do so to embarrass corporations or governments. Edward Snowden's publication of secret programs in the U.S. National Security Agency (NSA) was an example of hacktivism. Although hacktivists are often viewed favorably, their release of information can cause considerable damage. Wikileaks has been the most active hacktivist group. Its website, wikileaks.org, has a long list of files that still can be downloaded.

Test Your Understanding

15. a) What are cyberterror and cyberwar attacks? b) Why are cyberwar attacks especially dangerous?

PROTECTING DIALOGUES CRYPTOGRAPHICALLY

Having looked at the threat environment, we now begin to look at the tools that companies use to attempt to thwart attackers. One of these is cryptography. Formally, **cryptography** is the use of mathematics to protect information.

Cryptography is the use of mathematics to protect information.

Cryptography is important in and of itself. We begin with "crypto," however, because it is part of many other security protections. A knowledge of cryptography is necessary to understand how they work.

Encryption for Confidentiality

Encryption for Confidentiality When most people think of cryptography, they think of **encryption for confidentiality**, which Figure 4-9 illustrates. **Confidentiality** means that even if an eavesdropper intercepts a message, he or she will not be able to read it. The sender uses an encryption method, called a **cipher**, to create a message that an eavesdropper cannot read. However, the receiver can **decrypt** the message in order to read it.

Confidentiality means that even if an eavesdropper intercepts a message, he or she will not be able to read it.

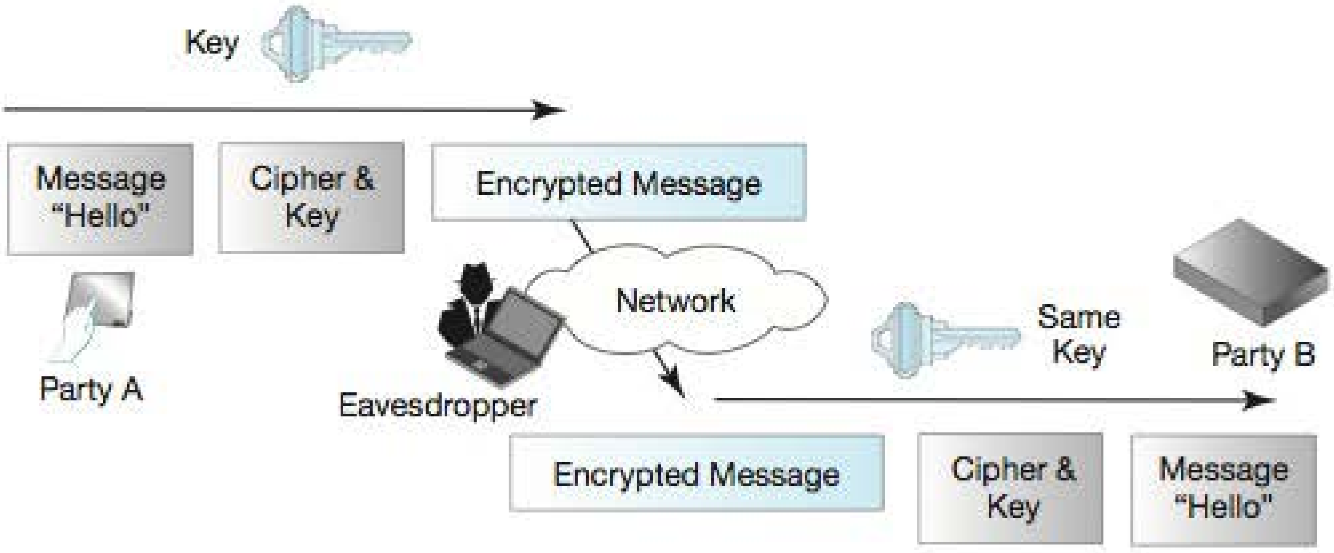

FIGURE 4-9 Encryption for Confidentiality

Keys It is impossible to keep ciphers secret. However, the message is encrypted with both the cipher and a key. Different keys produce different encrypted messages for the same cipher. Consequently, these keys must be kept secret. There is nothing mysterious about keys. They are simply strings of bits of a certain length. They should be selected randomly.

Key Length **Cryptanalysts** study encrypted messages in order to learn encryption keys. The normal way to do this is to try all possible keys to see which one produces an intelligible message. The way to defeat exhaustive key searches is to use long keys. Every bit that is added to the key doubles the number of keys that must be tried by cryptanalysts. For most encryption ciphers, key lengths must be 128 bits or greater to be considered strong. For some, however, strong keys must be 2,000 bits or longer.

Keys are long strings of bits.

Test Your Understanding

16. a) What protection does confidentiality provide? b) What is a cipher? c) In encryption for confidentiality, what must be kept secret? d) What is the minimum size for encryption keys to be considered strong in most encryption ciphers?

Electronic Signatures: Message Authentication and Integrity

Confidentiality is not the only goal of cryptology. In addition to encrypting each packet for confidentiality, cryptographic systems normally add **electronic signatures** to each packet. This is illustrated in Figure 4-10. Electronic signatures are small bit strings that provide message-by-message **authentication**, which ensures that the person or program you are communicating with is not an impostor. An electronic signature allows the receiver to detect a message added to the dialogue by an impostor.

Authentication ensures that the person or program you are communicating with is not an impostor.

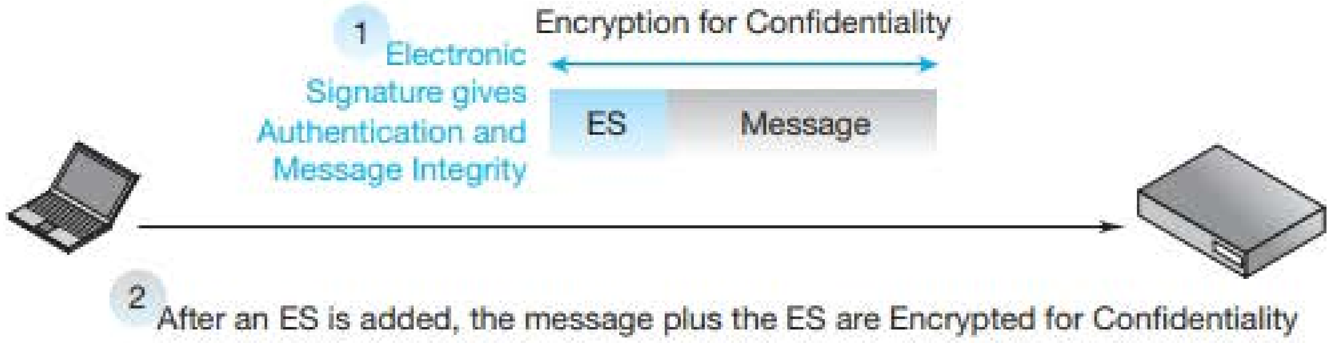
1 Electronic Signature gives Authentication and Message Integrity
Encryption for Confidentiality
ES
Message
2 After an ES is added, the message plus the ES are Encrypted for Confidentiality

FIGURE 4-10 Electronic Signature for Authentication (and Message Integrity)

Electronic signatures also provide a second cryptographic benefit, message integrity. **Message integrity** means that the receiver will be able to detect whether the packet is altered by an attacker while the packet is in transit. If the message integrity test fails, the receiver discards the message.

Overall, cryptographic systems provide three protections to every packet. Encryption for confidentiality provides message-by-message confidentiality, while electronic signatures provide message-by-message authentication and message integrity.[27]

Overall, cryptographic systems provide three protections to every packet. Encryption for confidentiality provides message-by-message confidentiality, while electronic signatures provide message-by-message authentication and message integrity.

Test Your Understanding

17. a) What two protections do electronic signatures provide? b) What three protections are typically given to each packet?

Host-to-Host Virtual Private Networks (VPNs)

Sometimes, transmission through untrusted networks is necessary. One of these untrusted networks is the Internet, which has no built-in security and is full of attackers. Other untrusted networks are wireless networks; in these networks, anyone can intercept your transmissions. The way to address a lack of security is to create a host-to-host **virtual private network (VPN)**. Figure 4-11 illustrates this concept. Of course, transmissions actually pass through a real network. In terms of security, however, the hosts are effectively communicating via a private network that connected just them. A VPN makes it appear that the two hosts are communicating via a private secure network.

A VPN makes it appear that the two hosts are communicating via a private secure network.

[27] Another common protection is anti-replay. In some cases, an attacker may be able to do damage by capturing an encrypted message. Although the attacker cannot read the encrypted message, he or she may be able to accomplish objectives by simply retransmitting the message later. Anti-replay protections prevent this.

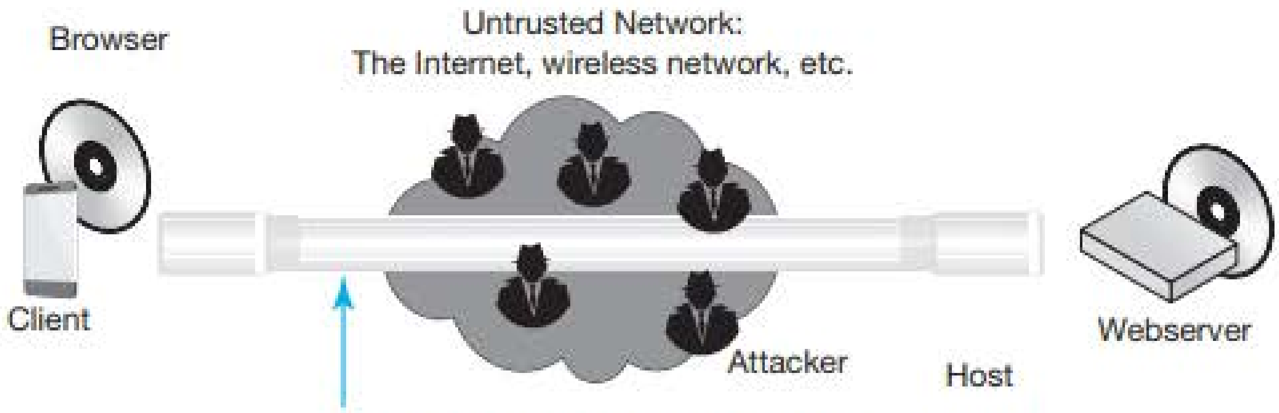
Browser
Untrusted Network:
The Internet, wireless network, etc.
Client
Attacker
Host
Webserver

SSL/TLS Virtual Private Network (VPN)
Cryptographically protected transmission path
Like having your own private network to connect to the webserver

FIGURE 4-11 SSL/TLS Host-to-Host Virtual Private Network (VPN)

A common cryptographic system for VPNs is SSL/TLS. SSL/TLS was created as Secure Sockets Layer (SSL) by Netscape. The Internet Engineering Task Force then took over the standard, renaming it Transport Layer Security (TLS). It is called by both names today, so we call it **SSL/TLS.**[28]

SSL/TLS is an attractive cryptographic system for webservice because SSL/TLS is built into every webserver and browser today, so the cost of adding SSL/TLS protection is negligible. Given security threats on the Internet, SSL/TLS should be used whenever possible.

Test Your Understanding

18. a) Distinguish between private networks and virtual private networks. b) Why is SSL/TLS attractive for VPNs to connect browsers to webservers?

AUTHENTICATION

Electronic signatures provide message-by-message authentication. However, there are many types of authentication in use today, each with strengths and weaknesses. Authentication is crucial to controlling access to resources so that adversaries can be prevented from reaching them.

Authentication Terminology and Concepts

Figure 4-12 illustrates the main terminology and concepts in authentication. The user trying to prove his or her identity is the **supplicant**. The party requiring the supplicant to prove his or her identity is the **verifier**. The supplicant claims to be a particular user,

[28] When you use SSL/TLS, the URL begins with https://. Although you will not notice it, the port number in TCP changes from 80 to 443, which indicates HTTP over SSL/TLS.

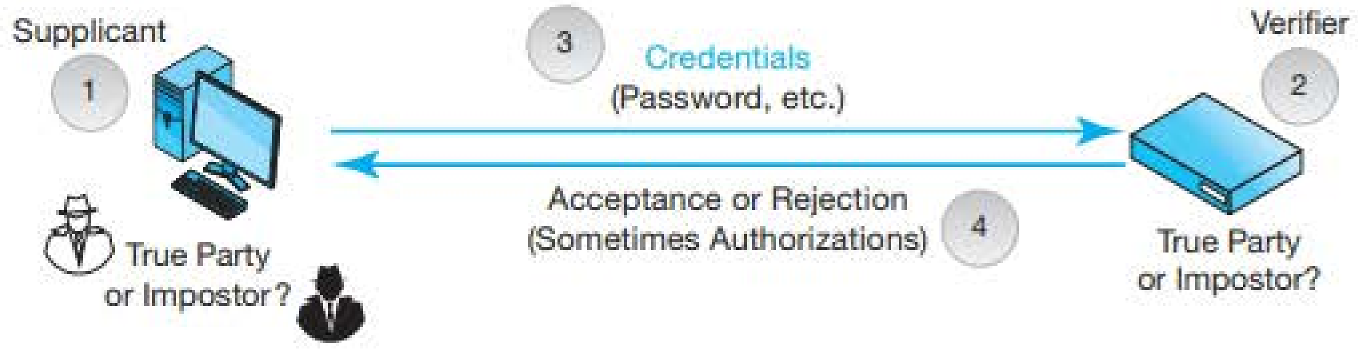

FIGURE 4-12 General Authentication Concepts

the **true party**. The supplicant tries to prove that he or she is the true party by providing **credentials** (proofs of identity) to the verifier.

The specific question that all authentication methods ask is whether the supplicant's credentials prove that he or she is the true party.

The type of authentication tool that is used with each resource must be *appropriate for the risks to that particular resource*. Sensitive information should be protected by very strong authentication methods. However, strong authentication is expensive and often inconvenient. For relatively nonsensitive data, weaker but less expensive authentication methods may be sufficient.

The type of authentication tool that is used with each resource must be appropriate for the risks to that particular resource.

Test Your Understanding

19. a) What is authentication? b) Distinguish between the supplicant and the verifier. c) What are credentials? d) Who is the true party? e) What is the specific goal of authentication? f) Is the supplicant the true party or is the supplicant an impostor? g) Why must authentication be appropriate for risks to an asset?

Reusable Passwords

The most common authentication credential is the **reusable password**, which is a string of characters that a user types to gain access to the resources associated with a certain **username** (account) on a computer. These are called reusable passwords because the user types the same password each time he or she needs access to the resource. Unfortunately, the reusable password is the weakest form of authentication, and it is appropriate only for the least sensitive assets.

The reusable password is the weakest form of authentication, and it is appropriate only for the least sensitive assets.

Reusable Passwords

- Passwords are strings of keyboard characters
- They are typed to authenticate the use of a username (account) on a computer
- They are used repeatedly and so are called reusable passwords

Benefits

- Ease of use for users (familiar)
- Inexpensive because they are built into operating systems

Often Weak (Easy to Crack)

- Common words and simple variations

Traditional Advice for Password Security

- Passwords should be long and complex
- Should be at least 8 to 12 characters long
- Should mix case, digits, and other keyboard characters ($, #, etc.)
- Such passwords are very strong but difficult to remember
- Such passwords are often written down or stored online or worked around
- In the end, long and complex passwords do not live up to their promise

2017 National Institute of Standards and Technology Guidance

- Use long phrases that are not easy to guess
- Using all lower case is fine if the phrase is long
- Still use a different password for each site, and not a simple variation of the phrase
- This approach is both more secure and easier on employees than traditional approaches to strong passwords

Perspective

- Even with improvements, reusable passwords are only strong enough for very unimportant assets
- The goal of new types of authentication is to allow firms to get rid of reusable passwords entirely

FIGURE 4-13 Reusable Password Authentication (Study Figure)

Ease of Use and Low Cost The popularity of password authentication is hardly surprising. For users, passwords are familiar and relatively easy to use. For corporate IT departments, passwords add no cost because operating systems and many applications have built-in password authentication.

Picking Poor Passwords Unfortunately, users tend to pick very poor passwords. The most common password is 123456, and many others are easily guessable. These weak passwords are common words, names of sports teams, and common variations (such as replacing an "s" by a dollar sign and the letter "l" by a 1). If any sizeable fraction of a company's employees uses weak passwords, the company will be at considerable risk, because an attacker can often jump easily from an employee computer (or mobile phone) to a more sensitive device such as a server in accounting.

Traditional Advice on Reusable Passwords Traditional advice has been to force employees to pick strong passwords by insisting that their passwords be long, have at least a single change of case (not in the first character), a digit, and a non-letter, non-number character such as a space or @ sign.

This enforces strong passwords, but people have a hard time remembering them. They write them down, store them in a file on the computer, or keep them in some other poorly protected form. They use the same password for many different hosts. They forget their passwords and get a password reset that usually involves them answering challenge questions that are easily guessed or allows an attacker taking over their e-mail account to do the password resets.

National Institute of Standards and Technology In the United States, the National Institute of Standards and Technology creates recommended security practices for the U.S. government. These recommended practices are implemented by many commercial and nonprofit firms as well because of the Institute's reputation for security excellence. In 2017, the Institute release a revolutionary set of recommendations for reusable passwords. They said that traditional approaches to enforcing strong passwords had backfired and led to bad security practices.

In its new recommendation, the Institute recommended a radical change. Forget case changes, digits, and other keys. Just create long phrases that cannot be easily guessed. This will give the same or better computational strength and yet will be easy to remember. One should still use different passwords at other sites, but the new National Institute of Standards and Technologies new recommendation for long but memorable passphrases should go a long way to change the way corporations use reusable passwords.

Ever-Smaller Scope of Usefulness Even with the new password guidelines from the National Institute of Standards and Technology, password will remain a very weak form of authentication suitable only for the least risky assets. Other forms of authentication are being created specifically to allow firms to completely eliminate reusable passwords.

Passwords are only useful for nonsensitive assets.

Test Your Understanding

20. a) What was the traditional recommendation for passwords? b) What is the U.S. National Institute of Standards and Technology's new recommendation? c) What two benefits should this new recommendation bring? d) Is it still important not to use the same password at multiple sites? e) Why is it undesirable to use reusable passwords for anything but the least sensitive assets? f) Why are other forms of authentication being created?

Other Forms of Authentication

Companies are beginning to look for stronger types of authentication for most of their resources. This will allow them to replace most or all of their reusable password access systems. We have space to mention only the few types of authentication shown in Figure 4-14.

Access Cards To get into your hotel room, you may have to swipe an **access card** through a card reader. Many bus systems let riders purchase access cards to pay for their travel. Many companies use access cards for door access control. In addition, simple access card readers can be plugged into USB ports on computers for computer access. Of course, the loss of access cards is a fundamental problem.

Perspective

- Goal is to replace reusable passwords

Access Cards

- Permit door access
- Need to control distribution and disable lost or stolen access cards

Biometrics

- The use of biological measurements to authenticate you
- Vary in cost, precision, and susceptibility to deception
- Fingerprint scanning
    - Inexpensive but poor precision, deceivable
    - Sufficient for low-risk uses
    - For a notebook, may be better than requiring a reusable password
- Iris scanning
    - Based on patterns in the colored part of your eye
    - Expensive but precise and difficult to deceive
- Facial scanning
    - Based on facial features
    - Increasingly used in computers and cellular phones
    - Controversial because can be done surreptitiously—without the supplicant's knowledge
    - Varies widely in strength

Digital Certificate Authentication

- Extremely strong
- See Figure 4-15

Two-Factor Authentication

- Supplicant needs two forms of credentials
- Example: debit card and pin
- Strengthens authentication

FIGURE 4-14 Other Forms of Authentication

Losses must be reported immediately, and the card must be disabled remotely and for all access doors.

Biometrics In biometric authentication, access control is granted based on something you always have with you—your body. **Biometrics** is the use of bodily measurements to authenticate you. There are several types of biometrics. They differ in cost, precision, and susceptibility to deception by someone wishing to impersonate a legitimate user.

Biometrics is the use of body measurements to authenticate you.

- At the low end on price, precision, and the ability to reject deception is **fingerprint recognition**, which looks at the loops, whorls, and ridges in a finger. Although not a strong form of authentication, its price makes it acceptable for low-risk resources such as most tablets, and smartphones. For such devices, fingerprint recognition may be preferable given the tendency of people to pick poor passwords and to forget them.
- At the high end of the scale on price, precision, and the ability to reject deception is **iris recognition**,[29] which looks at the pattern in the colored part of your eye. Iris scanners are normally used for access to sensitive rooms.
- A controversial form of biometrics is **facial recognition**, in which an individual is identified by his or her facial features. Facial recognition can be done **surreptitiously**—without the knowledge of the person being scanned. This raises privacy issues. On the positive side, we are beginning to see facial recognition scanning on computers and smartphones.

Digital Certificate Authentication The strongest form of authentication is **digital certificate authentication**.[30] Figure 4-15 illustrates this form of authentication.

- In this form of authentication, each party has a secret **private key** that only he or she knows.
- Each party also has a **public key**, which anyone can know. It is not kept secret.
- The public key of a person is available from a **certificate authority** in a document called a **digital certificate**.[31] A digital certificate is cryptographically protected for message integrity, so that it cannot be changed without this change being obvious in a way that causes the verifier to reject it.

[29] In science fiction movies, eye scanners are depicted as shining light into the supplicant's eye. This does not happen. Iris scanners merely require the supplicant to look into a camera. In addition, science fiction movies use the term retinal scanning. The retina is the back part of the eye; it has distinctive vein patterns. Retinal scanning is not used frequently because the supplicant must press his or her face against the scanner.

[30] It also good for authenticating software processes, which have no heads or fingers and have a difficult time swiping access cards.

[31] The true party creates a public key/private key pair on his or her own computer. The TP then sends the public key to the certificate authority. The CA creates the digital certificate and sends it to anyone who wishes it. The true party never transmits his or her private key to anyone.

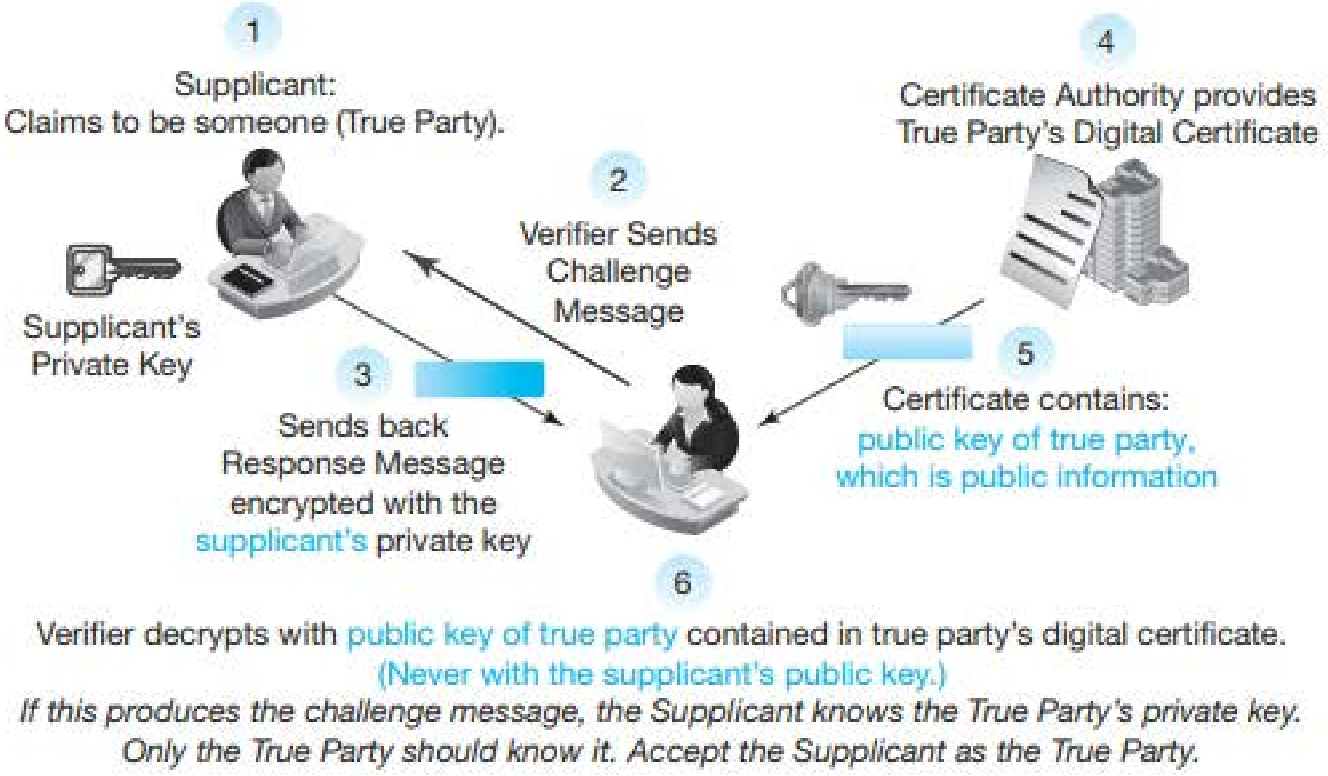
1
Supplicant:
Claims to be someone (True Party).
Supplicant's
Private Key
2
Verifier Sends
Challenge
Message
3
Sends back
Response Message
encrypted with the
supplicant's private key
4
Certificate Authority provides
True Party's Digital Certificate
5
Certificate contains:
public key of true party,
which is public information
6
Verifier decrypts with public key of true party contained in true party's digital certificate.
(Never with the supplicant's public key.)
If this produces the challenge message, the Supplicant knows the True Party's private key.
Only the True Party should know it. Accept the Supplicant as the True Party.

FIGURE 4-15 Digital Certificate Authentication

The supplicant claims to be someone, the true party. To test this claim, the verifier sends the subject a **challenge message**. This is just a random stream of bits. It is not even encrypted for confidentiality.

To prove its claim to being the true party, the supplicant encrypts the challenge message *with his or her private key* and sends this **response message** to the verifier. Again, there is no encryption for confidentiality.

The verifier gets the true party's digital certificate, which contains the true party's public key. The verifier tests the response message by decrypting it with the *public key of the true party*, which is contained in the digital certificate. If the decryption produces the original challenge message, then the supplicant has proven that he or she knows the private key of the true party. Only the true party should know this key. Therefore, it is reasonable to authenticate the supplicant as the true party. If the decrypted response message is not a match for the original challenge message, the supplicant is treated as an impostor.

Note that the verifier uses the *public key of the true party*—not the *supplicant's public key*. If the verifier used the supplicant's public key, the test would always succeed. The supplicant's public key would decrypt the message correctly. Impostors would *always* be authenticated.

Note that the verifier uses the public key of the true party—not the supplicant's public key.

There are three parties involved—the supplicant, the verifier, and the true party. Each has a public key and a private key. Therefore, you should never say *the* private key or *the* public key. Always say the supplicant's public or private key, the verifier's public or private key, or the true party's public or private key.

There are three parties involved—the supplicant, the verifier, and the true party. Each has a public key and a private key. Therefore, you should never say "the private key" or "the public key." Always say the supplicant's public or private key, the verifier's public or private key, or the true party's public or private key.

Two-Factor Authentication Debit cards are potentially dangerous because if someone finds a lost debit card, the finder might be able to use it to make purchases. Consequently, possession of the debit card is not enough to use it. To use a debit card, the user must type a **personal identification number (PIN)**, which usually is four or six digits long. Requiring two credentials for authentication is called **two-factor authentication**. Two-factor authentication is more difficult to defeat because the attacker must obtain both sets of credentials.[32]

Two-factor authentication requires two forms of authentication.

Test Your Understanding

21. a) How do you authenticate yourself with an access card? b) What is biometrics? c) Why may fingerprint recognition be acceptable for user authentication to a laptop? d) Why is iris recognition desirable? e) Why is face recognition controversial?

22. a) In digital certificate authentication, what does the supplicant do? b) What does the verifier do? c) Does the verifier decrypt with the true party's public key or the supplicant's public key? Why is this important? d) How does the verifier get the true party's public key?

23. a) What characteristic of the true party is used in access card authentication, iris authentication, and digital certificate authentication? b) Which form of authentication that we looked at depends on the supplicant proving that it knows something that only the true party should know? c) What if this information is learned by an attacker? d) Why is two-factor authentication desirable?

FIREWALLS AND INTRUSION DETECTION SYSTEMS

In hostile military environments, travelers must pass through checkpoints. At each checkpoint, guards will examine their credentials. If the guards find the credentials insufficient, the guard will stop the traveler from proceeding and note the violation in a checkpoint log.

[32] However, if a user's computer is compromised, the attacker typically controls both credentials, so two-factor authentication gives no security. Two-factor authentication also may fail if an eavesdropper can intercept authentication communication between the supplicant and the verifier. Two-factor authentication is desirable, but factors that limit its use must be understood.

Dropping and Logging Provable Attack Packets

Figure 4-16 shows that firewalls operate the same way. When a packet arrives, the **firewall** examines it. If the firewall identifies a packet as a **provable attack packet**, the firewall discards it. (Synonyms for provable are definite, certain, etc.) On the other hand, if the packet is not a provable attack packet, the firewall allows it to pass.

If a firewall identifies a packet as a provable attack packet, the firewall discards it.

If a packet is not a provable attack packet, the firewall passes it.

The firewall copies information about each discarded packet into a **firewall log file**. Firewall managers should read their firewall log files every day to understand the types of attacks coming against the resources that the firewall is protecting. This alerts the security staff to the kinds of attacks it is under at the time.

Note that firewalls pass *all* packets that are not provable (certain) attack packets, even if they are suspicious. By analogy, police cannot arrest someone unless they have probable cause, which is a reasonably high standard of proof. They cannot arrest someone for being suspicious.

Note that firewalls pass all packets that are not provable (certain) attack packets, even if they are suspicious. By analogy, police cannot arrest someone unless they have probable cause, which is a reasonably high standard of proof. They cannot arrest someone for acting suspiciously.

Consequently, firewalls never stop all attack packets. It is important to harden all internal hosts against attacks by adding firewalls, adding antivirus programs, installing all patches promptly, and taking other precautions. This chapter

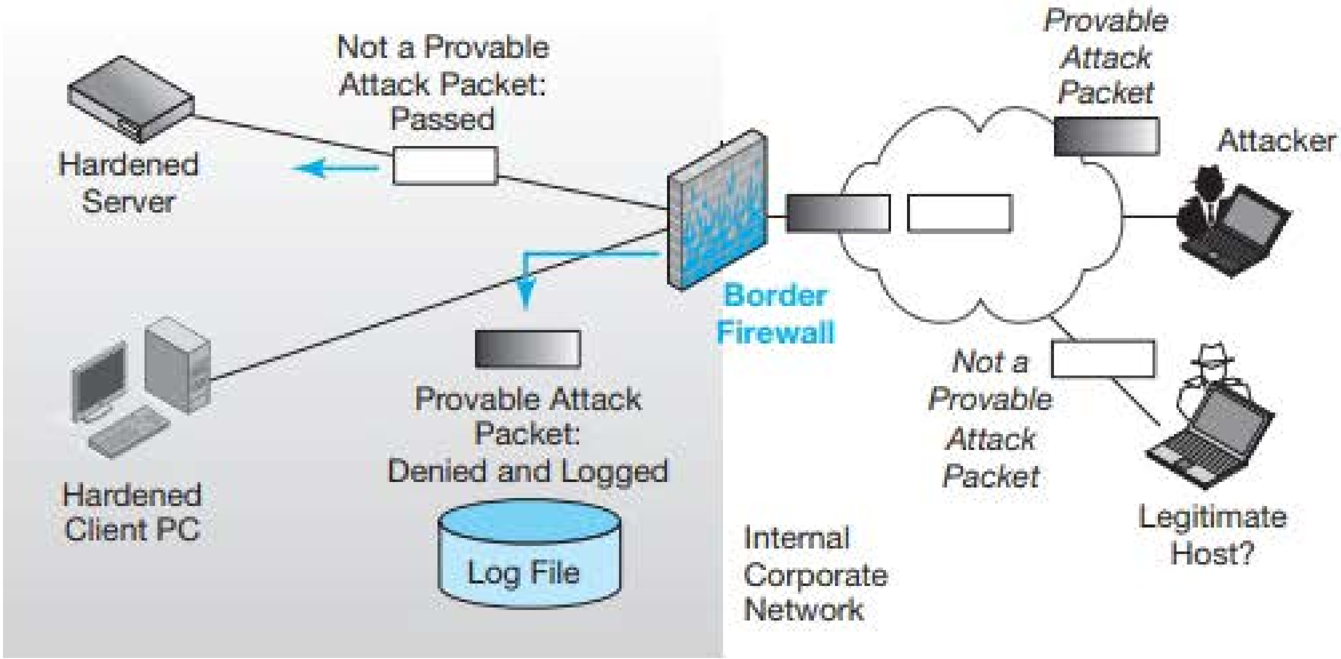

FIGURE 4-16 General Firewall Operation

focuses on network security, rather than host security, so we will not consider host hardening.

Because firewalls do not stop packets that are not provable attack packets, they never stop all attack packets.

Test Your Understanding

24. a) What does a firewall do when an arriving packet is definitely an attack packet? b) Does a firewall drop a packet if it probably is an attack packet? c) Why is it important to read firewall logs daily?

Stateful Packet Inspection (SPI) Firewalls

How do firewalls examine packets to see if they are attack packets? Actually, there are several **firewall filtering mechanisms**. We only look at two—stateful packet inspection (SPI) and next-generation firewalls.

Figure 4-17 shows that **stateful packet inspection (SPI)** firewalls recognize that there are two *states* (stages) in a dialogue between two parties. The first is the initial handshaking state, which is the initial interaction that takes place to authenticate the other party and other activities. This is crucial to security, so SPI firewalls spend a great deal of time on this and require high-quality authentication. If a connection between the two sides is not authenticated at this stage, the connection-opening attempt is terminated.

After this intense but brief handshaking state, everything else is ongoing communication. If a packet arrives that is part of an approved connection, it is given only cursory examination because authentication had been proven initially. It is passed through with little or no additional inspection. It is like an employee with an ID badge.

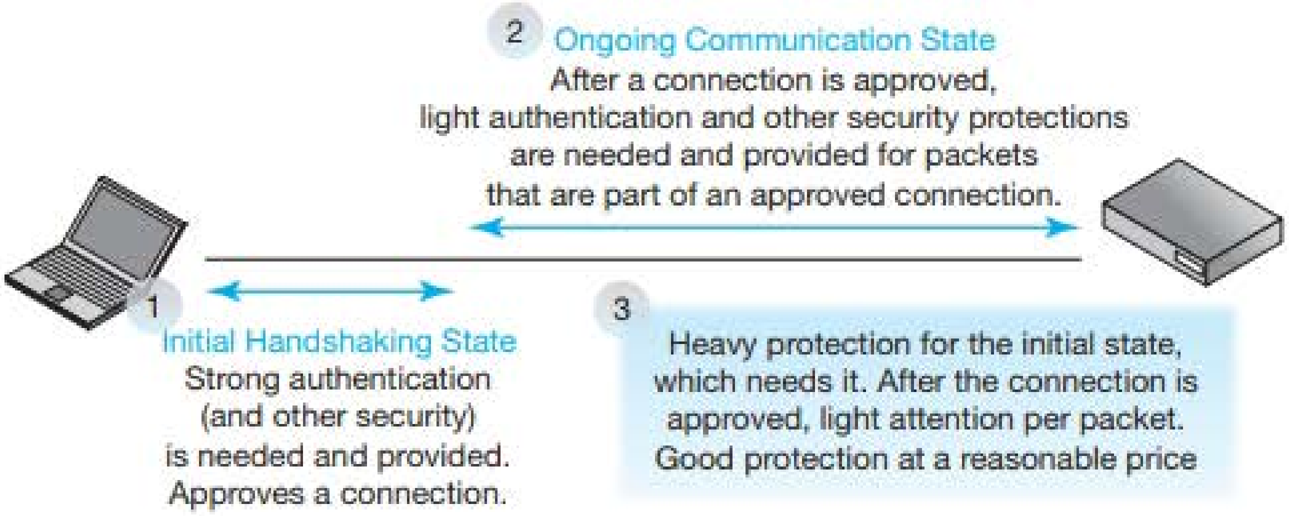
2 Ongoing Communication State
After a connection is approved,
light authentication and other security protections
are needed and provided for packets
that are part of an approved connection.
1 Initial Handshaking State
Strong authentication
(and other security)
is needed and provided.
Approves a connection.
3
Heavy protection for the initial state,
which needs it. After the connection is
approved, light attention per packet.
Good protection at a reasonable price

FIGURE 4-17 Stateful Packet Inspection (SPI) Firewall

To give an analogy, think of a telephone conversation. When someone calls you, you want to ensure that you know who you are talking to. Once you establish that, you ignore authentication for the remainder of the call. This is what SPI firewalls do with network connections.

If a packet attempts to open a connection, the SPI firewall compares it to the rules in its **access control list (ACL)**. ACL rules specify what to do with arriving packets. Figure 4-18 shows a simple SPI firewall ACL for connection-opening attempts.

An access control list (ACL) is a set of rules for determining what to do with arriving packets.

There are six columns. The fourth brings up something you saw in Chapter 2. This is the server port number. For webservers, the well-known port number is 80. For mail servers, it is 25. In this figure, the company that uses this ACL respects well-known port numbering.

Server port numbers usually specify the application involved in the connection.

- The first rule allows any device to open a connection to a particular webserver (IP address 60.44.2.17, server port number 80). This might be the company's public webserver.
- The second rule allows any device to open a connection to *any* webserver. This is a lazy rule that some firewall administrators use when they do not know what webservers they have but realize that blocking access to a legitimate webserver will cause problems.

Rule	Source IP Address	Destination IP Address	Server Port Number	Action on Connection	Remark
1	Any	60.3.47.138	80	Allow	Open access to this webserver.
2	Any	ANY	80	Allow	Open access to any webserver.
3	Any Internal	60.1.232.89	80	Authenti-cate, then allow	Open access for internal hosts to this webserver, following authentication.
4	Finance	Finance	Any	Authenti-cate, then allow	Any connection between Finance hosts with authentication.
5	Any Internal	60.44.2.17.	25	Allow	Open access for internal hosts to this mail server.
6	Any	Any	Any	Deny	Deny any connection not permitted by a previous rule.

FIGURE 4-18 Access Control List (ACL) for a Stateful Inspection Firewall to Apply to Packets Attempting to Create a Connection

- The third rule allows any *internal host* to connect to a particular webserver. This might be a webserver for human resources information that all internal employees should be able to reach. To add a level of protection, the supplicant must authenticate itself before the connection.
- The fourth rule permits connections between all finance department hosts. Again, authentication is required in this high-security environment.
- The fifth rule permits connections between any internal host and mail server 60.44.2.17 (server port number 25).
- The final rule categorically denies connections that have not been previously allowed by earlier rules. This is the "deny all" rule that typically ends this type of firewall. It enforces the policies that went into creating the ACL by disapproving any connection not envisioned by that policy. For example, if a client in marketing wishes to connect to a finance department server, this will be prohibited because it is not specifically approved.

Stateful packet inspection firewalls provide a good balance of security and economy. For the most sensitive part of a connection, SPI firewalls provide strong security. For ongoing communication, when less intense security is needed, SPI firewalls only do enough security for the situation. The latter saves money. Thanks to this balance of strong security when it is most needed and economical operation most of the time, most main border firewalls today are stateful packet inspection firewalls.

Test Your Understanding

25. a) Why are stateful packet inspection (SPI) firewalls attractive? b) What are the two states in connections for SPI firewalls? c) Which state needs the most security protection? Why? d) Why are SPI firewalls economical? e) What type of firewall do most corporations use for their main border firewalls?

26. a) In Figure 4-18, explain why Rule 1 brings more security than Rule 2. b) Explain why the last rule in an ACL should deny anything not previously approved by earlier rules. c) Why do you think authentication is sometimes required before accepting a connection? d) When a packet addressed to 60.1.232.89 arrives, what rule will the SPI firewall look at first? e) Why must Rule 2 come after Rule 1? f) Add a rule to permit access by hosts in accounting to server 60.3.4.67. Require authentication. What rule number would you give it?

Next-Generation (Application Aware) Firewalls (NGFWs)

The newest type of firewall is the **next-generation firewall (NGFW)**. The most important capability of NGFWs is that they are **application aware**. This means that, unlike stateful packet inspection firewalls, NGFWs can determine the specific application that is sending and receiving messages over a connection.

Stateful packet inspection firewalls might seem to be application aware, but they really are only aware of port numbers. This is problematic because an attacker can run an attack program over Port 80. Rule 2 in Figure 4-18 would enable all connections to this malware. This is called **port spoofing**. It is difficult to thwart with SPI firewalls.

Application-aware NGFWs prevent port spoofing by identifying any application that does not behave like a webserver. Further, NGFWs can usually identify the specific program (like Shazam does for music).

Identifying applications, while valuable, is expensive.

- First, the application-aware firewall must collect all traffic in a connection.
- If TCP is being used by the application, all of the packets delivering an application message must be brought together, their TCP segments extracted, and the application message reconstructed.
- This must be done for multiple application messages, sometimes many.
- Finally, the pattern across many application messages must be examined and compared to the fingerprints of various known applications. This may cause the NGFW to terminate the connection and report its findings to a security administrator.[33]

Application awareness allows firewall administrators to create rules for individual applications. Going back to Figure 4-18, an NGFW firewall would add a column for *Application*. For example, companies may not permit YouTube traffic or may limit it to avoid overloading the network. And, of course, if a malware program is identified, its traffic can be stopped automatically. By providing visibility into application traffic patterns, NGFWs also provide information for better network management.

Test Your Understanding

27. a) Why are SPI firewalls limited in their ability to detect attack packets? b) How do NGFWs address this problem? c) Think of at least two specific examples of how application information can be used to increase security. d) Why are NGFWs more expensive than SPI firewalls? (The answer is not in the text.)

Next-Generation Firewalls Are Application-Aware

Can base decisions on the actual application that is creating traffic

Advantages of Being Application Aware

Detects port spoofing (running a different application on a well-known port number).

Create accept/reject rules based on specific applications.

Recognize specific malware applications, providing laser-focused firewall decisions.

For network and security management, can see which applications are using which percentages of network traffic

Tradeoffs

Processing requirements make NGFWs more expensive than SPI firewalls per packet

Richer information requires greater management effort to create and implement policies and rules

FIGURE 4-19 Next-Generation (Application Aware) Firewalls (Study Figure)

[33] This complex process used to require a prohibitive amount of processing power. However, application-specific integrated circuits (ASICs) can now be built to handle application identification. ASICs put many calculations normally done in software into hardware instead. This is far faster than software processing, giving ASICs the power to handle their loads.

Intrusion Detection Systems (IDSs)

It would be naïve to expect firewalls to stop all attack packets. Most obviously, they do not drop suspicious packets—only definite attack packets. **Intrusion detection systems (IDSs)** were created to supplement firewalls by focusing specifically on identifying *suspicious* transmissions. When they find suspicious packet streams that may create a problem, they log them for security administrators to examine. If a threat appears to be very serious, the IDS will send an alarm to security administrators. Next-generation firewall processing grew out of IDS processing methods, but NGFWs are still limited to stopping definite attacks. NGFWs are like locks, whereas IDSs are like burglar alarms.

> *Intrusion detection systems (IDSs) were created to supplement firewalls by focusing specifically on identifying suspicious transmissions.*

False Positives (False Alarms) Have you ever had a neighbor with a twitchy car alarm that went off whenever a cat walked near the car? Unfortunately, intrusion detection systems also generate many **false alarms**, which are called **false positives**. An IDS sends the average business about 10,000 security alerts per day, only a handful of which are real threats. Finding these real attacks is literally like finding a needle in a haystack.[34] Companies must invest sufficient human resources into handling alarms and reading log files. Even then, finding the one or two true attacks in a long stream of false alarms leads to frustration and flagging vigilance.

No Alternatives Although IDSs create many problems, firms today realize that it is impossible to prevent breaches automatically and that they must be able to identify and stop attacks if they succeed initially. Learning to look for meaningful patterns in IDS alerts and log files is enormously difficult, but there is no alternative.

Test Your Understanding

28. a) Do IDSs stop packets? b) Why are they painful to use? c) How do they offer a broader picture of the threat environment than NGFWs?

The Need for Intrusion Detection

- Firewalls only stop provable attack packets
- Some way is needed to identify suspicious transmissions

Intrusion Detection System (IDS) Characteristics

- Like car alarms for security
- If detect suspicious activity, send a warning or shut down the threat
- Problem of many false alarms (false positives)

FIGURE 4-20 Intrusion Detection Systems (IDSs) (Study Figure)

[34] John E. Dunn, "Average US Business Fields 10,000 Security Alerts per Day, Damballa Analysis Finds," Terchworld.com, May 13, 2014. http://news.techworld.com/security/3516426/average-us-business-fields-10000-security-alerts-per-day-damballa-analysis-finds/.

IN MORE DEPTH

Antivirus Protection

Both firewalls and antivirus programs attempt to stop attacks. However, they work at different levels. Firewalls examine packets and groups of packets. **Antivirus (AV) programs**, in contrast, examine entire files. When an e-mail message arrives at a mail server, the server may pass any attachment to an AV program for vetting.

Firewalls examine packets and groups of packets. Antivirus (AV) programs, in contrast, examine entire files.

Antivirus programs do not simply check for viruses. They also examine the attached file for worms, Trojan horses, and other forms of malware. These programs were named antivirus programs when "malware" was roughly synonymous with "virus." Although the scope of detection has broadened, the name antivirus has stuck.

Traditionally, AV programs only looked for malware signatures, which are snippets of code that let the antivirus program identify particular malware programs. Signature detection is still widely used, but it is no longer sufficient. First, the number of malware programs is now so large that the processing power to detect all known malware via signature detection would drive any computer to its knees. More fundamentally, many malware programs now mutate constantly, rewriting their code in a way that maintains functionality while making the matching of strings of characters useless.

Today, AV programs also look for behavioral patterns—things the file is attempting to do. To give an extreme example, if the file is a program that will try to reformat a computer's hard drive, that is an undeniable indication that the program is malware. Some AV programs even run the suspect program in a sandbox (environment it cannot escape from) to watch it operate.

Test Your Understanding

29. a) Distinguish between what firewalls look at and what antivirus programs look at. b) Are AV programs used to detect more than viruses? Explain. c) Distinguish between signature detection and behavioral pattern detection. d) Why is signature detection not enough?

Firewalls versus Antivirus Programs

- Firewalls analyze packets and streams of packets
- Antivirus programs analyze files

Search for All Malware, Not Only Viruses

Signatures versus Behavior

- Traditionally looked for signatures (characteristic bit patterns) for specific malware
- Malware writers now create code that mutates slightly each time it runs
- Defeats most signature detection
- Now, also look at behavioral patterns: What programs do

FIGURE 4-21 Antivirus Protection (Study Figure)

END-OF-CHAPTER QUESTIONS

Thought Questions

4-1. What are your choices if you are hit by ransomware? Which would you recommend?

4-2. a) What form of authentication would you recommend for relatively unimportant resources? Justify your answer. b) What form of authentication would you recommend for your most sensitive resources?

4-3. What is the promise of newer authentication systems?

4-4. Is the supplicant the true party or an impostor?

4-5. In digital certificate authentication, the supplicant could impersonate the true party by doing the calculation with the true party's private key. What prevents impostors from doing this?

4-6. What are the implications for digital certificate authentication if the true party's private key is stolen?

4-7. a) If someone in your firm gives you his or her password and you log into that person's account, is this hacking? Justify your answer in terms of the definition of hacking. b) If you think someone in your office is sending slanderous e-mail about you, is it hacking if you break into that person's e-mail account to see if this is true? Justify using the definition. c) If you log into a server at your bank to test their security, is this hacking? Justify using the definition.

Harder Thought Questions (You May Not Get These, but Try)

4-8. When a sales clerk accepts a credit card payment, he or she should type the last four digits of the credit card into the terminal in order for the terminal to verify that the last four digits on the card are the same as on the magnetic stripe. Why should the sales clerk not ask the customer what the last four digits are?

4-9. Keys and passwords must be long. Yet most personal identification numbers (PINs) that you type when you use a debit card are only four or six characters long. Yet this is safe. Why?

Perspective Questions

4-10. What was the most surprising thing you learned in this chapter?

4-11. What was the most difficult part of this chapter for you?

This page intentionally left blank

Chapter **5**

Ethernet (802.3) Switched LANs

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Explain basic Ethernet terminology and how Ethernet is standardized.
- Describe basic physical propagation concepts: digital and binary signaling and why they reduce transmission errors; full-duplex transmission, and parallel transmission.
- Explain the technologies of 4-pair UTP and optical fiber; compare their relative strengths and weaknesses, including cost and transmission distances.
- Design a physical network based on knowledge of transmission requirements and Ethernet physical link standards, including link aggregation.
- Describe the Ethernet II Frame. Explain basic Ethernet data link layer switch operation.
- Describe security threats to Ethernet and ways to deal with them.

ETHERNET BEGINS

Bob Metcalfe, a PhD student at Harvard University, wrote his dissertation on the new ARPANET (which would later morph into the Internet). His committee rejected it as insufficiently theoretical. Metcalfe was devastated. He had been offered a position at the Xerox Palo Alto Research Center, which was doing cutting-edge computer and network research. In particular, PARC had just built the Alto, which looked like a PC but was far more powerful. It had a full-page display and a graphical user interface using the mouse, which PARC adopted from Doug Engelbart's Augmentation Research Center at Stanford Research Institute. Apple later popularized this input device with the Macintosh.

When Metcalfe told Xerox that he would not be graduating, PARC told him to come anyway and finish his dissertation while he worked there. Metcalfe asked for a brief delay so he could first visit the University of Hawai`i's ALOHANET project. PARC accepted the delay. ALOHANET did packet transmission using radio. If two stations transmitted at the same time, their colliding packets would be garbled and would not be retransmitted. During his visit, Metcalfe analyzed the ALOHANET protocol and found ways to reduce collisions. He added the analysis to his dissertation. This time, his committee accepted it.

At Xerox, his job was to network the Altos. Metcalfe realized that his improvements to the ALOHANET protocol would permit him to run a similar network over physical transmission media. There were several physical media that he could use. To keep his options open, he referred to physical media generically as *the ether*, after a discredited nineteenth-century theory about how light propagated. He wrote software and hand-soldered printed circuit boards to make his vision real. When his **Ethernet** network became operational, it ran at 2.94 Mbps, which was enormous speed for that time.

When Xerox decided not to commercialize Ethernet, Metcalfe started his own company. In those days, there were several wired local area network standards. However, the brilliant simplicity of Metcalfe's protocol meant that Ethernet products were substantially cheaper and quicker to develop than products following competing protocols. Ethernet quickly blew the competition out of the water. Since then, Ethernet has continued its dominance in wired local area networks and has grown remarkably in speed. In this chapter, you will see many examples of how Ethernet continues to be a relatively inexpensive technology that still provides the speed and other affordances that companies need.

Ethernet is inexpensive but does what corporations need. This is its formula for dominance in wired local area networks.

INTRODUCTION

Local Area Networks

Local area networks (LANs) are networks that operate on the **customer premises**,[1] which is the property owned by the organization that uses the network. This might be a home, an entire building, a university campus, or an industrial park. On its own premises, the company can use whatever technology standards it wishes.

Local area networks (LANs) are networks that operate on a customer premises—the property owned by the organization that uses the network.

[1] "Customer premises" is always spelled as plural, although it is used as if it is singular. It's a legal jargon thing.

Operate on a customer premises

The property owned by the person or organization that uses the network

Companies can use whatever technology standards they wish

FIGURE 5-1 Local Area Networks (LANs) (Study Figure)

Test Your Understanding

1. What is a local area network (LAN)?

Perspective: Layer 1 and Layer 2 Standards

Let's begin with a brief recap of some distinctions made in the first two chapters. Figure 5-2 shows a switched Ethernet network. Ethernet is a single-network standard, so it is governed by physical and data link layer standards. The messages that travel from the source host to the destination host are *frames*, forwarding devices are *switches*, and Ethernet uses *EUI-48* data link layer addresses instead of IP addresses. The path that a frame travels through an Ethernet network is a *data link*. The transmission links that connect pairs of devices are *physical links*. *Ethernet signaling standards* govern physical layer transmission.

Today, LAN standards come from the IEEE Standards Association, through its **802 LAN/MAN Standards Committee**. Ethernet standards specifically come from the Committee's **802.3 Working Group** (Figure 5-3). Figure 5-4 notes that almost all physical and data link layer standards are open system interconnection (OSI) standards. It also notes that Ethernet standards are submitted to the International Organization for Standards (ISO) by the IEEE for acceptance as an official OSI standard. However, ISO always accepts these submissions. In fact, as soon as the 802.3 Working Group finishes a standard (and sometimes even before), vendors begin developing products.

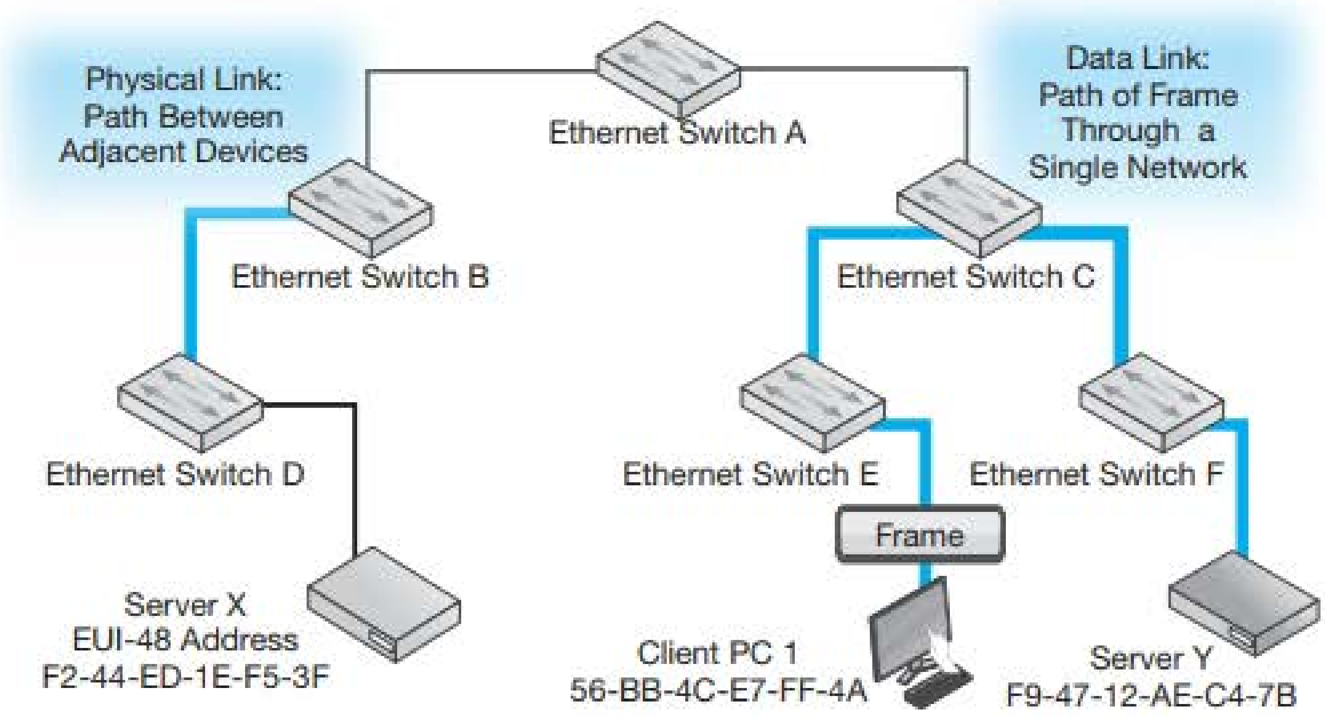

FIGURE 5-2 Switched Ethernet Network: Physical and Data Links

Requires standards at Layers 1 (wires and signals) and 2 (frames and switches)

OSI standards dominate at these layers

Ethernet standards are created by the IEEE 802.3 Working Group of the IEEE Standards Association's 802 LAN/MAN standards committee

Called 802.3 standards

Other Working Groups exist (e.g., the 802.11 WG creates Wi-Fi standards)

Submitted to ISO, which ratifies them as OSI standards

FIGURE 5-3 Ethernet Origins

Test Your Understanding

2. a) At what layers are Ethernet standards defined? b) Are Ethernet messages packets or frames? c) Are Ethernet forwarding devices switches or routers? d) Is the path an Ethernet message takes from the source host to the destination host a physical link, a data link, or a route? e) Does Ethernet use EUI-48 addresses or IP addresses? f) Why are Ethernet standards formally called 802.3 standards?

Basic Physical Layer Terminology

Workgroup Switches and Core Switches Figure 5-5 shows that Ethernet networks have two types of switches.

- **Workgroup switches** connect individual hosts to the network.
- In turn, **core switches** connect switches to other switches. An Ethernet network's collection of core switches is called the network's **core**.

Workgroup switches connect individual hosts to the network.

Core switches connect switches to other switches.

Layer	Ethernet (802.3 Standards)	Dominant Standards Architecture	Comments
Application		None	
Transport		Internet (IETF)	
Internet			
Data Link	Yes	OSI (ISO and ITU-T)	Ethernet standards created by the 802.3 Working Group, submitted to ISO for final acceptance as standards.
Physical	Yes		

FIGURE 5-4 Ethernet Standards and OSI Standards

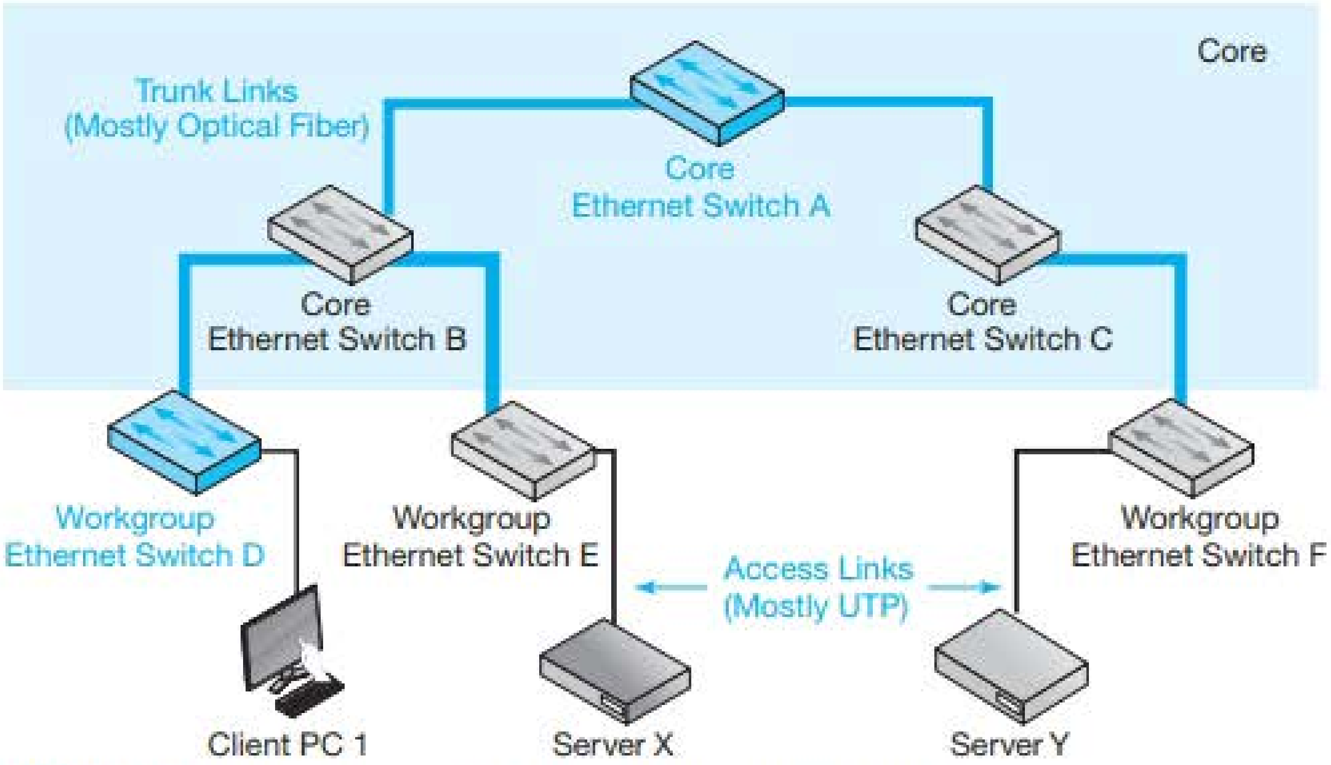

FIGURE 5-5 Types of Ethernet Switches and Transmission Links

Figure 5-6 shows a typical workgroup switch. It is 48 cm (19 inches) wide to fit into a standard equipment rack. It is 9 cm (3.5 inches) tall. Core switches are the same width and depth, but their heights typically range from about 18 cm (7 inches) to a meter (39.37 inches) tall.

Access Links and Trunk Links Just as there are two types of Ethernet switches, there are two types of physical links.

- **Access links** connect individual hosts to their workgroup switches.
- **Trunk links** connect switches to other switches.

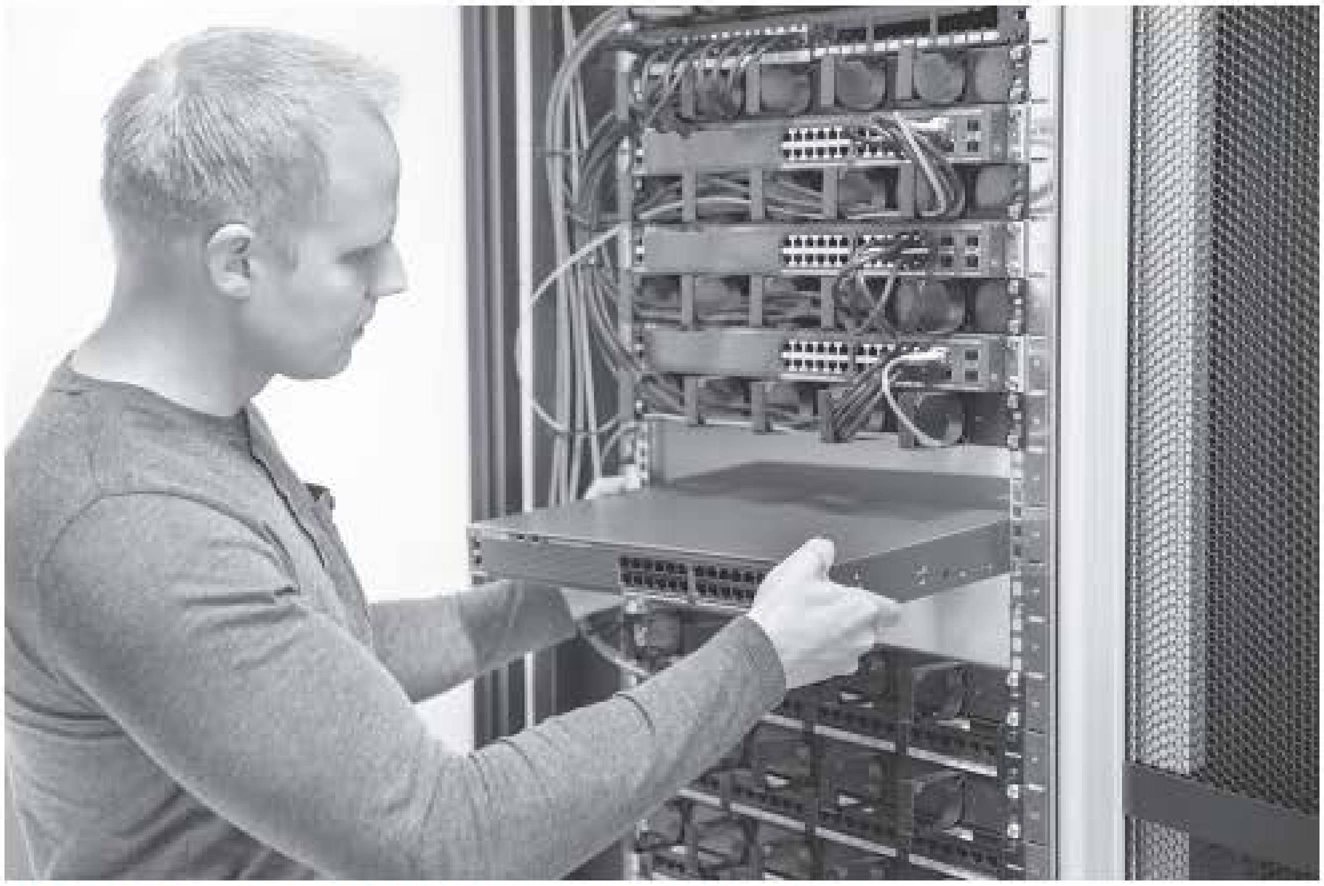

FIGURE 5-6 Ethernet Workgroup Switch with 48 Ports

Test Your Understanding

3. a) ______ switches connect users to the network. b) ______ switches connect switches to other switches. c) _____ links connect users to workgroup switches. d) ____ links connect switches to other switches.

ETHERNET PHYSICAL LAYER STANDARDS

Physical layer standards govern physical links between devices. This includes connectors, plugs, transmission media, and signaling. We look at signaling first because it introduces concepts we will need when we look at transmission media.

Test Your Understanding

4. What four things do physical layer standards govern?

Signaling

Bits and Signals A frame is a long series of bits (1s and 0s). To transmit the frame over a physical link, the sender converts these 1s and 0s into physical signals. These signals **propagate** (travel) down the transmission link to the device at the other end. That device converts the signal back into the 1s and 0s of the frame.

Binary Signaling Figure 5-7 illustrates the two main types of signaling, binary and digital signaling. **Binary signaling** has two **states** (conditions), which may be two voltage levels or light being turned on or off. One state represents a 0. The other state represents a 1. In the figure, a 1 is represented as a high voltage, and a 0 is represented as a 0 voltage. In optical signaling, a light being turned on typically represents a 1, and light being turned off typically represents a 0.

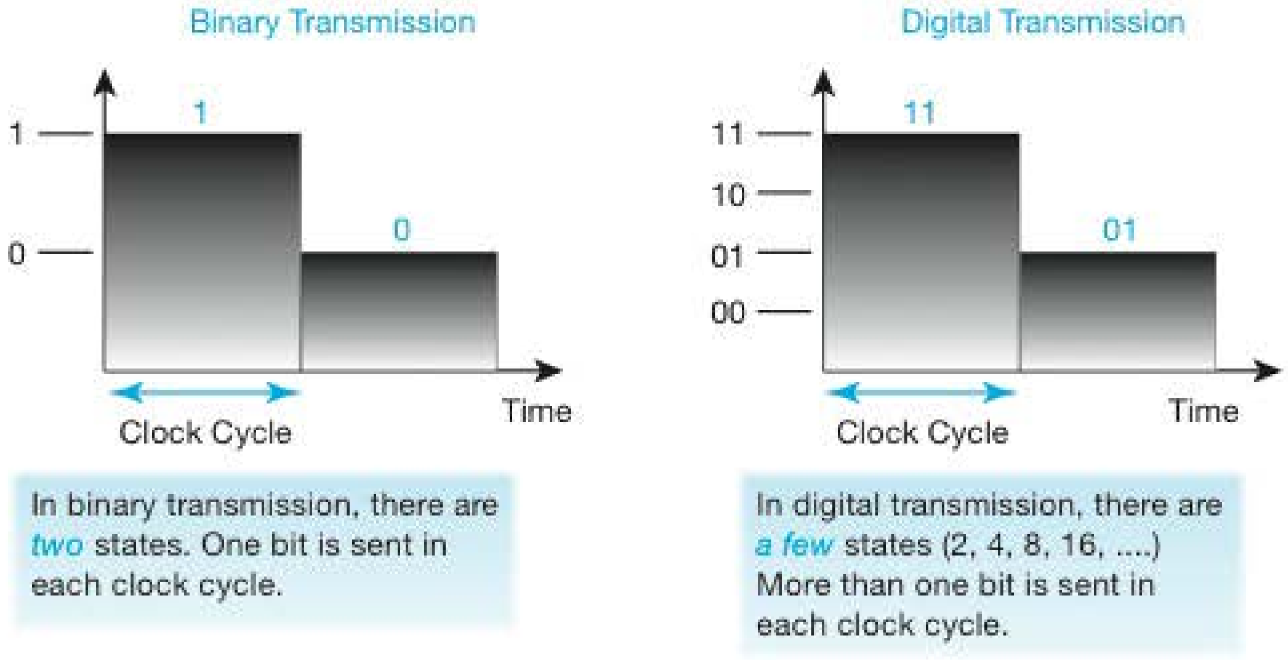

In binary transmission, there are *two* states. One bit is sent in each clock cycle.

In digital transmission, there are *a few* states (2, 4, 8, 16, ...) More than one bit is sent in each clock cycle.

FIGURE 5-7 Binary and Digital Signaling

Digital Signaling The figure also shows **digital signaling**, in which there are a *few* **alternative states** (2, 4, 8, etc.).[2] How many is "a few?" In some systems, there can be 64 or even 256 states, but the number of states is usually much lower. The number of alternative states is always a power of two—two, four, eight, sixteen, and so forth.

In binary signaling, there are two possible states.

In digital signaling, there are a few *possible alternative states (2, 4, 8, etc.).*

Adding states increases the complexity and cost of signaling. However, Figure 5-7 shows that if you have multiple states, you can send multiple bits in a single clock cycle. With two states, you can only represent a single 1 or a 0. With four states, however, the lowest state might represent 00, the next lowest state might represent 01, the next 10, and the highest 11. With four states, then, you can send two bits at a time.

In Chapter 2, we saw that the number of alternative states is two to the power of the number of bits. In symbols, this is $a = 2^b$. In digital signaling, a is the number of possible alternative states, and b is the number of bits transmitted in each clock cycle. For example, if you transmit one bit per clock cycle, then b is one, and a is 2. This is binary signaling. The two alternatives are 1 and 0. If you wish to transmit three bits per clock cycle, you need 2^3 (8) alternative states.

$$\textit{Equation 1: } a = 2^b$$

Binary Is a Special Case of Digital We have talked about binary and digital transmission systems as if they are different. Actually, binary transmission is a *special case* of digital transmission. In binary transmission, *few* means "two." All transmission is digital.

Binary transmission is a special case of digital signaling. Not all signaling is binary, but all signaling is digital.

Clock Cycles Note the term *clock cycle* in the figure. When the sender transmits, it holds the transmission state constant for a brief period. This period is the **clock cycle**. The receiver can read the signal at any time within the clock cycle. As clock cycles get shorter, more state signals can be transmitted per second, but it becomes more difficult to read them at the receiving end. To get a feeling for this, note that if the transmission speed is one gigabit per second (1 Gbps) and binary transmission is used, each clock cycle is only *one billionth* of a second!

[2] If "bi" means two, where does "digital" come from? It comes from the fact that we call our 10 fingers *digits*. In fact, some early computer systems operated on Base 10 arithmetic, the same arithmetic that we 10-fingered people use. Very quickly, however, the advantages of building computers and transmission systems that used two or a multiple of two states brought about binary and digital computation and also binary and digital signaling.

Test Your Understanding

5. a) Distinguish between binary and digital signaling. b) if you wish to transmit three bits per clock cycle, how many states must the system have? (Answer: 8) c) If you want to transmit five bits per clock cycle, how many states must the system have? d) If you know that a system has 16 states, how many bits can it send per clock cycle? (Answer: 4) e) The 802.11ac Wi-Fi standard uses 256 states. How many bits can it send per clock cycle? f) Every time you double the number of states, how many more bits can you transmit? (The answer is not in the text.) g) Why is the signal held constant over each clock cycle? h) How long is the clock cycle if I transmit at 100 Mbps per second using signaling with four states?

4-Pair Unshielded Twisted Pair (UTP) Physical Links

Physical links connect adjacent devices along the data link in a single network. Physical layer standards specifically govern transmission media, connectors, and plugs. Ethernet uses two types of cabling today. These are 4-pair UTP and optical fiber.

4-Pair Unshielded Twisted Pair (UTP) Cables Ethernet copper wire is called **4-pair unshielded**[3] **twisted pair (UTP)** cabling because the cord contains eight wires arranged in four pairs. Figure 5-8 shows that two wires of each pair are twisted around each other several

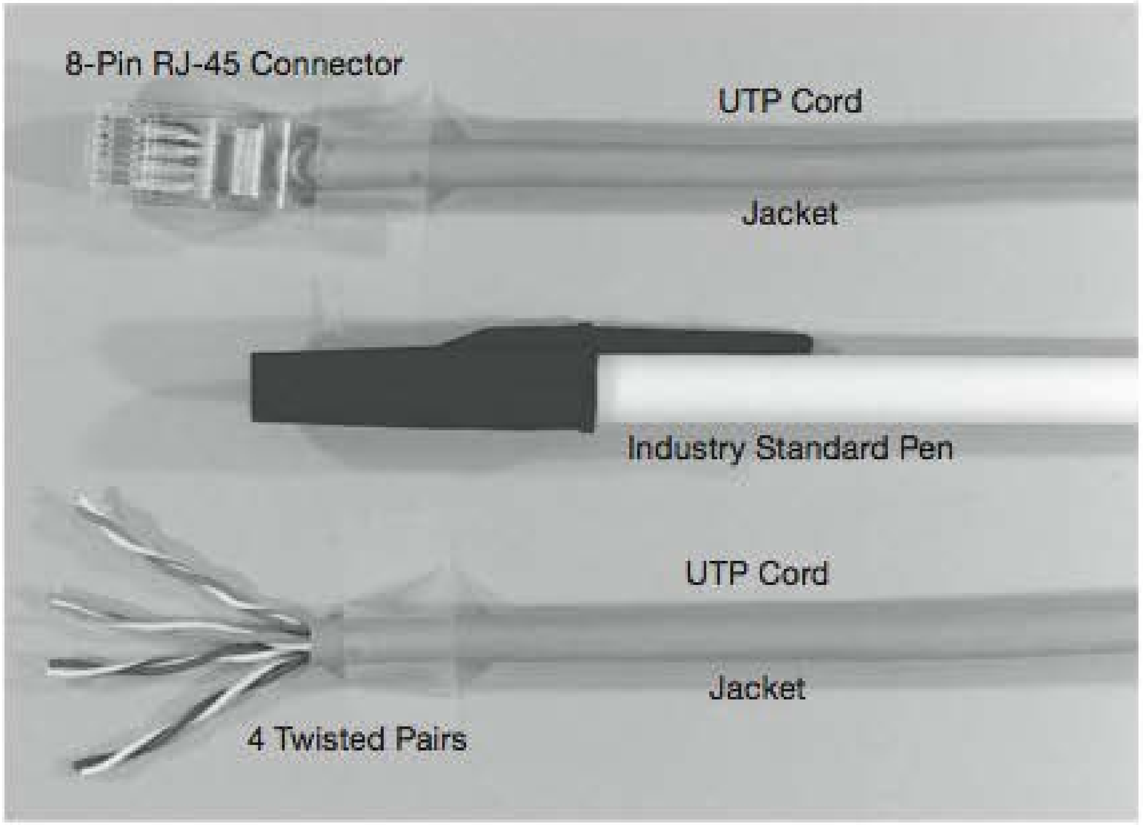

FIGURE 5-8 4-Pair Unshielded Twisted Pair (UTP) Ethernet Cable

[3] Ethernet cable is unshielded. To operate in harsh electromagnetic environments, cords may be protected by surrounding the entire cord and perhaps even individual pairs with metal foil shielding. Although placing tinfoil on your head will not protect you from the government eavesdropping on your thoughts, metal foil prevents the electromagnetic background energy from interfering with the transmitted signal. Today, shielded cabling is rare.

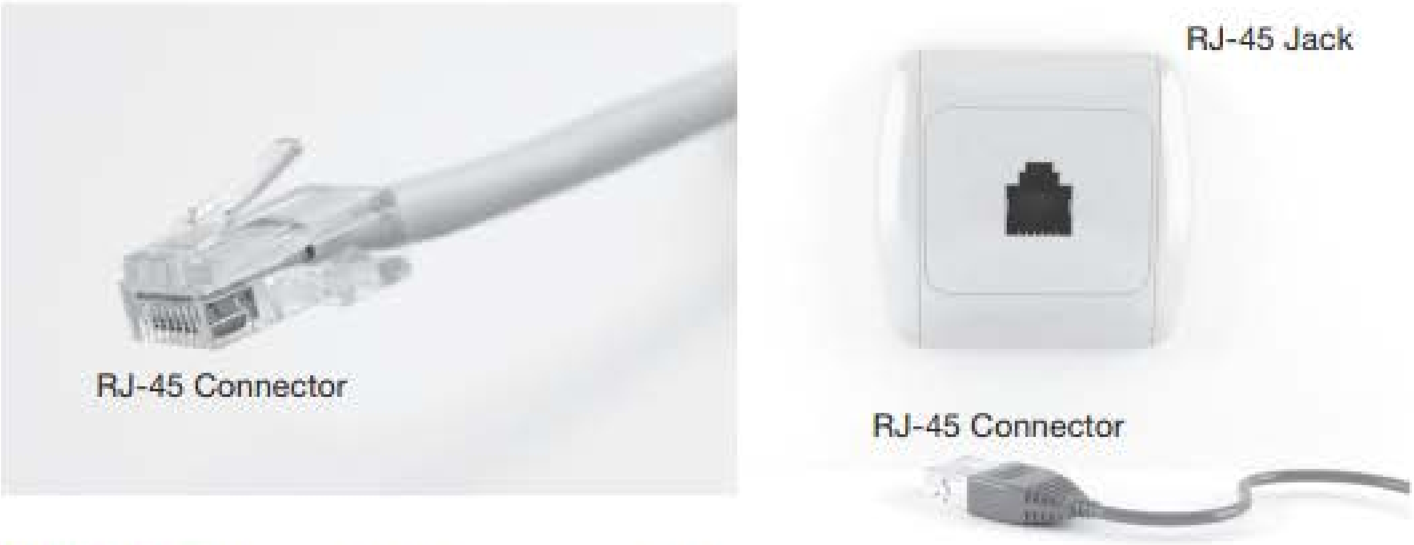

FIGURE 5-9 Ethernet (RJ-45) Connector and Jack

times per inch.[4] Figure 5-9 shows the **RJ-45 connectors** and **RJ-45 jacks** that 4-pair UTP uses. These cords are popularly called **Ethernet cords** because Ethernet is their main use today, and their connectors and jacks are popularly called **Ethernet connectors** and **Ethernet jacks**.

Parallel Transmission Signals are sent by changing voltage or other characteristics of an electrical signal. Ethernet transmits on all four pairs in each direction simultaneously. This is **parallel transmission**. As Figure 5-10 shows, Ethernet transmits four times as fast as it could if it only had a single pair.[5] The benefit of parallel transmission is higher speed.

The benefit of parallel transmission is higher transmission speed.

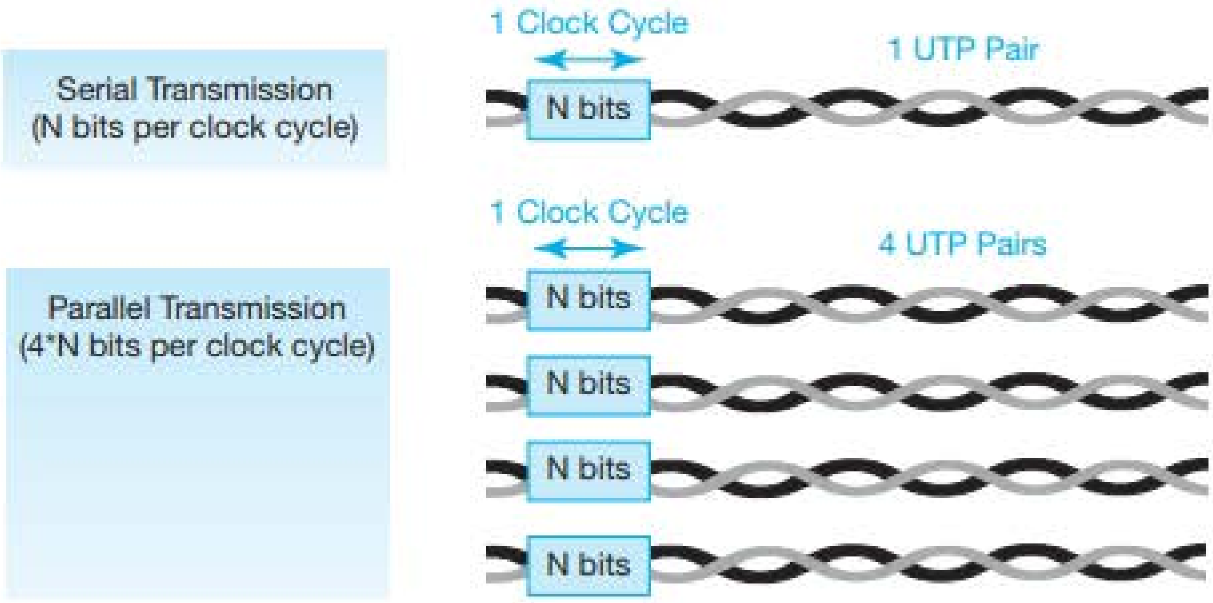
Serial Transmission (N bits per clock cycle)
1 Clock Cycle
1 UTP Pair
N bits
Parallel Transmission (4*N bits per clock cycle)
1 Clock Cycle
4 UTP Pairs
N bits
N bits
N bits
N bits

FIGURE 5-10 Parallel Transmission in Ethernet

[4] The two wires of each pair are twisted around each other because it limits the effects of nearby electromagnetic interference from lights, electrical motors, and other wire pairs, even in the same 4-pair cable. In the nineteenth century, Alexander Graham Bell realized that if you twist the two wires in a pair, interference adds to the signal on half of the twist and subtracts from the signal on the other half. The two will cancel out. Does it work this perfectly? No, but it works quite well.

[5] Ethernet is not the only transmission technology to use parallel transmission. In the past, many printer interfaces used eight or more transmission lines in each direction. Most computers, in turn, connect their components with a transmission bus that has 100 or more wires in parallel.

Each pair radiates radio signals, dissipating the signal.
This causes attenuation, which increases with propagation distance.

FIGURE 5-11 Radiative Attenuation in 4-Pair UTP

Radiative Attenuation Ethernet cable consists of long copper wires. This makes it an excellent antenna. As the signal travels down the cable, some of the signal radiates away, dissipating the signal's energy. Dissipation grows with distance (Figure 5-11). Beyond some distance, the signal becomes unreadable.

Test Your Understanding

6. a) How many wires are there in Ethernet cable? b) How is each pair organized? c) What are the two names for connectors and jacks? d) How does Ethernet use parallel transmission? e) What is the benefit of parallel transmission? f) What propagation problem limits transmission distance in 4-pair UTP?

Maximum Cord Distance How far can a UTP cord carry a signal? Figure 5-12 shows that maximum transmission distance depends on two things. One is the quality of the UTP cable. In increasing order of quality, there are **Category 5e**, **Category 6**, and **Category 6A**.[6] These are normally called **Cat 5e**, **Cat 6**, and **Cat 6A**.[7] Today, nearly all

Ethernet Signaling Standard	Transmission Speed	Cable Quality Category	Maximum Cord Length
100BASE-TX	100 Mbps	Cat 5e, 6, 6A	100 meters
1000BASE-T	1 Gbps	Cat 5e, 6, 6A	100 meters
2.5GBASE-T[a]	2.5 Gbps	Cat 5e, 6, 6A	100 meters
5GBASE-T[a]	5 Gbps	Cat 5e, 6, 6A	100 meters
10GBASE-T	10 Gbps	Cat 6, Cat 6A	55 meters
10GBASE-T	10 Gbps	Cat 6A	100 meters

[a] New standards designed in response to faster Wi-Fi Access Points.

FIGURE 5-12 Ethernet Signaling Standards, Transmission Speed, UTP Cable Quality, and Maximum Cord Length for 4-Pair UTP

[6] The 802.3 Working Group does not create wiring quality standards. These standards come from the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC). The 802.3 Working Group adds signaling standards.

[7] Category 5e is Category 5 enhanced. Category 6A is Category 6 Augmented. Why not make both enhanced or Augmented? The answer is that the standards were created at different times by different standards agencies.

installed Ethernet cable is Cat 5e and Cat 6. As the figure shows, both can carry Ethernet signals 100 meters at 1 Gbps.

The 10 Gbps Problem For 10 Gbps (10GBASE-T), however, Cat 5e cannot be used at all, and Cat 6 quality cable can only span 55 meters. This is too short for many situations. Cat 6A will carry 10GBASE-T up to 100 meters, but not much Cat 6A is installed.

2.5 Gbps and 5 Gbps for Wi-Fi Access Points There was not much demand for speeds beyond 1 Gbps until recently. However, wireless access points have been growing in speed. In Chapter 6, we will see that the newest Wi-Fi standard can exceed speeds of 6 Gbps. Companies need to connect these new access points via 4-pair UTP to an Ethernet switch. In response, the 802.3 Working Group is developing two new standards. They are designed to carry signals at 2.5 Gbps and 5 Gbps. More importantly, they can use existing installed Cat 5e and Cat 6 cabling. These two standards are 2.5GBASE-T and 5GBASE-T.[8]

Test Your Understanding

7. a) If you need to transmit 600 million bits per second 90 meters, what signaling standard and UTP quality standard would you use? b) If you need to transmit 7 Gbps 40 meters, what signaling standard and UTP quality standards could you use? c) Which would you probably use if you wanted to send the signal over installed cabling? d) If you need to transmit 7 Gbps over 120 meters, what signaling standard and UTP quality standard would you use?
8. For what specific purpose were the 2.5GBASE-T and 5GBASE-T standards developed?

Optical Fiber (Fiber)

In **optical fiber**, light signals travel through glass. Typically, light is turned on in a clock cycle to represent 1 and off to represent 0. Figure 5-13 shows that a **fiber cord** consists of two fiber **strands**—one for propagation in each direction. Two strands permit simultaneous two-way transmission, which is called **full-duplex transmission**.

Optical Fiber Cord with Two Strands for Full Duplex Communication

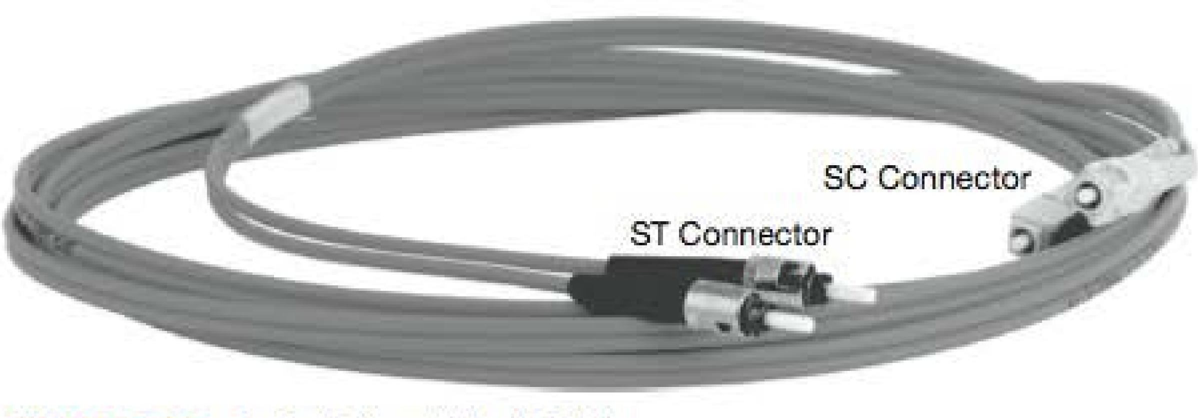

FIGURE 5-13 Optical Fiber (Fiber) Cable

[8] USB provides both data transmission and a little electrical power. Power over Ethernet (POE) optionally provides limited power to Ethernet devices, saving the cost of running power to them separately. POE gave enough wattage for early access points, but some new access point can be powered by Ethernet.

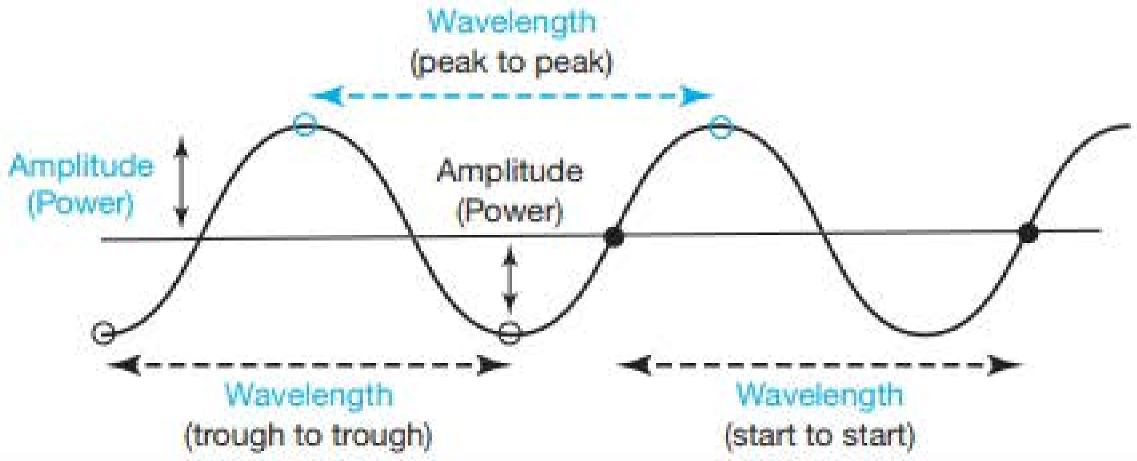

Light travels in waves.
A wave's amplitude is its power.
Optical fiber transmission is described in terms of wavelength.
Wavelength is the physical distance between comparable points on adjacent cycles.
Wavelengths for optical fiber are measured in nanometers (nm).
LAN fiber uses 850 nm almost exclusively because it is cheap and usually sufficient.
Wide area networks use 1,310 and 1,550 nm light to support longer distances.

FIGURE 5-14 Light Transmission Metrics

There is only a single UTP connector standard, but there are many types of optical fiber connectors.[9]

Figure 5-14 shows that light waves are measured in terms of **wavelength**. This is the physical distance between comparable parts of two consecutive waves. This might be the beginning of one cycle to the beginning of the next, two consecutive peaks, two consecutive troughs, and so forth. The **amplitude** of the wave, in turn, is its power (brightness).

Wavelengths are measured in **nanometers (nm)**. In optical fiber, there are three wavelength "windows" in which light travels especially well. These are centered on 850 nm, 1,310 nm, and 1,550 nm. Each window is about 50 nm wide.

Test Your Understanding

9. a) How does fiber usually transmit a 1? b) How do fiber cords typically provide full-duplex transmission? c) In what units are light wavelengths measured? d) What are the three wavelength windows used in fiber transmission? e) What is amplitude?

Multimode Fiber Propagation Limitations We saw in Figure 5-11 that propagation distance is limited by radiative attenuation. Figure 5-15 shows that this is not true in optical fiber. Signals travel through an inner glass core covered by an outer glass cladding. There is total internal reflection when a light ray hits the core/cladding boundary. There is no signal loss.

However, there is another propagation problem in the multimode fiber used in LANs. For technical reasons, light rays can only enter the core at a few angles. These

[9] Surprisingly, media standards do not specify connectors and plugs. However, Figure 5-13 shows that fiber can use different connectors at each end. Fiber can work with whatever type of fiber port a switch or router uses.

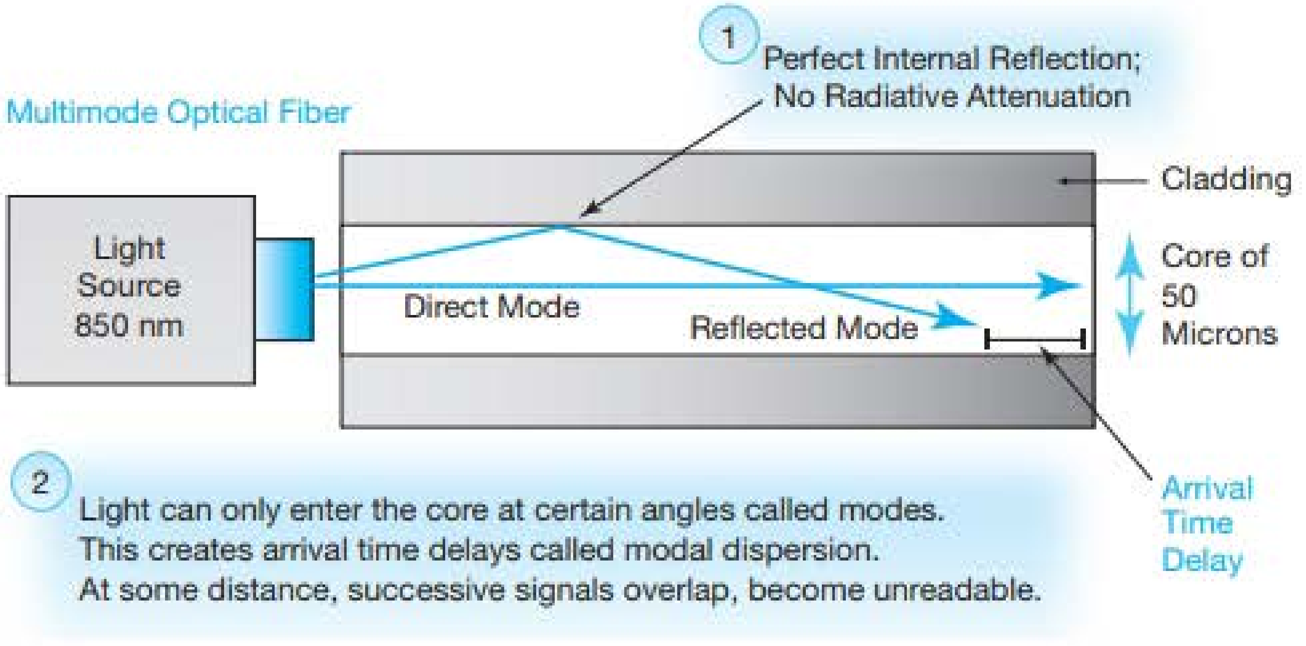
Multimode Optical Fiber
1
Perfect Internal Reflection;
No Radiative Attenuation
Cladding
Light
Source
850 nm
Core of
50
Microns
Direct Mode
Reflected Mode
Arrival
Time
Delay
2
Light can only enter the core at certain angles called modes.
This creates arrival time delays called modal dispersion.
At some distance, successive signals overlap, become unreadable.

FIGURE 5-15 Modal Dispersion in Multimode Fiber

light rays are called **modes**. LAN fiber cores are 50 microns in diameter—about the diameter of a human hair. Cores of this diameter can admit several modes, giving rise to the name **multimode fiber**.

In Figure 5-15, two modes are shown. One travels straight down the middle of the core. The other reflects at a core/cladding boundary. The reflected mode travels a longer distance and so takes slightly longer to reach the end. This time gap is called **modal dispersion**. If modal dispersion becomes too large, the modes of successive light pulses will overlap too much for the receiver to understand the signal. (The Box "Fiber Modes and Fiber Wavelengths" shows that fiber with a much smaller diameter only allows the direct mode. This eliminates modal dispersion, allowing signals to travel greater distances without becoming unreadable. However, this single-mode fiber is expensive. Multimode fiber distances are generally fine in LANs.)

Test Your Understanding

10. a) What is a mode? b) What is multimode fiber? c) What limits transmission distance in multimode fiber?

Maximum Optical Fiber Transmission Distances Earlier, we saw that limits on propagation distance for 4-pair UTP depend on Ethernet signaling standard, speed, and UTP cord quality. Figure 5-16 does the same for multimode optical fiber. For multimode fiber, the quality standards are **optical multimode (OM)** designations. OM3 and OM4 fiber are sold today. Figure 5-16 shows that at speeds up to 10 Gbps, both quality levels of multimode fiber easily span the 200 to 300 meters commonly needed in LANs. For very high speeds, however, maximum transmission distances become uncomfortably short.

In multimode fiber, quality standards are optical multimode (OM) standards. OM3 and OM4 multimode fiber are sold today.

Multimode Fiber Quality Standard	Ethernet Signaling Standard	Light Wavelength	Transmission Speed	Maximum Transmission Distance
OM3	1000BASE-SX	850 nm	1 Gbps	550 meters
	10GBASE-SR	850 nm	10 Gbps	300 meters
	40GBASE-SR4	850 nm	40 Gbps	100 meters
	100GBASE-SR10	850 nm	100 Gbps	100 meters
	100GBASE-SR4	850 nm	100 Gbps	70 meters
OM4	1000BASE-SX	850 nm	1 Gbps	1,000 meters
	10GBASE-SR	850 nm	10 Gbps	440 meters
	40GBASE-SR4	850 nm	40 Gbps	125 meters
	100GBASE-SR10	850 nm	100 Gbps	150 meters
	100GBASE-SR4	850 nm	100 Gbps	100 meters

FIGURE 5-16 Ethernet Multimode Fiber Signaling Speed, Optical Fiber Quality Standards, and Maximum Transmission Distance[10]

To see how to use the figure, suppose you need to provide 10 Gbps signaling over 250 meters.

- On OM3 cabling, 10GBASE-SR has a maximum transmission distance of 300 meters, and OM4 fiber raises this to 440 meters. Both would work, but OM3 fiber is less expensive and so would be the preferred choice.
- In turn, if the required distance is 330 meters, only more expensive OM4 fiber would work because the maximum propagation distance for OM3 at 10 Gbps is only 300 meters.
- Now suppose that the maximum transmission distance is 85 meters, which is close to the length of a football field. OM3 and OM4 fiber would both work. However, from Figure 5-12, so would Cat 6A UTP. UTP is less expensive than fiber, so when UTP can do the job, it is the correct choice.

In Figure 5-16, all choices use multimode fiber and 850 nm light. At 100 Gbps, maximum cord distances fall below the traditional minimums for LAN physical links. Companies may now have to use more expensive 1,310 nm and 1,550 nm light sources that carry signals farther in some of their links. They may even begin to implement some links using single-mode fiber, which is described in the Box, "Fiber Modes and Light Wavelengths."

[10] Earlier, we looked at parallel transmission in the case of 4-pair UTP. The 40GBASE-SR4 and 100GBASE-SR10 and 100BASE-SR4 optical fiber standards also use parallel transmission. The 40GBASE-SR4 standard uses four fiber strands in each direction, each operating at 10 Gbps; this gives a total of 40 Gbps. The 100GBASE-SR10 standard uses ten strands in each direction, also transmitting at 10 Gbps per strand. The 100GBASE-SR4 standard uses four strands operating at 25 Gbps. SR, by the way, means Short Range, indicating that it is designed for LANs.

The Coming Explosion in Multimode Fiber Speed Standards In its first 27 years, the 802.3 Working Group only produced six speed standards. We are in the middle of a very brief period in which Ethernet will get an explosion in new standard speeds. Earlier, we saw 2.5GBASE-SX and 5BASE-SX, which were created for the new faster Wi-Fi access points. New higher-speed fiber standards, in turn, are needed for hyperscale (very large) server farms. The 40GBASE-SX and 100GBASE-SX standards are already in place, but companies want even faster standards as well as more standards between the highest and the lowest. The new standards will probably focus on 25 Gbps, 200 Gbps, and 400 Gbps. Fortunately, few corporations have hyperscale data centers, so the relevance of these new extremely high-speed standards should be limited.

Test Your Understanding

11. a) If I wish to run Ethernet over fiber using 1000BASE-SX signaling over 500 meters, what are my options? (Answer: Both OM3 or OM4 cabling will be sufficient.) b) Which should I choose? Justify your answer. c) If I wish to run Ethernet over fiber using 100GBASE-SR signaling over 100 meters, what options do I have? d) Which should I choose? Justify your answer. e) If I wish to run Ethernet over fiber using 100GBASE-SR10 signaling over 70 meters, what options do I have? f) What is the farthest I can transmit a signal with 100GBASE-SR signaling? g) What is the quality designator for multimode optical fiber?

Link Aggregation (Bonding)

Ethernet speeds have traditionally increased by factors of 10 (10 Mbps, 100 Mbps, 1 Gbps, 10 Gbps). What should you do if you only need slightly more speed than a certain standard specifies? For example, what if you have a pair of 1000BASE-SX switches that you need to connect at 1.8 Gbps? You could replace the switches with 10GBASE-SR switches. (In some cases, you can replace ports or groups of ports on a switch.) However, upgrading by a factor of 10 can be expensive.

Figure 5-17 illustrates that a company can also *install two or more* UTP or fiber trunk links to connect a pair of 1000BASE-SX switches. The IEEE calls this **link aggregation**.

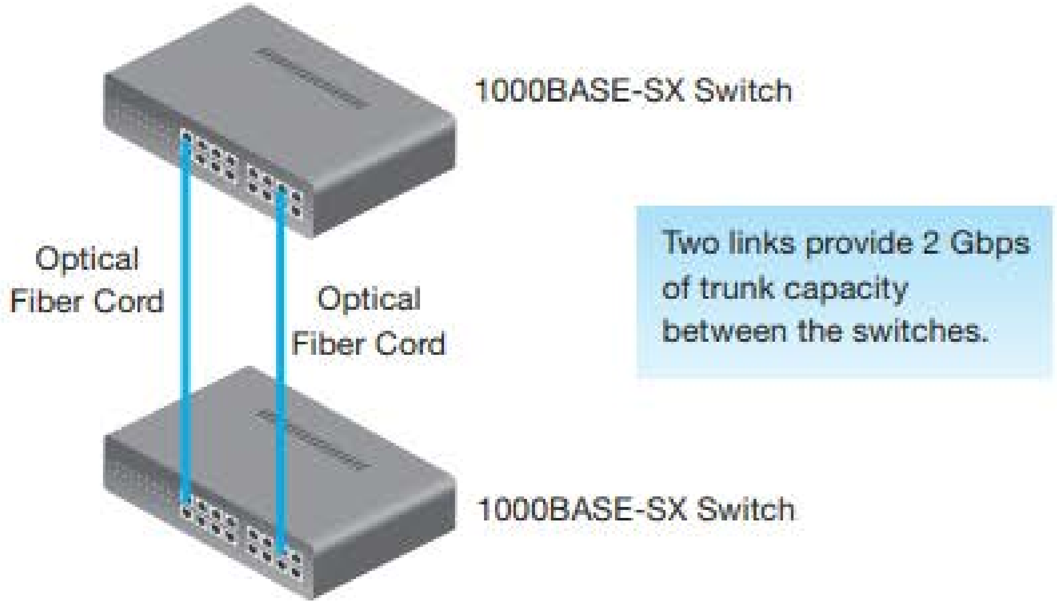

FIGURE 5-17 Link Aggregation (Bonding)

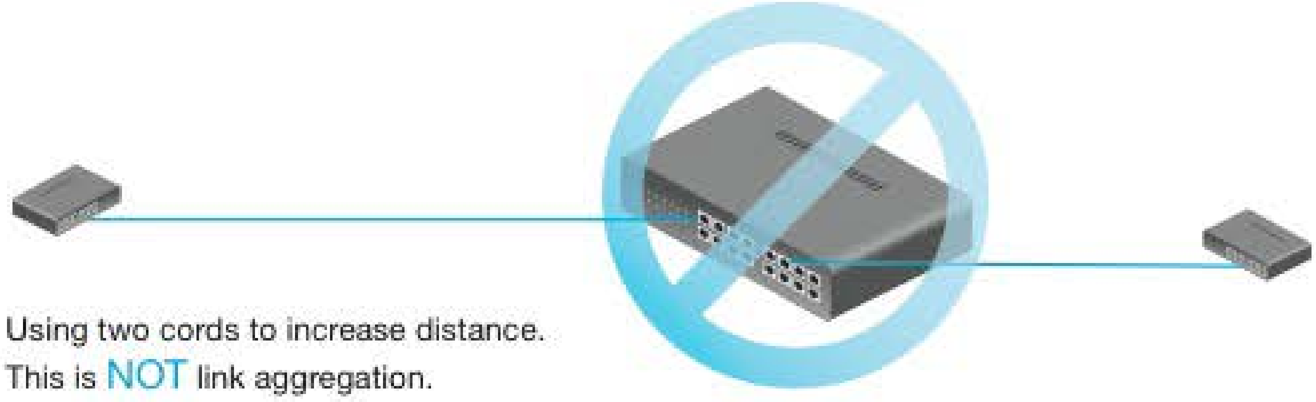

Using two cords to increase distance.
This is NOT link aggregation.

FIGURE 5-18 Link Aggregation Increases Speed, not Distance

Networking professionals also call it **bonding**. Ethernet supports link aggregation for both UTP and fiber ports. If you need 1.8 Gbps of capacity to connect two switches, you can use two bonded fiber cords and 1000BASE-SX signaling.

A common mistake in understanding link aggregation is to confuse increasing speed—which link aggregation is designed to do—and increasing distance, which it distinctly *does not* do. Figure 5-18 shows this mistaken understanding.

Link aggregation uses existing ports, so it usually costs much less than purchasing new faster switches. You do have to add another physical link or two, but cords are cheap compared to switch upgrades. Still, after two or three aggregated links, the company should compare the cost of further link aggregation with the cost of a bigger increase in capacity by moving up to the next Ethernet speed. A 10-fold increase is likely to be a solution that lasts a long time.

Test Your Understanding

12. a) If I wish to connect two switches with fiber at a speed of 30 Gbps, what options do I have? b) Which would you choose? Justify your answer.

Perspective on Purchasing Physical Links in Ethernet

Network budgets are growing slowly if at all. At the same time, demand for service is growing rapidly. Companies must spend their money very carefully because overspending on some physical links will deny funds for others. The key general principle of network design is, "Choose the least expensive option that will achieve the required speed." If 4-pair UTP can do it, don't use fiber. If OM3 fiber can do it, don't buy OM4. If multimode fiber can do it, don't use single-mode fiber. If link aggregation is cheaper than going up to the next speed standard, use it. Of course, requirements should reflect not only today's traffic but also the likely increase in demand over a reasonable future time frame. There may be additional considerations in some cases. For example, a firm may have a policy of only using OM4 fiber for future links. However, learn the key general principle.

The key general principle of network design is, "Choose the least expensive option that will achieve the required speed."

Test Your Understanding

13. Explain the key general principle of network design and why you should follow it.

IN MORE DEPTH

Fiber Modes and Light Wavelength

Single-Mode and Multimode Fiber

Optical fiber offers single-mode and multimode fiber technology. Figure 5-19 shows that the main difference is core diameter. Multimode fiber has a "thick" core diameter of 50 microns (millions of a meter). This is roughly the thickness of human hair. Single-mode fiber has a core diameter of only 8.3 microns.

As we saw earlier in the chapter, the "thick" 50-micron core diameter of multimode fiber permits multiple modes to enter the fiber. This leads to modal dispersion, which limits distance. As core diameter decreases, fewer angle modes can enter the core. At 8.3 microns, only a single mode can enter the core. Consequently, fiber with a diameter of 8.3 microns is called **single-mode fiber**. In single-mode fiber, there is no modal dispersion, so signals travel much farther than they do in multimode fiber. The only remaining limitation is absorption of the light by the core's glass. This glass is very pure, so this **absorptive attenuation** is small. Single-mode fiber can often span many kilometers.

On the downside, single-mode fiber is more expensive than multimode fiber to buy and install. In addition, single-mode fiber transmission normally eschews inexpensive 850 nm light sources in favor of more expensive 1,310 nm and 1,550 nm light sources, which can send the signal farther. Until now, multimode fiber using 850 nm light sources has been fine for LANs. In hyperscale server farms, however, we can expect more single-mode fiber operating with greater light wavelengths.

Test Your Understanding

14. a) Compare relative cost and maximum propagation distance for multimode and single-mode fiber. b) Why does multimode fiber dominate LAN installations? c) For LAN fiber, what is the dominant signal wavelength? d) Why does this wavelength dominate?

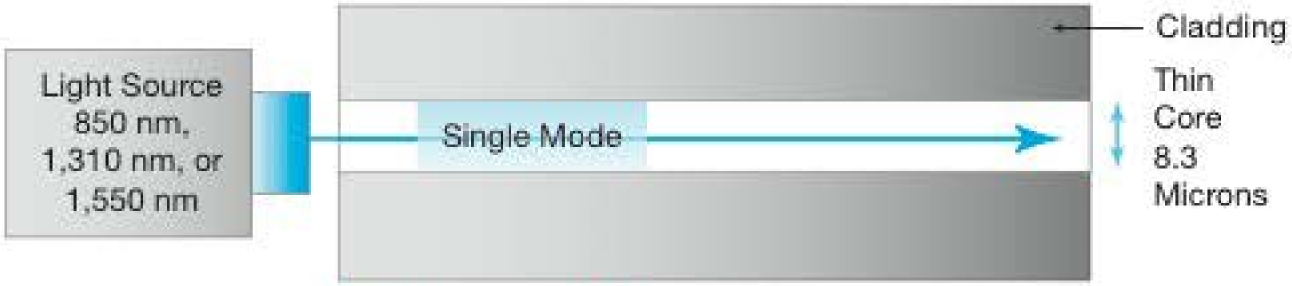

There is no modal dispersion
There is only the absorptive attenuation of the glass core
This is very small, so distance limits are very large

FIGURE 5-19 Single-Mode Fiber

THE ETHERNET DATA LINK LAYER SWITCHING AND FRAME SYNTAX STANDARD

Single switched networks, like all single networks, require standards at the physical and data link layers. We have just seen that Ethernet has *many physical layer* standards. However, Figure 5-20 shows that Ethernet only has a single major data link layer standard. This is the **802.3 MAC Layer Standard**.

Physical Link and Data Link Length Restrictions

In the previous section, we focused on distance limitations for *physical links* in Ethernet. Fortunately, network design focuses on *data links* between two hosts or a host and a router. For data links, there is no distance limitation.

Figure 5-21 shows how multiple physical links are organized into a data link.

- The source sends a signal that is "High-Low." It transmits using 1000BASE-T signaling over the Cat 5e UTP link to the first switch. The signal gets distorted, but it is still readable as a 1 or 0 up to 100 meters.
- The first switch does not merely amplify the distorted signal. It completely **regenerates** the signal. It sends a brand-new High-Low signal. The link between the first and second switches uses OM4 fiber. With 1000BASE-SX, the signal can travel up to 1,000 meters.
- The second switch, like the first, regenerates the arriving signal so that it can travel up to 100 meters to the destination host using 1000BASE-T signaling over Cat 6 UTP.

In the figure, the maximum length of the data link between the source and destination host is 1,200 meters. This can easily be lengthened by adding more switches and physical links. The maximum distances we saw earlier in the chapter were only for physical links. There is no maximum distance for data links in switched networks, so there is no limit to the size of switched networks.

The maximum distances we saw earlier in the chapter were only for physical links. There is no maximum distance for data links in switched networks, so there is no limit to the size of switched networks.

Switches versus Transmission Lines Suppose you must connect two end points, given a speed and distance requirement. Suppose that you can use a single run of expensive optical fiber. Suppose that you can use two runs of less expensive 4-pair

Layer	Standard				
Data Link Layer	802.3 MAC Standard				
Physical Layer	1000BASE-T	1000BASE-SX	10GBASE-SR	100GBASE-SR4	. . .

FIGURE 5-20 Ethernet Standards for Signaling and Frame Forwarding

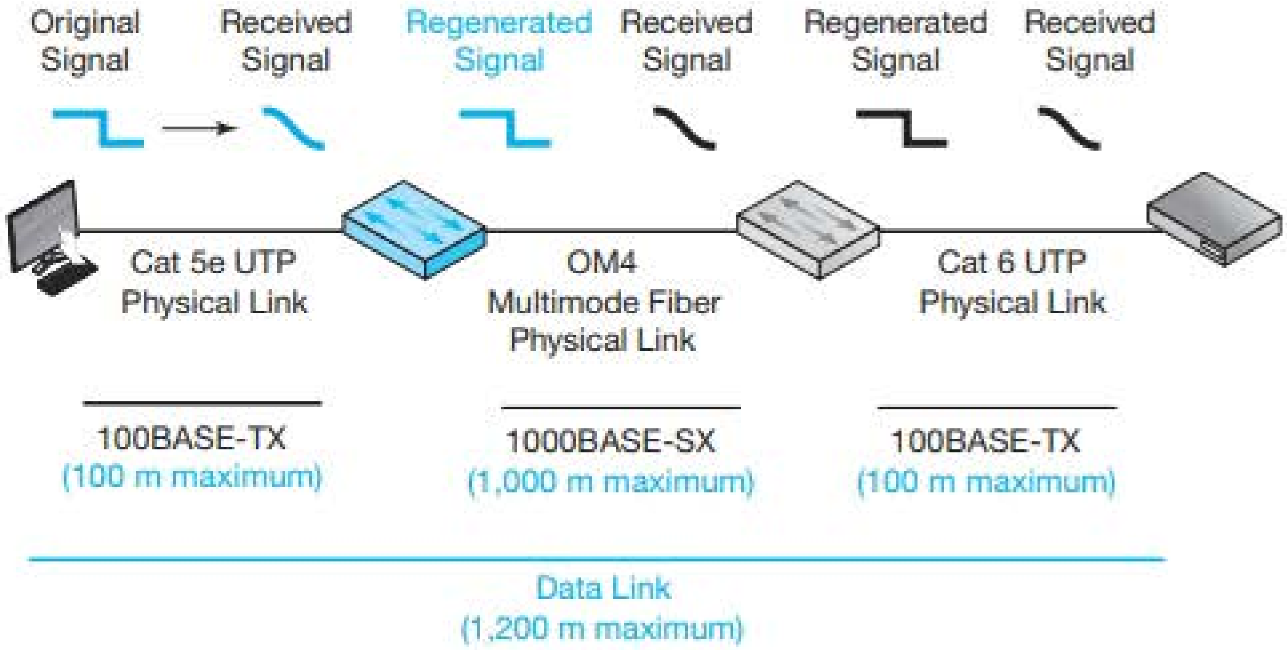

FIGURE 5-21 Distances for Physical Links versus Data Links

UTP plus an intermediate switch. Figure 5-22 illustrates this situation. Which should you select? The answer is that if a single physical link will do, adding an intermediate switch almost never makes sense economically. Switches are much more expensive than physical links.

Test Your Understanding

15. a) Are the maximum distances for UTP and optical fiber transmission shown in Figure 5-12 and Figure 5-16 distance limits for physical links or data links? b) In Figure 5-21, what would the maximum data length be if the physical link on the left was OM4 fiber? c) If you need to span 600 meters at 1 Gbps, what options do you have? (Include the possibility of using an intermediate switch.) d) How would you decide which option to choose? e) If a distance can be spanned by UTP or optical fiber, why would you almost never add an intermediate switch?

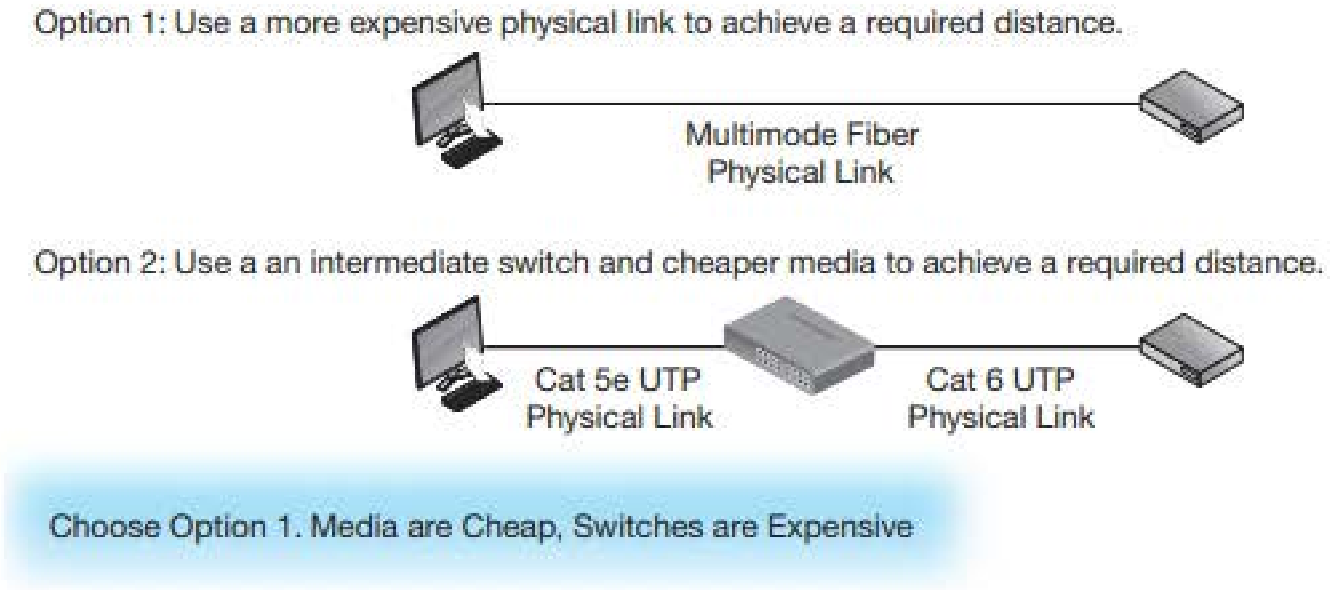
Option 1: Use a more expensive physical link to achieve a required distance.
Multimode Fiber Physical Link
Option 2: Use a an intermediate switch and cheaper media to achieve a required distance.
Cat 5e UTP Physical Link
Cat 6 UTP Physical Link
Choose Option 1. Media are Cheap, Switches are Expensive

FIGURE 5-22 The Undesirability of Attaining the Required Distance with an Intermediate Switch

Ethernet Data Link Layer Switch Operation

In this section, we discuss the basic data link layer operation of Ethernet switches. This is also governed by the 802.3 MAC Layer Standard. In the section after this one, we discuss other aspects of Ethernet switching that a firm may or may not use.

Frame Forwarding Figure 5-23 shows an Ethernet LAN with three switches. Larger Ethernet LANs have hundreds of switches, but the operation of individual switches is the same no matter how many switches there are. Each individual switch makes a **switching decision** about which port to use to send the frame back out to the next switch or to the destination host.

In the figure, Host A1 (we abbreviate the address) wishes to send a frame to Host E5. This frame must go to Switch 1, then Switch 2, and then Switch 3. Switch 3 will send the frame to Host E5.

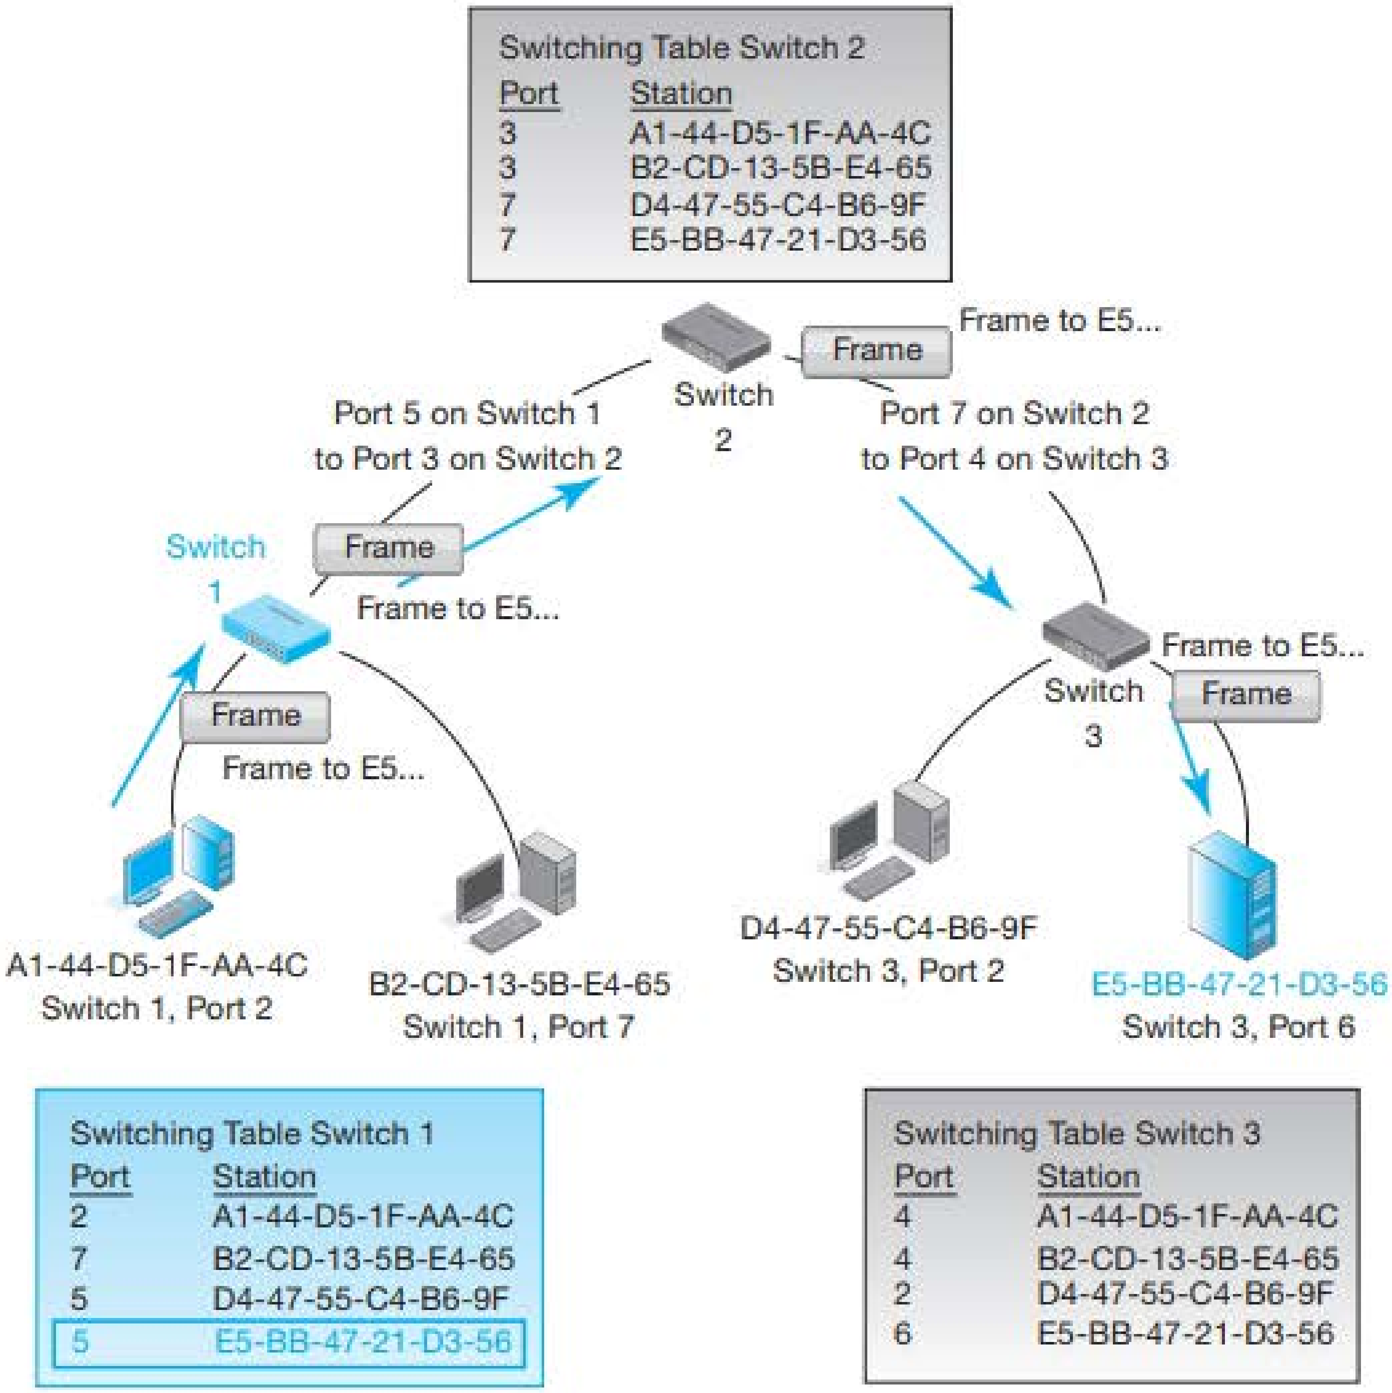

Switching Table Switch 2

Port	Station
3	A1-44-D5-1F-AA-4C
3	B2-CD-13-5B-E4-65
7	D4-47-55-C4-B6-9F
7	E5-BB-47-21-D3-56

Switching Table Switch 1

Port	Station
2	A1-44-D5-1F-AA-4C
7	B2-CD-13-5B-E4-65
5	D4-47-55-C4-B6-9F
5	E5-BB-47-21-D3-56

Switching Table Switch 3

Port	Station
4	A1-44-D5-1F-AA-4C
4	B2-CD-13-5B-E4-65
2	D4-47-55-C4-B6-9F
6	E5-BB-47-21-D3-56

FIGURE 5-23 Multi-Switch Ethernet Operation

To begin this process, Host A1 puts E5-BB-47-21-D3-56 in the Destination Address Field of the frame. It sends the frame to Switch 1, into Port 2.

- Switch 1 looks up the address E5-BB-47-21-D3-56 in its switching table. It sees that E5 is associated with Port 5, so it sends the frame out Port 5. This is a very simple process, using little processing power. Ethernet switches are inexpensive for the volume of traffic they carry.
- Port 5 on Switch 1 connects to Port 3 on Switch 2. Switch 1 sends the frame out Port 5 to Switch 2. Switch 2 now looks up the address E5-BB-47-21-D3-56 in its switching table. This address is associated with Port 7, so Switch 2 sends the frame out Port 7.
- The frame arrives at Switch 3 through Port 4. Switch 3 now looks up the address E5-BB-47-21-D3-56 in its own switching table. This time, the address is associated with Port 6. Switch 3 sends the frame out Port 6. This takes it to the destination Host E5-BB-47-21-D3-56.

Note that each switch only knows the information in its switching table.[11] More specifically, it only knows what port to use to send the frame back out. Switches do not know the entire data link between the source host and the destination host.

Test Your Understanding

16. a) Do switches know the entire data link between the source and destination host? b) What does a switch know?

Hierarchical Switch Organization Note that the switches in Figure 5-24 form a **hierarchy**, in which each switch has only one parent switch above it. In fact,

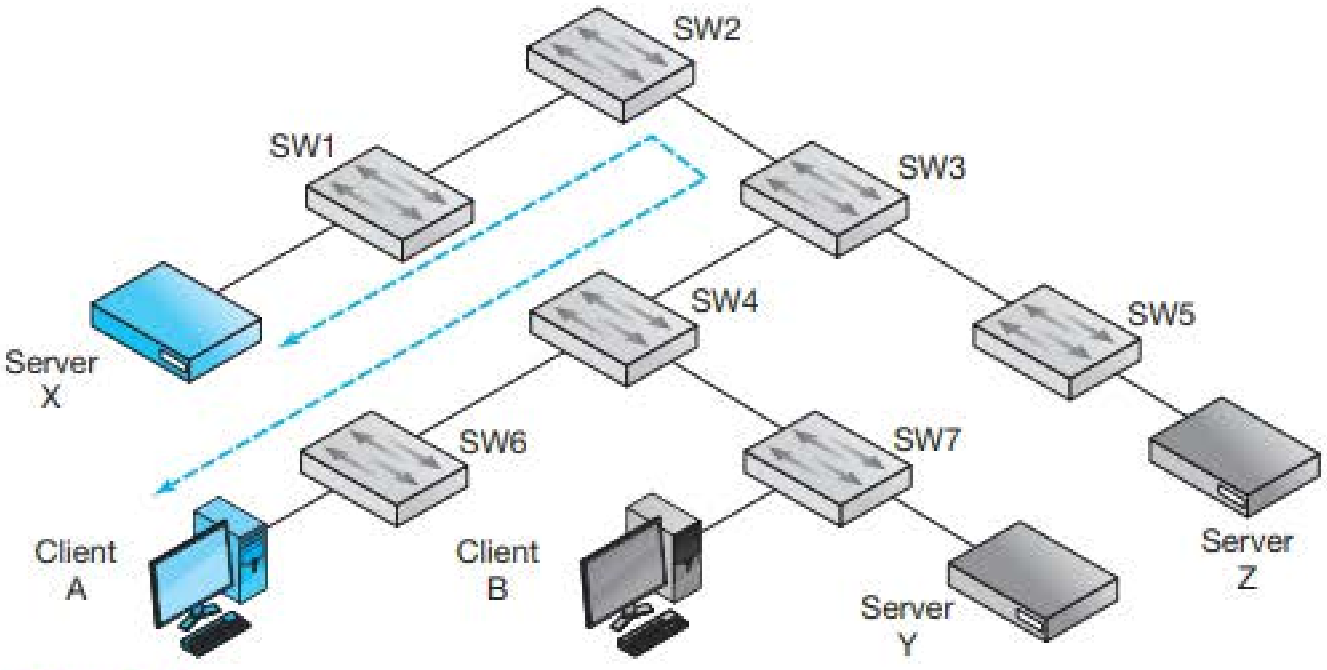

FIGURE 5-24 Hierarchical Ethernet Topology

[11] How does an Ethernet switch build its switching table? It notes the source address in every frame that arrives. If the source port address is not in its table, the switch adds it to the table.

the Ethernet standard *requires* a strict **hierarchical topology** (topology is the physical organization of switches and transmission links). Otherwise, loops would exist, and a single loop will cause the network to shut down. Figure 5-2 earlier showed a larger switched Ethernet LAN organized in a hierarchy.

Ethernet requires a hierarchical switch topology.

In a hierarchy, there is only a single possible path between any two hosts. To see this, look at the data link between Client Host A and Server X. The frame *must* pass through Switch 6, Switch 4, Switch 3, Switch 2, and Switch 1.

In a hierarchy, there is only a single possible path between any two hosts.

If there is only a single possible path between any two hosts, then there is only one possible port to send an arriving frame back out. Thanks to this simple rule, an Ethernet switch only needs a little computation power per frame handled—a simple table lookup. This makes Ethernet switches inexpensive per frame handled. During the 1970s and 1980s, there were competitors in the LAN switching market, but Ethernet's low cost, combined with adequate performance, made it dominant in the market.

Test Your Understanding

17. a) How are switches in an Ethernet LAN organized? b) Because of this organization, how many possible paths can there be between any two hosts? c) In Figure 5-2, what is the single possible path between Client PC 1 and Server X? Just give the letters of the switches. d) Between Client PC 1 and Server Y? Just give the letters of the switches. e) In Figure 5-24, list the switches on the path of frames from Client B to Server Z. f) Repeat for Client A to Client B. (Yes, clients do talk to one another.) g) From Server X to Server Z. (Yes, servers do talk to one another.)

18. a) What is the benefit of having a single possible path through an Ethernet network? b) Why has Ethernet become the dominant LAN technology?

Core Fields in the Ethernet Frame

In Chapter 2, we saw the most important fields in the Ethernet II frame. The box, "Secondary Fields in the Ethernet II Frame," discusses the remaining fields. We will add one piece of information to the discussion of the core fields here. We look at how to represent EUI-48 addresses in hexadecimal notation, which is how they are usually depicted.

Recall that Ethernet addresses are EUI-48 addresses (formerly MAC addresses). Although computers work with this raw 48-bit form, humans normally express these addresses in Base 16 **hexadecimal (hex) notation**. Figure 5-25 shows how to convert a 48-bit Ethernet address to hex notation. In the figure, the address begins with 10100011.

- First, divide the 48 bits into twelve 4-bit units, which computer scientists call nibbles. The first nibble is 1010. The second is 0011.

4 Bits	Decimal (Base 10)	Hexadecimal (Base 16)
0000	0	0 hex
0001	1	1 hex
0010	2	2 hex
0011	3	3 hex
0100	4	4 hex
0101	5	5 hex
0110	6	6 hex
0111	7	7 hex
1000	8	8 hex
1001	9	9 hex
1010	10	A hex
1011	11	B hex
1100	12	C hex
1101	13	D hex
1110	14	E hex
1111	15	F hex

Divide a 48-bit Ethernet address into 12 four-bit "nibbles." (1010, 0001, etc.)
Convert each group of 4 bits into a Hex symbol. (A, 1, etc.)
Combine two hex symbols into pairs and place a dash between pairs (A1-etc.)
For example, 10100001 becomes 1010 0001, which becomes A 1, which becomes A1 (followed by a dash)
The finished hex expression might be: A1-36-CD-7B-DF-01 hex

FIGURE 5-25 Hexadecimal Notation

- Second, convert each nibble into a hexadecimal symbol, using Figure 5-25. For example, 1010 is a A. The next nibble is 0001, which is 1.
- Third, write the hex symbols as six pairs of symbols separated by dashes. In this case, the first pair is A1. The entire address in "hex" might be A1-CC-66-0D-5E-BA.

To convert a hex address back to binary, change each symbol pair back to its 8-bit pattern. For example, if a hex pair is 2E, 2 is 0010, and E is 1110, so 2E is equivalent to the octet 00101110. Note that you must keep the two leading 0s in 0010 because the two symbols represent eight bits.[12]

Test Your Understanding

19. a) Are Ethernet EUI-48 addresses expressed in hex for humans, devices, or both? b) Which letters may appear in a hex EUI-48 address? c) What is 5 hex in binary? (Answer: 0101) d) What is 9 hex in binary? e) What is F hex in binary? (Answer: 1111) f) What is A hex in binary? g) What is binary 0011 in hex? (Answer: 3) h) What is binary 0000 in hex? i) What is binary 1111 in hex? j) What is A6 in binary? (Answer: 01101101) k) What is 6D hex in binary? l) Convert A1-B2-C3-44-5D-3C to binary. Leave a space between each octet. As a check, there must be 48 bits.

[12] Excel offers the bin2hex and hex2bin functions. Many advanced calculators can also do the calculation.

IN MORE DEPTH

Secondary Fields in The Ethernet Frame

Figure 5-26 lists three secondary fields in an Ethernet II frame.

Tag Fields (Optional) In Chapter 4, we saw that companies may give frames priority levels so that high-priority frames for latency-intolerant applications can go first. This was not in the original 802.3 standard. If a company wishes to use priority, it must configure its equipment to recognize two optional **tag fields**. These fields, when used, are inserted just before the length field. The first tag field, the Tag Protocol ID Field, merely indicates that the frame is tagged. The second gives the tagged information. Note again that these tag fields are optional. If priority and the other matters they handle are not used, there are no tag fields in the frame.

Three bits in the Tag Control Information Field are for priority level. With 3 bits, there can be eight priority levels. Another 12 bits are used for Ethernet VLAN numbers, which we will see in the Security section.

PAD Field In early versions of Ethernet, if the total length the Ethernet Data Field was less than 46 octets, it could cause problems for network operations. Consequently, if the Data Field is less than 46 octets long, a PAD Field is added to make the Data Field plus the PAD 46 octets long. For example, if the Data Field without a PAD is 40 octets long, a six-octet PAD Field is added. However, if the Data Field is 100 octets long, no PAD is added. The PAD Field, if needed, is placed before the Frame Check Sequence Field.

Test Your Understanding

20. a) What information do the two tag fields give? b) When is the PAD Field added?

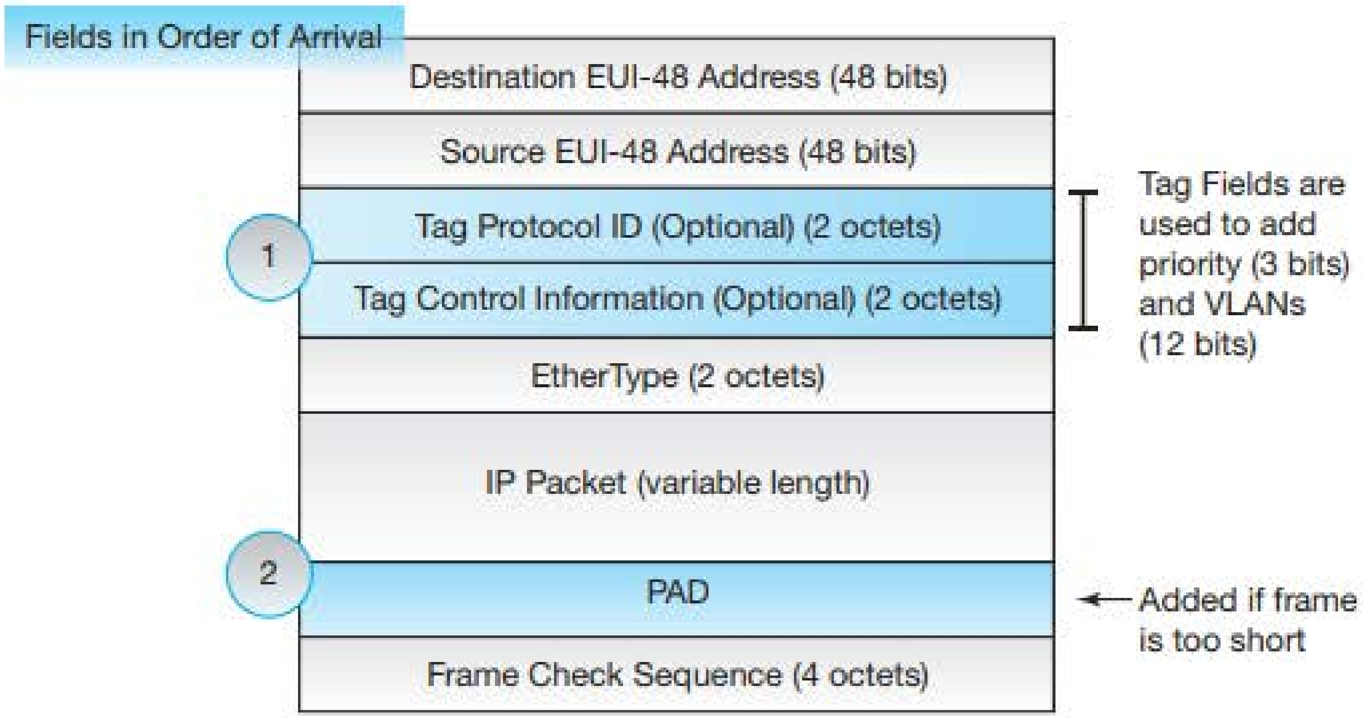
Fields in Order of Arrival
Destination EUI-48 Address (48 bits)
Source EUI-48 Address (48 bits)
1
Tag Protocol ID (Optional) (2 octets)
Tag Control Information (Optional) (2 octets)
Tag Fields are used to add priority (3 bits) and VLANs (12 bits)
EtherType (2 octets)
IP Packet (variable length)
2
PAD
Added if frame is too short
Frame Check Sequence (4 octets)

FIGURE 5-26 Secondary Fields in an Ethernet II Frame

MANAGEMENT

SNMP

In Chapter 3, we discussed the Simple Network Management Protocol (SNMP) for managing remote devices. SNMP was developed expressly for Ethernet and TCP/IP devices. As noted in Chapter 3, companies use network visualization programs to analyze data from the SNMP management information base and to send commands to devices to change how they operate. To be visible to the SNMP manager, an Ethernet switch must be a **manageable switch**, meaning that it must have an SNMP agent. It also needs the electronics to gather the data the SNMP manager asks for in Get commands and to make changes indicated in Set commands. Manageable switches are much more expensive than ordinary nonmanageable switches like the one you may have at home.

Test Your Understanding

21. a) What protocol do companies use to manage their Ethernet networks? b) What are manageable switches? c) Are all Ethernet switches manageable?

Reliability

We have seen that Ethernet switches must be arranged in a hierarchy. In a hierarchy, there can only be one possible data link between two hosts. We saw that this makes Ethernet switches inexpensive. However, if there is a break in a transmission link or switch, there is no way around it. In Figure 5-27, what if the link fails between Switch 1 and Switch 2? When Host A transmits a frame, it will stop at Switch 2 because it cannot travel on. (For the moment, ignore the backup link.) This does not mean that the entire network goes down. Host C can still reach Host B because the data link between them does not pass through the failed link.

Now look at the backup link that can connect Switch 2 and Switch 5. What if this backup link is plugged into the two switches? In that case, the entire network will be unable to work. Ethernet technology is very serious about strict hierarchies. (The first

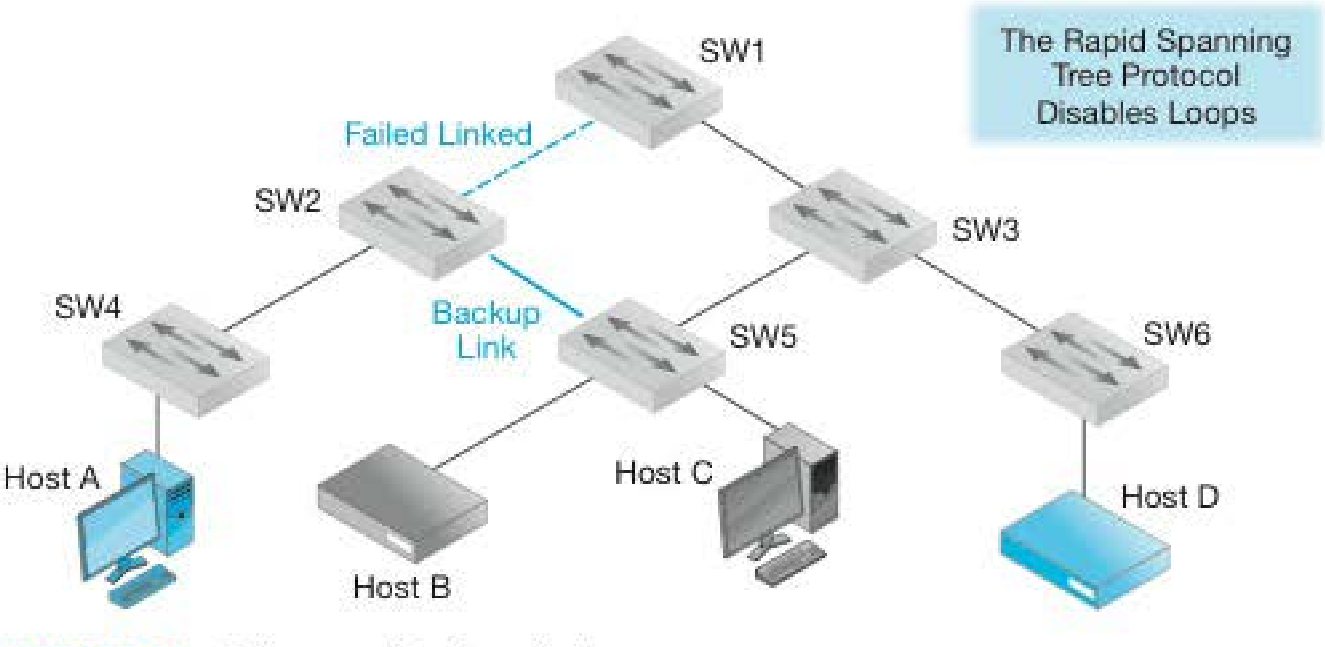

FIGURE 5-27 Failures and Backup Links

author sometimes does this as a class demonstration with cheap 4-port Ethernet switches to show what happens when a loop is introduced into a functioning network.)

However, if the switches are commercial grade, the problem only lasts a moment. The switches quickly realize that something is wrong. They begin sending supervisory frames to one another using the **Rapid Spanning Tree Protocol (RSTP)**. The switches then break the loop by closing the two ports of one of the physical links creating the loop. The network is hierarchical again. Transmission restarts.

Network engineers quickly saw that they could also use RSTP for an unintended purpose, to create backup links like the one shown in Figure 5-27. They could set it up so the backup link would be automatically disabled by RSTP but left in place. Then, if the link between Switch 1 and Switch 2 failed, the switches would engage in RSTP exchanges. They would open the two ports connecting the backup link. The network would be whole again, and all communication would continue. This seems simple. It is not. Creating backup links in a way that the network will reconfigure itself as the hierarchy the network manager wants turns out to be complicated. Beyond a handful of backup links, the effort begins to be prohibitive.

Test Your Understanding

22. a) What reliability problem does Ethernet have? b) How can some redundant backup links be installed without creating loops? c) Is this easy to do?

ETHERNET SECURITY

Ethernet Security in Perspective

Ethernet was designed to be simple, low in equipment price, and extremely low in management labor. This required the extensive use of trust. Devices can add themselves to basic Ethernet networks without proving their rights to join them. Corporations have tended to downplay Ethernet security because you have to get inside the corporate walls to exploit these weaknesses by plugging in physically to the LAN. However, if a computer within the network is compromised, the attacker is effectively inside the walls and can attack freely. These attacks on LANs are not widespread, but they are potentially dangerous. We look at four security threats and countermeasures the companies should assess.

Virtual LANs (VLANs) for Network Segregation

Even within a corporate site, a company does not want to give every employee access to every resource. Ethernet can enforce this segregation of network resources. It does this through **virtual LANs (VLANs)**, which are clusters of servers and hosts that are allowed to communicate with one another.[13]

[13] Although VLANs are now thought of as security tools, they were originally created to reduce broadcasting. In some situations, hosts or switches will broadcast frames to all other hosts on the LAN. To give an example, if the destination address in an arriving packet is not in a switch's switching table, the switch does not know what port to send it out. The switch broadcasts it out all ports. In small Ethernet LANs, this does not create problems. In very large Ethernet LANs, broadcasting is more frequent and can be a serious capacity hog.

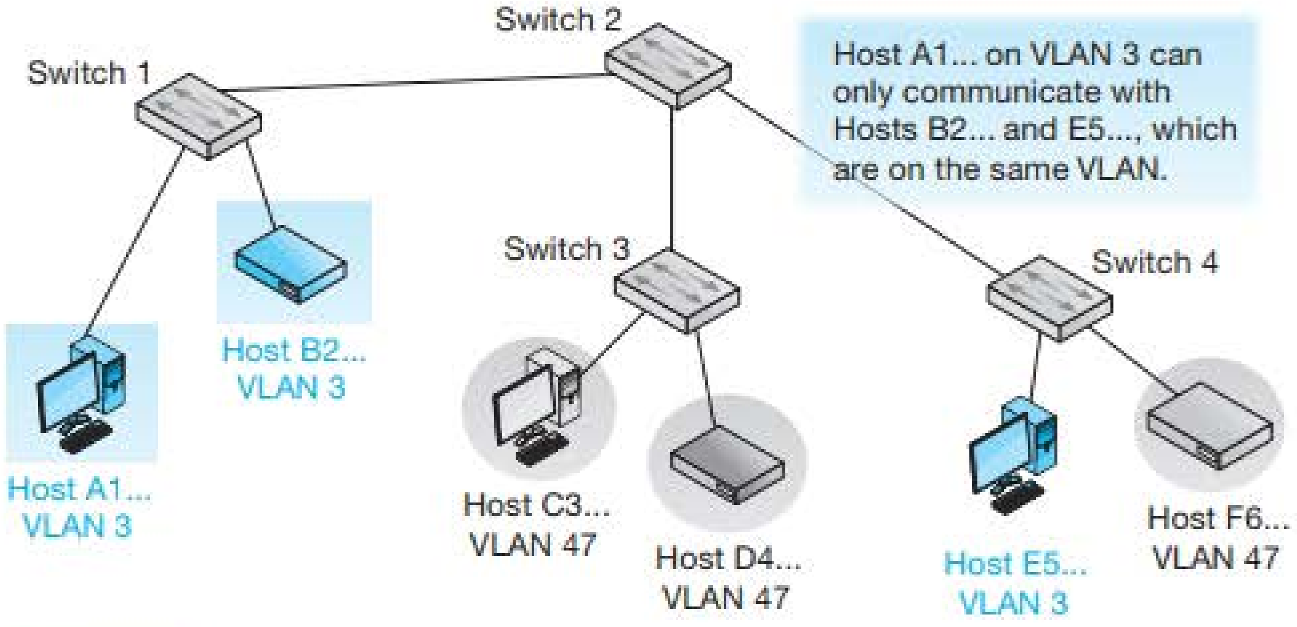

FIGURE 5-28 Virtual LANs (VLANs)

Figure 5-28 illustrates VLANs in Ethernet. In the figure, there are four switches and six hosts (two clients and four servers). These hosts are assigned to one of two virtual LANs: VLAN 3 or VLAN 47.

Host A1 is on VLAN 3. If it sends frames to Host B2 or Host E5, the switches will permit the transmission because Hosts A1, B2, and E5 are on the same VLAN. However, if Host A1 tries to send a frame to Host D4, the switches will not deliver the frame because Host D4 is on a different virtual LAN, VLAN 47.

If the user of Host A1 is in the marketing department, VLAN 3 might consist of the clients and servers in marketing. VLAN 47, in turn, might be the VLAN for the accounting department. Putting these departments on different VLANs ensures that people in the marketing department have no access to the accounting servers, which may hold sensitive information.

Test Your Understanding

23. a) What is the security benefit of Ethernet VLANs? b) In Figure 5-28, to which hosts can Host D4 send frames?

Initial User Authentication Through 802.1X

One way to reduce risks is to authenticate a user before he or she is authorized to use a switch port. In **802.1X Port-Based Network Access Control**, the switch initially permits frames to be exchanged only between the supplicant host and a central authentication server (see Figure 5-29). The authentication server asks the supplicant for specific credentials. The supplicant responds. If the server accepts the credentials and authenticates the host, it authorizes the switch to authorize access to the port. Otherwise, the port remains unauthorized and the supplicant is locked out of the network.

Test Your Understanding

24. a) What security threat is 802.1X designed to protect against? b) When 802.1X is being used, what happens if an attacker plugs his or her host into a switch?

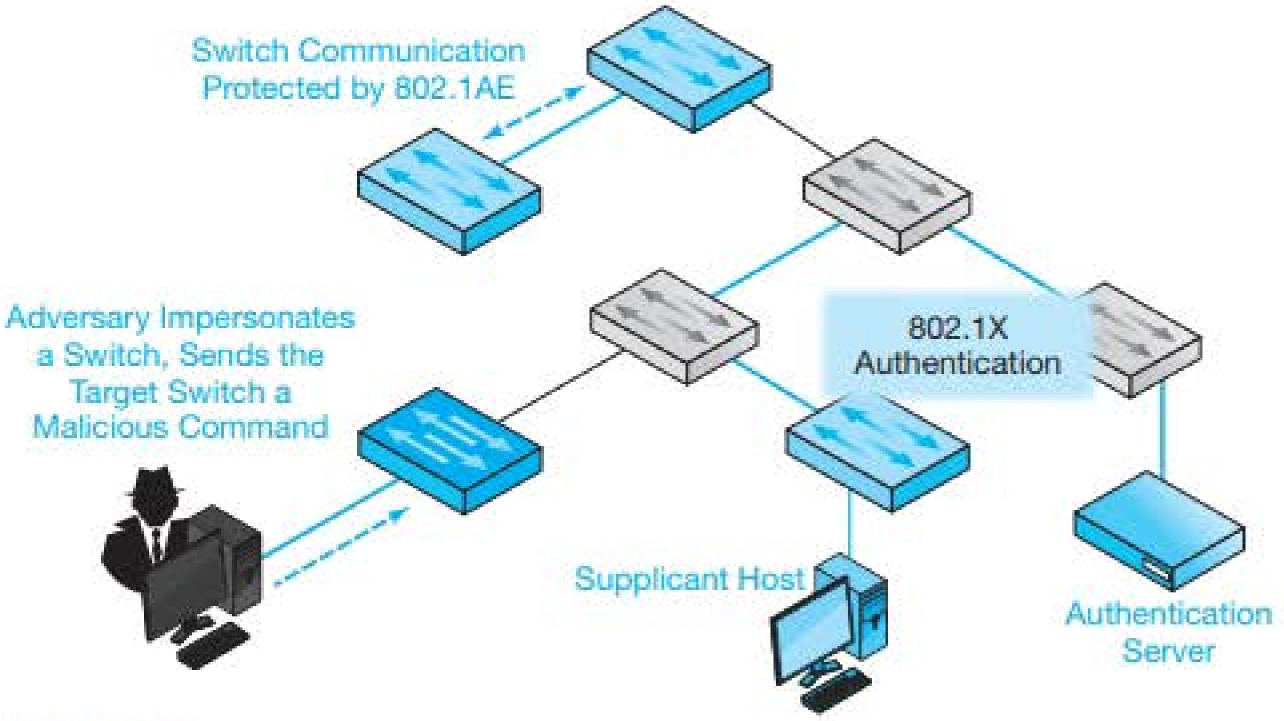
Switch Communication Protected by 802.1AE
Adversary Impersonates a Switch, Sends the Target Switch a Malicious Command
802.1X Authentication
Supplicant Host
Authentication Server

FIGURE 5-29 Initial User Authentication with 802.1X and Switch-to-Switch Security with 802.1AE

802.1AE Switch-to-Switch Protection

Of course, an authenticated host may still be a malicious user or an attacker who has taken over the host. As noted earlier, the host can imitate a switch and send management frames to real switches. It can tell them to shut down, direct all traffic through them so that they can read everything going through the network, and do many other things.

Attacks on switches are serious because they affect the overall operation of the Ethernet network. Normally, a switch will accept any management frame sent by another switch (or a host pretending to be a switch). The way to reduce trust is to require switches to authenticate themselves before another switch will listen to them. The 802.1AE standard shown in Figure 5-29 was created to do this. It also encrypts traffic between switches.

Test Your Understanding

25. What type of attack does 802.1AE protect against?

ARP Cache Poisoning

Another possibility is a **man-in-the-middle** attack[14] using **ARP cache poisoning**. Every host has an ARP cache that associates known IP destination addresses with their known EUI-48 address. It is easy for an attacker's host to send an **ARP update** message to other hosts it can reach via Ethernet. It tells them that the EUI-48 address of the router to which outgoing packets will be sent is actually the attacker's EUI-48 address. If hosts allow these unsolicited updates, as they often do, then every time they send packets believing that they are sending them to the router, they are actually sending them to the attacker's host. The attacker can read them and pass them on.

[14] This is the only common case in networking today in which gender-neutral terminology is not used.

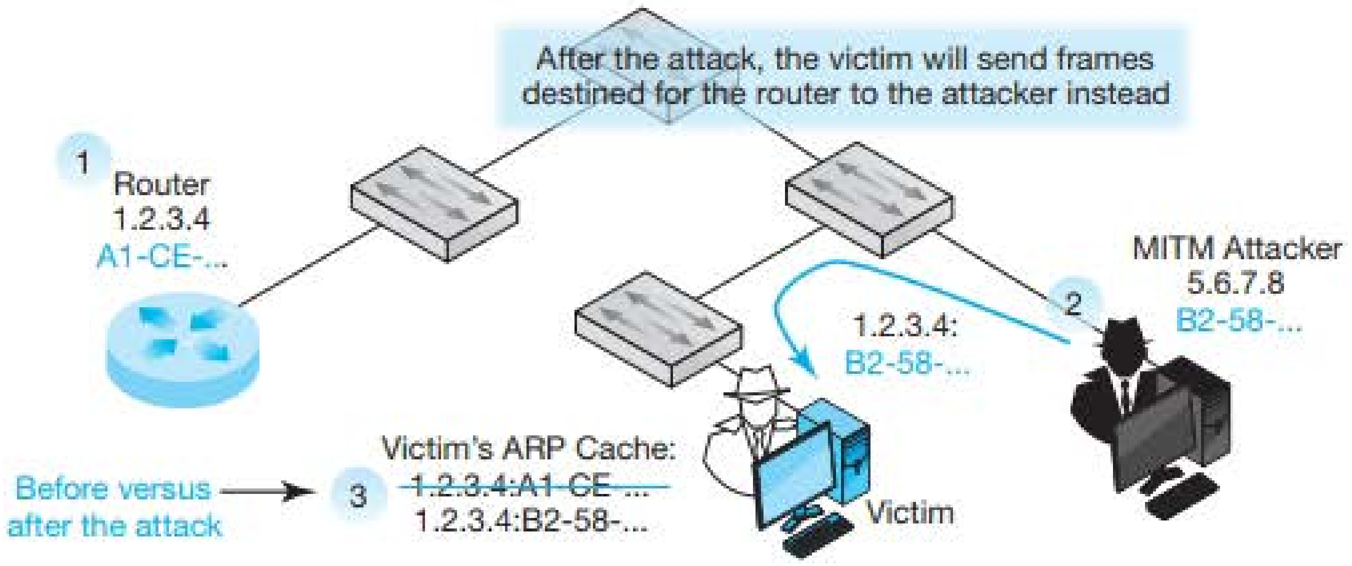

FIGURE 5-30 ARP Cache Poisoning (Study Figure)

Test Your Understanding

26. a) Before the attack, where does the ARP cache tell the victim to send a frame carrying a packet to the router? b) Where does it tell the victim to send such frames after the attack? c) What harm can the attacker do?

END-OF-CHAPTER QUESTIONS

Thought Questions

5-1. **a)** If both UTP and optical fiber can be used for a particular physical link, which should I choose? Why? b) If both OM3 fiber and OM4 fiber can be used for a particular physical link, which should I choose? Why? c) Some companies are now installing OM4 even when OM3 can do the job. Why do you think they do that?

5-2. What Work Group of the 802 LAN/MAN Standards Committee developed the 802.1X and 802.1AE standards? This working group, by the way, creates security and management standards for use in all other Working Groups.

5-3. In ARP cache poisoning, the attacker poisons the victim's ARP cache. This allows the attacker to read frames that the victim sends to the router. How can it read the frames that the victim receives from the router?

5-4. Can you think of a way to allow a client on one virtual LAN to communicate with a server on another VLAN? *Hint*: Switches typically have no way of doing it. *Hint 2*: It involves an access control list or some other type of device. The ACL permits specific client-server host pairs to communicate.

Design Questions

5-5. Design an Ethernet network to connect a single client PC to a single server. Both the client and the server will connect to their workgroup switches via UTP. The two devices are 900 meters apart. They need to communicate at 800 Mbps. Your design will specify the locations of any switches and the transmission link between the switches.

5-6. Add to your design in the previous question. Add another client next to the first client. Both connect to the same switch. This second client will also communicate with the server and will also need 800 Mbps in transmission speed. Again, your design will specify the locations of switches and the transmission link between the switches.

Troubleshooting Question

5-7. You connect two switches in a large Ethernet switch with 113 switches. You are using 4-pair UTP. Immediately after you make the connection, the network stops transmitting traffic. What do you think might have happened?

Perspective Questions

5-8. What was the most surprising thing you learned in this chapter?

5-9. What was the most difficult part of this chapter for you?

Chapter **5a**

Hands-On: Cutting and Connectorizing UTP[1]

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Cut, connectorize, and test 4-pair UTP cabling.
- Explain the difference between solid wire and stranded-wire UTP.
- Know when to use patch cables.

INTRODUCTION

Chapter 5 discussed UTP wiring in general. This chapter discusses how to cut and connectorize (add connectors to) solid UTP wiring.

SOLID AND STRANDED WIRING

Solid-Wire UTP versus Stranded-Wire UTP

The TIA/EIA-568 standard requires that long runs to wall jacks use **solid-wire UTP**, in which each of the eight wires really is a single solid wire.

However, patch cords running from the wall outlet to a NIC usually are **stranded-wire UTP**, in which each of the eight "wires" really is a bundle of thinner wire strands. So stranded-wire UTP has eight bundles of wires, each bundle in its own insulation and acting like a single wire.

[1] This material is based on the author's lab projects and on the lab project of Professor Harry Reif of James Madison University.

Solid-Wire UTP

- Each of the eight wires is a solid wire
- Low attenuation over long distances
- Easy to connectorize
- Inflexible and stiff—not good for runs to the desktop

Stranded-Wire UTP

- Each of the eight "wires" is itself several thin strands of wire within an insulation tube
- Flexible and durable—good for runs to the desktop
- Impossible to connectorize in the field (bought as patch cords)
- Higher attenuation than solid-wire UTP—Used only in short runs
  - From wall jack to desktop
  - Within a telecommunications closet (see Chapter 3)

FIGURE 5a-1 Solid-Wire and Stranded-Wire UTP (Study Figure)

Relative Advantages

Solid wire is needed in long cords because it has lower attenuation than stranded wire. In contrast, stranded-wire UTP cords are more flexible than solid-wire cords, making them ideal for patch cords—especially the one running to the desktop— because they can be bent more and still function. They are more durable than solid-wire UTP cords.

Adding Connectors

It is relatively easy to add RJ-45 connectors to solid-wire UTP cords. However, it is very difficult to add RJ-45 connectors to stranded-wire cords. Stranded-wire patch cords should be purchased from the factory precut to desired lengths and preconnectorized.

In addition, when purchasing equipment to connectorize solid-wire UTP, it is important to purchase crimpers designed for solid wire.

CUTTING THE CORD

Solid-wire UTP normally comes in a box or spool containing 50 meters or more of wire. The first step is to cut a length of UTP cord that matches your need. It is good to be a little generous with the length. This way, bad connectorization can be fixed by cutting off the connector and adding a new connector to the shortened cord. Also, UTP cords should never be subjected to pulls (strain), and adding a little extra length creates some slack.

STRIPPING THE CORD

Now the cord must be stripped at each end using a **stripping tool** such as the one shown in Figure 5a-2. The installer rotates the stripper once around the cord, scoring (cutting into) the cord jacket (but not cutting through it). The installer then pulls off the scored end of the cord, exposing about 5 cm (about 2 inches) of the wire pairs.

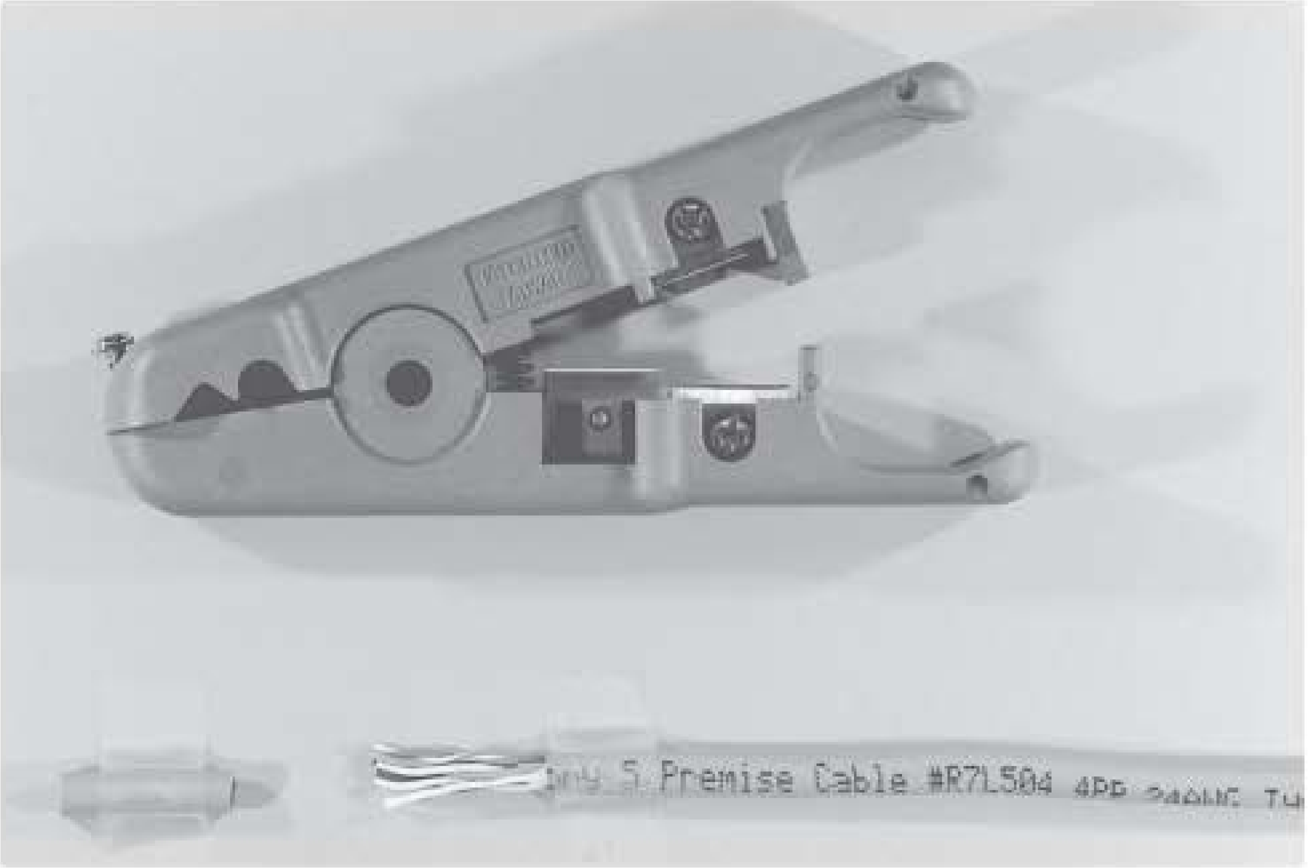

FIGURE 5a-2 Stripping Tool

It is critical not to score the cord too deeply, or the insulation around the individual wires may be cut. This creates short circuits. A really deep cut also will nick the wire, perhaps causing it to snap immediately or later.

WORKING WITH THE EXPOSED PAIRS

Pair Colors

The four pairs each have a color: orange, green, blue, or brown. One wire of the pair usually is a completely solid color. The other usually is white with stripes of the pair's color. For instance, the orange pair has an orange wire and a white wire with orange stripes.

Untwisting the Pairs

The wires of each pair are twisted around each other several times per inch. These must be untwisted after the end of the cord is stripped.

Ordering the Pairs

The wires now must be placed in their correct order, left to right. Figure 5a-3 shows the location of Pin 1 on the RJ-45 connector and on a wall jack or NIC.

Which color wire goes into which connector slot? The two standardized patterns are shown in Figure 5a-4. The T568B pattern is much more common in the United States.

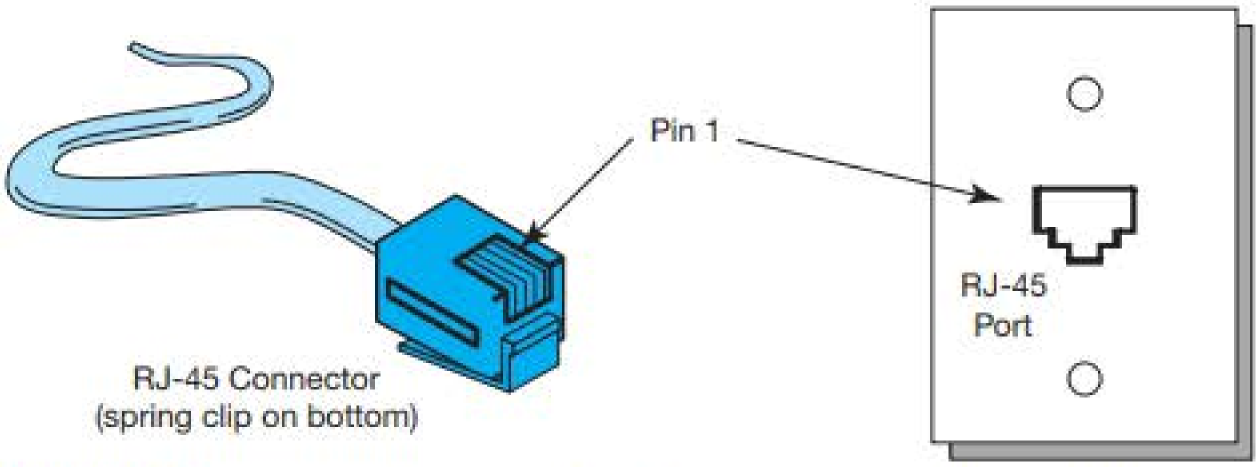

FIGURE 5a-3 Location of Pin 1 on an RJ-45 Connector and Wall Jack or NIC

The connectors at both ends of the cord use the same pattern. If the white-orange wire goes into Pin 1 of the connector on one end of the cord, it also goes into Pin 1 of the connector at the other end.

Cutting the Wires

The length of the exposed wires must be limited to 1.25 cm (0.5 inch) or slightly less. After the wires have been arranged in the correct order, a cutter should cut across the wires to make them this length. The cut should be made straight across so that all wires are of equal length. Otherwise, they will not all reach the end of the connector when they are inserted into it. Wires that do not reach the end will not make electrical contact.

ADDING THE CONNECTOR

Holding the Connector

The next step is to place the wires in the RJ-45 connector. In one hand, hold the connector, clip side down, with the opening in the back of the connector facing you.

Pin*	T568A	T568B
1	White-Green	White-Orange
2	Green	Orange
3	White-Orange	White-Green
4	Blue	Blue
5	White-Blue	White-Blue
6	Orange	Green
7	White-Brown	White-Brown
8	Brown	Brown

*Do not confuse T568A and T568B pin colors with the TIA/EIA-568 Standard.

FIGURE 5a-4 T568A and T568B Pin Colors

Sliding in the Wires

Now, slide the wires into the connector, making sure that they are in the correct order (white-orange on your left). There are grooves in the connector that will help. Be sure to push the wires all the way to the end or proper electrical contact will not be made with the pins at the end.

Before you crimp the connector, look down at the top of the connector, holding the tip away from you. The first wire on your left should be mostly white. So should every second wire. If they are not, you have inserted your wires incorrectly.[2]

Some Jacket Inside the Connector

If you have shortened your wires properly, there will be a little bit of jacket inside the RJ-45 connector.

CRIMPING

Pressing Down

Get a really good **crimping tool** (see Figure 5a-5). Place the connector with the wires in it into the crimp and push down firmly. Good crimping tools have ratchets to reduce the chance of your pushing down too tightly.

Making Electrical Contact

The front of the connector has eight pins running from the top almost to the bottom (spring clip side). When you **crimp** the connector, you force these eight pins through the insulation around each wire and into the wire itself. This seems like a crude electrical connection, and it is. However, it normally works very well. Your wires are now

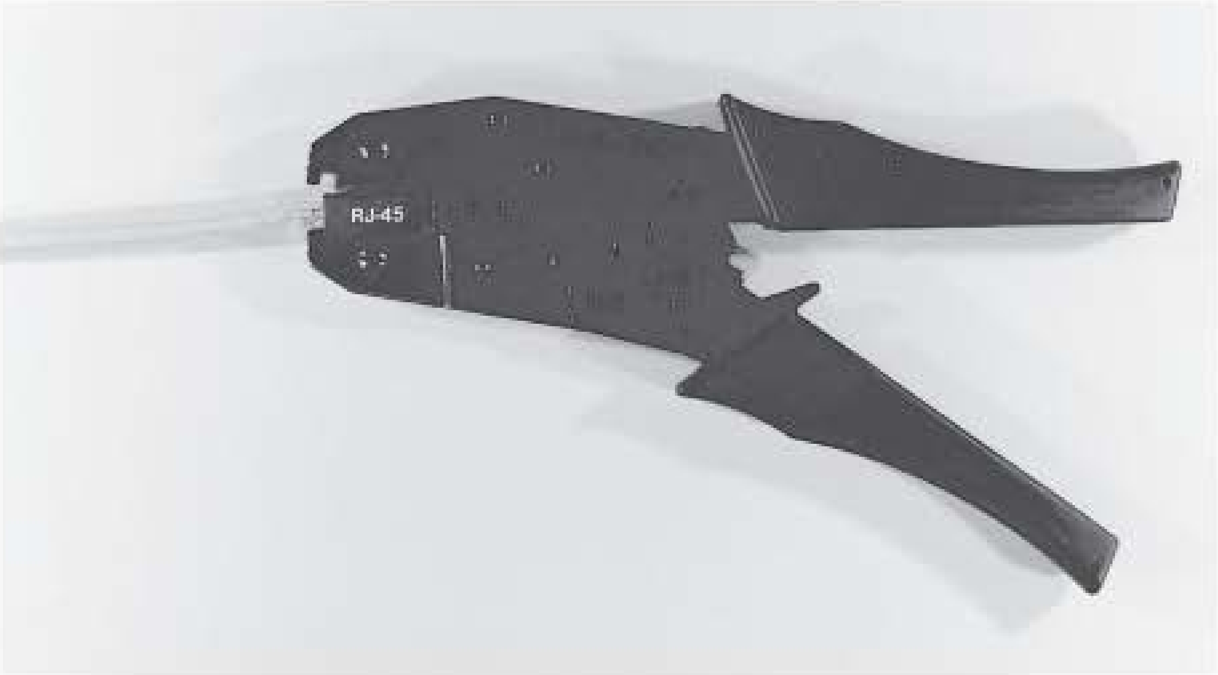

FIGURE 5a-5 Crimping Tool

[2] Thanks to Jason Okumura, who suggested this way of checking the wires.

connected to the connector 's pins. By the way, this is called an **insulation displacement connection (IDC)** because it cuts through the insulation.

Strain Relief

When you crimp, the crimper also forces a ridge in the back of the RJ-45 connector into the jacket of the cord. This provides **strain relief**, meaning that if someone pulls on the cord (a bad idea), they will be pulling only to the point where the jacket has the ridge forced into it. There will be no strain where the wires connect to the pins.

TESTING

Purchasing the best UTP cabling means nothing unless you install it properly. Wiring errors are common in the field, so you need to test every cord after you install it. Testing is inexpensive compared to troubleshooting subtle wiring problems later.

Testing with Continuity Testers

The simplest testers are **continuity testers**, which merely test whether the wires are arranged in correct order within the two RJ-45 connectors and are making good electrical contact with the connector. They cost only about $100.

Testing for Signal Quality

Better testers cost $500–$2,000 but are worth the extra money. In addition to testing for continuity problems, they send **test signals** through the cord to determine whether the cord meets TIA/EIA-568 signal-quality requirements. Many include **time domain reflectometry (TDR)**, which sends a signal and listens for echoes in order to measure the length of the UTP cord or to find if and where breaks exist in the cord.

Test Your Understanding

1. a) Explain the technical difference between solid-wire UTP and stranded-wire UTP. b) In what way is solid-wire UTP better? c) In what way is stranded-wire UTP better? d) Where would you use each? e) Which should only be connectorized at the factory?
2. If you have a wire run of 50 meters, should you cut the cord to 50 meters? Explain.
3. Why do you score the jacket of the cord with the stripping tool instead of cutting all the way through the jacket?
4. a) What are the colors of the four pairs? b) If you are following T568B, which wire goes into Pin 3? c) At the other end of the cord, would the same wire go into Pin 3?
5. After you arrange the wires in their correct order and cut them across, how much of the wires should be exposed from the jacket?
6. a) Describe RJ-45's insulation displacement approach. b) Describe its strain relief approach.
7. a) Should you test every cord in the field after installation? b) For what do inexpensive testers test? c) For what do expensive testers test?

Chapter **6**

Wireless LANs I

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Explain basic Wi-Fi 802.11 terminology and the role of access points.
- Explain basic radio signal propagation concepts, including frequencies, antennas, and wireless propagation problems. These are physical layer concepts.
- Explain the frequency spectrum, service bands, channels, bandwidth, licensed versus unlicensed service bands, and the type of spread spectrum transmission used in 802.11 Wi-Fi LANs. These are also physical layer concepts.
- Describe 802.11 Wi-Fi WLAN operation with access points and a switched Ethernet distribution system to link the access points. Distinguish between BSSs, ESSs, and SSIDs. Discuss communication between access points. These are data link layer concepts.
- If you read the box, "Media Access Control (MAC)," compare CSMA/CA+ACK and RTS/CTS for media access control. These are data link layer concepts.
- Compare and contrast the 802.11n and 802.11ac transmission standards. Discuss emerging trends in 802.11 operation, including channels with much wider bandwidth, MIMO, beamforming, and multiuser MIMO. These are physical layer concepts.
- If you read the box, "802.11/Wi-Fi Notes," Be able to know what happens when devices follow different Wi-Fi standards, explain how devices that follow new Wi-Fi standards get released in profile waves, and describe emerging 802.11 standards and what they will bring.

INTRODUCTION

OSI Standards

In Chapter 5, we looked at wired Ethernet networks. Technologies for networks require both physical and data link layer standards. Consequently, they use OSI standards. In this chapter and in Chapter 7, we look at wireless LANs. Like wired LANs, wireless LANs are single networks, which require physical and DLL standards. So they too use OSI standards.

Wireless LANs are governed by standards at the physical and data link layers. OSI dominates at this layer. This tells you that wireless LAN standards are OSI standards rather than IETF standards.

Test Your Understanding

1. a) At what layers do wireless LANs operate? b) Do wireless LAN standards governed by OSI or TCP/IP standards? Justify your answer.

802.11 = Wi-Fi

802.11 **Wireless LANs (WLANs)** use radio for physical layer transmission on the customer premises. In the last chapter, we saw that the 802.3 Working Group of the IEEE's 802 LAN/MAN Standards Committee creates Ethernet standards. Other working groups create other standards. The dominant WLAN standards today are the 802.11 standards, which are created by the **IEEE 802.11 Working Group**.

Wireless LANs (WLANs) use radio for physical layer transmission on the customer premises.

Wi-Fi It is common to call the 802.11 standards "**Wi-Fi**" standards. In fact, the terms have become almost interchangeable, and we use them that way too. However, as an IT professional, you should understand the technical difference

Wireless LANs

- Require standards at the physical and data link layer
- So OSI standards
- Standards created by the IEEE 802.11 Working Group

Wi-Fi

- Certification system managed by the Wi-Fi Alliance
- Wi-Fi now synonymous for 802.11

FIGURE 6-1 802.11 / Wi-Fi Wireless LAN (WLAN) Technology (Study Figure) between 802.11 and Wi-Fi. The term Wi-Fi stems from the **Wi-Fi Alliance**, which is an industry consortium of 802.11 product vendors. When the 802.11 Working Group creates standards, it often creates many options. The Wi-Fi Alliance creates subsets of 802.11 standards with selected options. The alliance conducts interoperability tests among products that claim to meet these "profiles." Only products that pass interoperability tests may display the Wi-Fi logo on their products. Products that do not pass are rarely sold, so when someone picks up a box containing an 802.11 product, they almost always see the Wi-Fi logo. This is why Wi-Fi has come to be more widely known than 802.11.

It is common to call the 802.11 standards "Wi-Fi" standards. In fact, the terms have become almost interchangeable, and we use them that way too.

Test Your Understanding

2. a) Distinguish between 802.3 standards and 802.11 standards. b) What is the actual difference between 802.11 and Wi-Fi? c) Do we use the two terms interchangeably?

Basic Access Point Operation

Figure 6-2 shows access point operation. First, it shows what happens when a wireless host sends a frame to another wireless host using the same access point (1 and 2). The source host transmits the frame to the access point. The access point then retransmits the frame to the destination host. We show this interaction as a pair of point-to-point transmissions.[1]

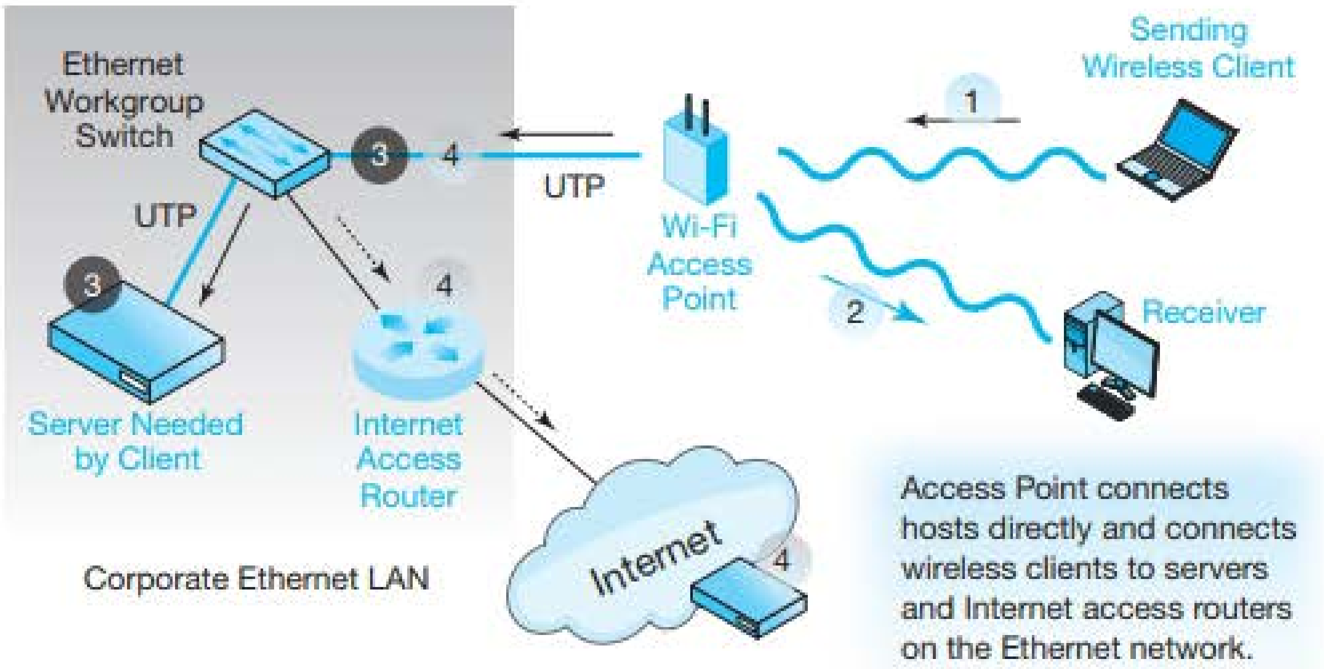

Access Point connects hosts directly and connects wireless clients to servers and Internet access routers on the Ethernet network.

FIGURE 6-2 Access Point Operation

[1] Actually, each device broadcasts its signal, so the signal spreads in all directions from the transmitter. The arrows indicate that only the receiver, to which the frame is addressed, pays attention to the frame (or at least should).

In most situations, however, the client needs to connect to a server that is elsewhere, on the corporate Ethernet LAN or outside the organization on the Internet. As Figure 6-2 shows, to reach corporate servers and to reach the site's Internet access router, the client needs to communicate over the corporate Ethernet LAN (3). Consequently, the access point connects via UTP to an Ethernet workgroup switch, which connects the wireless client to the rest of the site network.

Test Your Understanding

3. a) In a Wi-Fi LAN, do two wireless hosts usually send frames directly to one another? Explain. b) Why does the access point connect to the corporate Ethernet LAN?

RADIO SIGNAL PROPAGATION

Perfidious Radio

Chapter 5 discussed propagation effects in wired transmission media (UTP and optical fiber). Propagation effects in wired transmission can be well controlled by respecting cord distance limits and taking other installation precautions. This is possible because wired propagation is predictable. If you input a signal, you can estimate precisely what it will be at the other end of a cord. A wired network is like a faithful, obedient dog.

Propagation effects in wired transmission can be well controlled by respecting cord distance limits and taking other installation precautions.

In contrast, radio propagation is very unreliable. Radio signals bounce off obstacles, fail to pass through walls and filing cabinets, and have other problems we look at in this section. Consequently, Wi-Fi networks, which use radio to deliver signals, are more complex to implement than wired networks. They do not have a few simple installation guidelines that can reduce propagation effects to nonissues. Therefore, we will spend more time on wireless propagation effects than we did on wired propagation effects. We are dealing with cats.

Propagation effects in wireless networks are complex and difficult to solve.

Test Your Understanding

4. a) In 802.3 Ethernet networks, can simple installation rules usually reduce propagation effects to nonissues? b) In 802.11 Wi-Fi networks, can simple installation rules usually reduce propagation effects to nonissues?

Frequencies

Radios for data transmission are called **transceivers** because they both transmit and receive. When transceivers send, their wireless signals propagate as waves, as we saw in Chapter 5. Figure 6-3 again notes that waves have amplitude and wavelength.

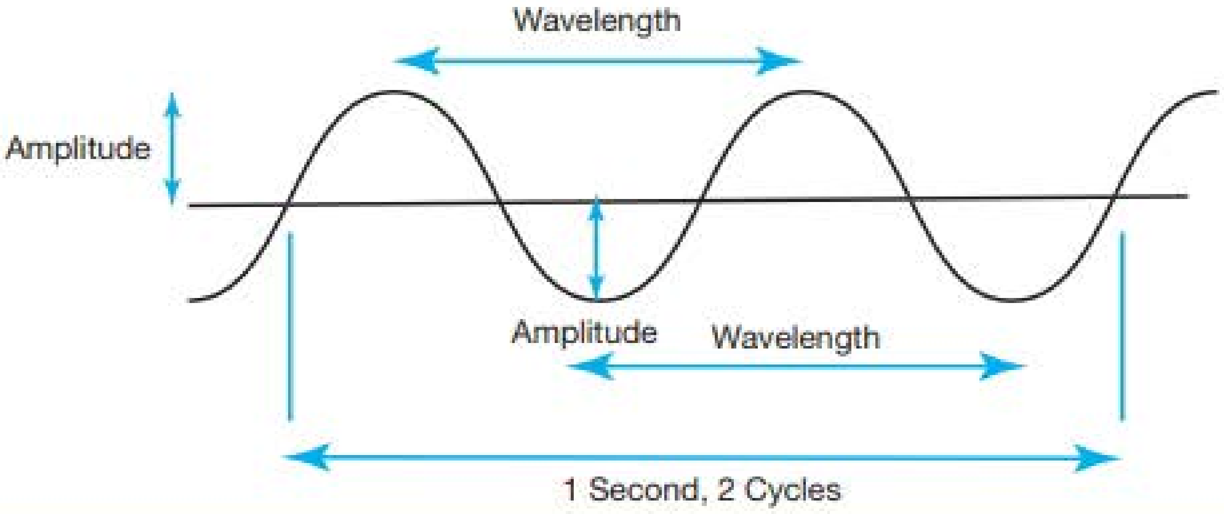

Wavelength is the physical distance between comparable points on adjacent cycles.
Optical fiber transmission is described in terms of wavelength.

Frequency is the number of cycles per second.
In this case, there are two cycles in 1 second, so the frequency is two hertz (2 Hz).
Radio transmission is measured in terms of frequency.

Amplitude is the power of the wave.

FIGURE 6-3 Electromagnetic Wave

Optical fiber waves are described in terms of wavelength, but radio waves are described in terms of another wave characteristic, frequency.

Frequency is used to describe the radio waves used in WLANs.

In waves, **frequency** is the number of complete cycles per second. One cycle per second is one **hertz (Hz)**. Metric designations are used to describe frequencies. In the metric system, frequencies increase by a factor of 1,000 rather than 1,024. The most common radio frequencies for wireless transceivers range between about 500 megahertz (MHz) and 10 gigahertz (GHz).

Test Your Understanding

5. a) What is a transceiver? b) Is wireless radio transmission usually described in terms of wavelength or frequency? c) What is a hertz? d) At what range of frequencies do most wireless systems operate?

Antennas

A transceiver uses an **antenna** to transmit its signal. Figure 6-4 shows that there are two types of radio antennas: omnidirectional antennas and dish antennas.

- **Omnidirectional antennas** transmit signals equally strongly in all directions and receive incoming signals equally well from all directions. Consequently, the antenna does not need to point in the direction of the receiver. However, because the signal spreads in all three dimensions, only a small fraction of the energy transmitted by an omnidirectional antenna reaches the receiver. Omnidirectional

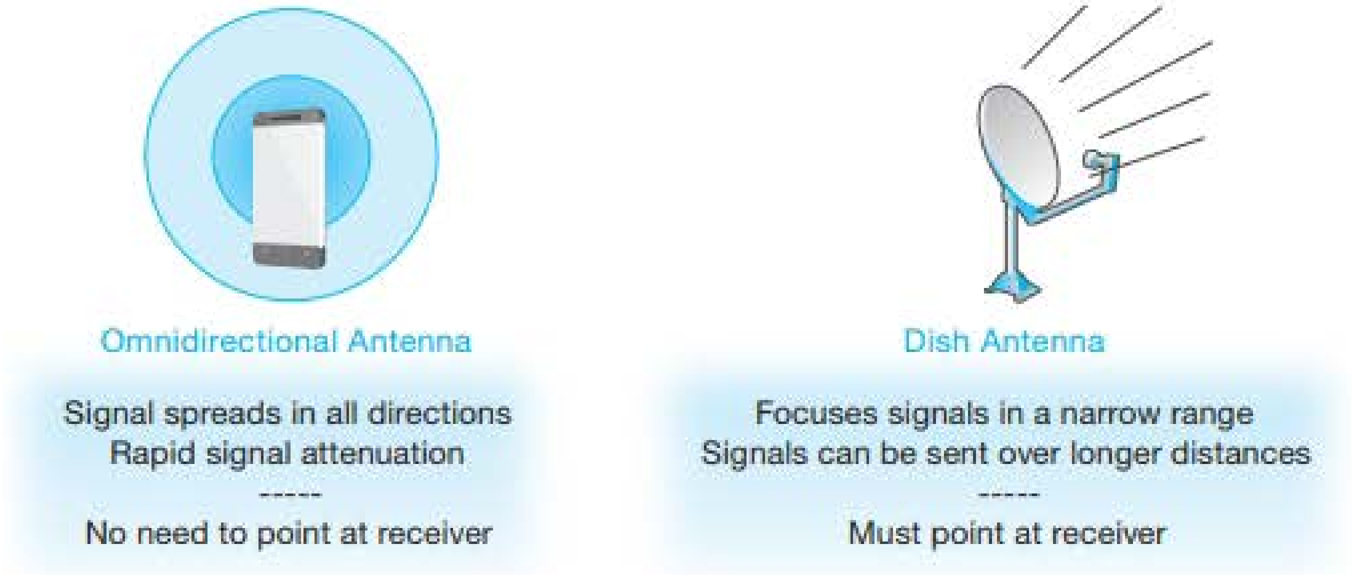
Omnidirectional Antenna
Signal spreads in all directions
Rapid signal attenuation

No need to point at receiver
Dish Antenna
Focuses signals in a narrow range
Signals can be sent over longer distances

Must point at receiver

FIGURE 6-4 Omnidirectional and Dish Antennas

antennas are used for short distances, such as those found in a wireless LAN or a cellular telephone network.

- **Dish antennas**, in contrast, concentrate signals in a particular direction, which allows signals to travel farther for the same transmission power. (A dish antenna is like the reflector in a flashlight.) It also allows them to receive weaker incoming signals from that direction. Dish antennas are used for longer distances because of their focusing ability, although users need to know the direction of the other radio. In addition, dish antennas are bulky. (Imagine if you had to carry a dish with you whenever you carried your cellular phone. You would not even know where to point the dish!)

Test Your Understanding

6. a) Distinguish between omnidirectional and dish antennas in terms of operation. b) Under what circumstances would you use an omnidirectional antenna? c) Under what circumstances would you use a dish antenna? d) What type of antenna normally is used in WLANs? Why?

Wireless Propagation Problems

We have already noted that, although wireless communication gives mobility, it is not very predictable, and there often are serious propagation problems. Figure 6-5 illustrates five common wireless propagation problems.

Inverse Square Law Attenuation Compared to signals sent through wires and optical fiber, radio signals attenuate very rapidly. When a signal spreads out from any kind of antenna, its strength is spread over the area of a sphere. (In omnidirectional antennas, power is spread equally over the sphere, whereas in dish antennas, power is concentrated primarily in one direction on the sphere.)

The area of a sphere is proportional to the square of its radius, so signal strength in any direction weakens by an **inverse square law** rule. If distance is doubled, signal strength falls to a quarter of its original value (1/2 squared). For example, if a signal is 100 watts at 10 meters, it will only be 25 W at 20 meters. If the distance is increased 10-fold, then signal strength will be only 1/100 its original value (1/10 squared), 1 watt.

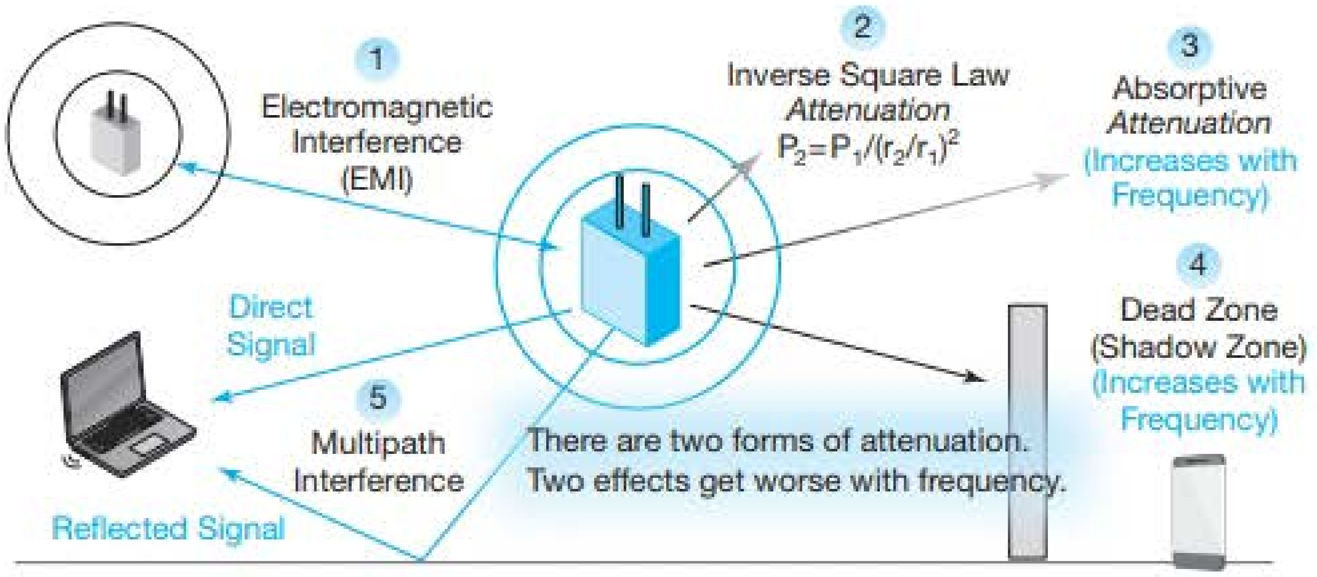
1
Electromagnetic Interference (EMI)
2
Inverse Square Law Attenuation
$P_2=P_1/(r_2/r_1)^2$
3
Absorptive Attenuation
(Increases with Frequency)
4
Dead Zone (Shadow Zone)
(Increases with Frequency)
Direct Signal
5
Multipath Interference
Reflected Signal
There are two forms of attenuation.
Two effects get worse with frequency.

FIGURE 6-5 Wireless Propagation Problems

Inverse square law attenuation is very rapid attenuation—far more rapid than attenuation in 4-Pair UTP and optical fiber.

Absorptive Attenuation As a radio signal travels, it is partially absorbed by the air molecules, plants, and other things it passes through. This **absorptive attenuation** is especially bad because water is an especially good absorber of radio signals. Rain and moisture in plants can reduce power substantially.

Absorptive attenuation can be confusing because we have already seen inverse square law attenuation. Yes, wireless propagation suffers from two forms of attenuation.

Distance Ratio	Distance Ratio Squared	Signal Strength Compared to Original	Initial Power (P2) (watts)	Final Power (P2) (watts)
1	1	100.0%	100	100
2	4	25.0%	100	25
3	9	11.1%	100	11.1
4	16	6.3%	100	6.3
5	25	4.0%	100	4.0
6	36	2.8%	100	2.8
7	49	2.0%	100	2.0
8	64	1.6%	100	1.6
9	81	1.2%	100	1.2
10	100	1.0%	100	1.0

Note: if the original distance is 10 meters and the final distance is 30 meters, the distance ratio will be 3. The signal strength ratio will be 11.15%. If the original power at 10 meters is 100 watts, the signal at 30 meters will be 11.1% of 100 W or 11.1 W.

FIGURE 6-6 Inverse Square Law Attenuation (Study Figure)

Inverse square law attenuation is due to the signal spreading out as a sphere and so becoming weaker at each point on the sphere as the sphere expands. Absorptive attenuation is signal loss through energy absorption.

Wireless transmission suffers from two forms of attenuation—inverse square law attenuation and absorptive attenuation.

Dead Zones To some extent, radio signals can go through and bend around objects. However, if there is a dense object (e.g., a thick wall) blocking the direct path between the sender and the receiver, the receiver may be in a **dead zone**, also called a shadow zone. In these zones, the receiver cannot get the signal. If you have a mobile phone and often try to use it inside buildings, you may be familiar with this problem.

Multipath Interference In addition, radio waves tend to bounce off walls, floors, ceilings, and other objects. As Figure 6-7 shows, this may mean that a receiver will receive two or more signals—a direct signal and one or more reflected signals. The direct and reflected signals will travel different distances and so may be out of phase when they reach the receiver. For example, one may be at its highest amplitude and the other at its lowest, giving an average of zero. If their amplitudes are the same, they will completely cancel out. In a real situation, multiple signals traveling different paths will interfere, so we call this type of interference **multipath interference**.

Multipath interference may cause the signal to range from strong to nonexistent within a few centimeters. If the difference in time between the direct and reflected signal is large, some reflected signals may even interfere with the next direct signal. We will see later that this is controlled by spread spectrum transmission, which spreads the signal over a wide range of frequencies so that multipath interference effects average out to zero. Multipath interference is the most serious propagation problem at WLAN frequencies.

Multipath interference is the most serious propagation problem at WLAN frequencies.

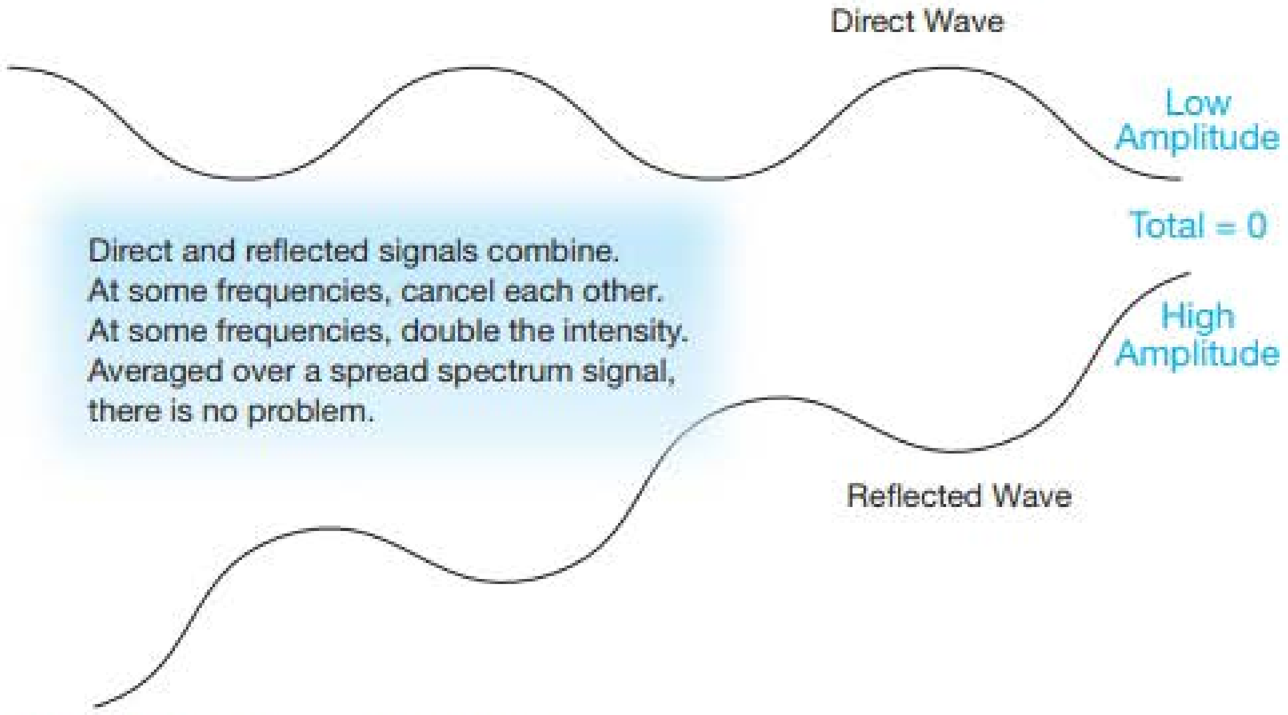

Direct and reflected signals combine.
At some frequencies, cancel each other.
At some frequencies, double the intensity.
Averaged over a spread spectrum signal, there is no problem.

FIGURE 6-7 Multipath Interference

Electromagnetic Interference (EMI) A final common propagation problem in wireless communication is **electromagnetic interference (EMI)**. Many devices produce electromagnetic radiation at frequencies used in wireless data communications. Among these devices are cordless telephones, microwaves, and nearby access points. We will see later in this chapter that placing access points so that they give good coverage without creating excessive mutual interference is an important but difficult task in WLAN management.

Frequency-Dependent Propagation Problems To complicate matters, two wireless propagation problems intensify as frequency increases.

- First, higher-frequency waves suffer more rapidly from absorptive attenuation than lower-frequency waves because they are absorbed more rapidly by moisture in the air. Consequently, as we will see in this chapter, WLAN signals around 5 GHz attenuate more rapidly than signals around 2.4 GHz.
- Second, dead zones become worse with increasing frequency. Radio waves become less able to bend around objects in their paths.

Test Your Understanding

7. a) If you quadruple propagation distance, how much will signal intensity change at the receiver? (Answer: 1/16) b) If you increase propagation distance by a factor of 100, how much will signal intensity change at the receiver? c) If the signal strength from an omnidirectional radio source is 8 mW at 30 meters, how strong will it be at 150 meters, ignoring absorptive attenuation? Show your work. (Answer: 0.32 mW) d) What will it be at 200 meters? e) If the signal strength from an omnidirectional radio source is 20 mW at 10 meters, how strong will it be at 70 meters, ignoring absorptive attenuation? Show your work.
8. a) Contrast inverse square law attenuation and absorptive attenuation. b) What causes dead zones? c) What is the most serious propagation problem in WLANs? d) How is it controlled? e) What two propagation problems become worse as frequency increases?

SERVICE BANDS AND BANDWIDTH

Service Bands

The Frequency Spectrum The **frequency spectrum** is the range of all possible frequencies from zero hertz to infinity, as Figure 6-8 shows.

Service Bands Regulators divide the frequency spectrum into contiguous spectrum ranges called **service bands,** which are dedicated to specific services. For instance, in the United States, the AM radio service band lies between 535 kHz and 1,705 kHz. The FM radio service band, in turn, lies between 87.5 MHz and 108.0 MHz. Wi-Fi uses the 2.4 GHz service band that we will see later in this chapter; this band extends from 2.4 GHz to 2.4835 GHz. Wi-Fi also uses the 5 GHz service band, which ranges from 5.25 GHz to 5.725 GHz (with some gaps in between that are used

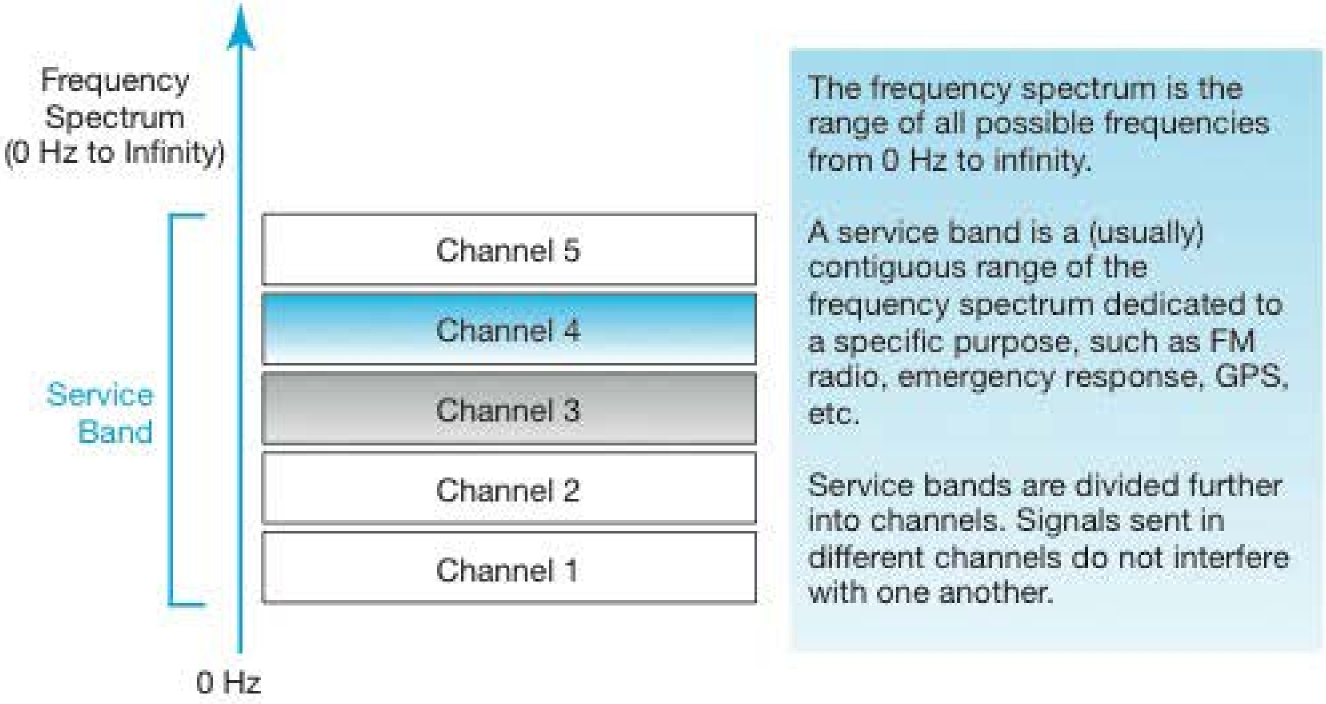

The frequency spectrum is the range of all possible frequencies from 0 Hz to infinity.

A service band is a (usually) contiguous range of the frequency spectrum dedicated to a specific purpose, such as FM radio, emergency response, GPS, etc.

Service bands are divided further into channels. Signals sent in different channels do not interfere with one another.

FIGURE 6-8 The Frequency Spectrum, Service Bands, and Channels

for other services). There are hundreds of other service bands for police and fire departments, amateur radio operators, communication satellites, and many other purposes.

Channels Service bands are subdivided further into smaller frequency ranges called **channels**. A different signal can be sent in each channel because signals in different channels do not interfere with one another. This is why you can receive different television channels successfully. In FM radio, channels are 200 kHz wide. So the first channel extends from 87.5 MHz to 88.5 MHz.

Test Your Understanding

9. a) Distinguish among the frequency spectrum, service bands, and channels. b) In radio, how can you send multiple signals without the signals interfering with one another? c) How many channels are there in the FM band? (You can compute this from information in the text.) d) Are the set of frequencies used for police communication in a city channels or a service band? Explain. e) An FM radio station is called Moldy Oldies 101.1. Is this a channel or a service band? f) Wi-Fi operates in the 2.4 GHz ____and the 5 GHz ____.

Signal and Channel Bandwidth

Figure 6-3 showed a wave operating at a single frequency. In reality, Figure 6-9 shows that real signals do not travel at a single frequency. Rather, real signals spread over a range of frequencies. This range is called the signal's bandwidth. **Signal bandwidth** is measured by subtracting the lowest frequency from the highest frequency.

A channel also has a bandwidth. For instance, if the lowest frequency of an FM channel is 89.0 MHz and the highest frequency is 89.2 MHz, then the **channel bandwidth** is 0.2 MHz (200 kHz). AM radio channels are 10 kHz wide, FM channels are 200 kHz wide, and television channels are 6 MHz wide.

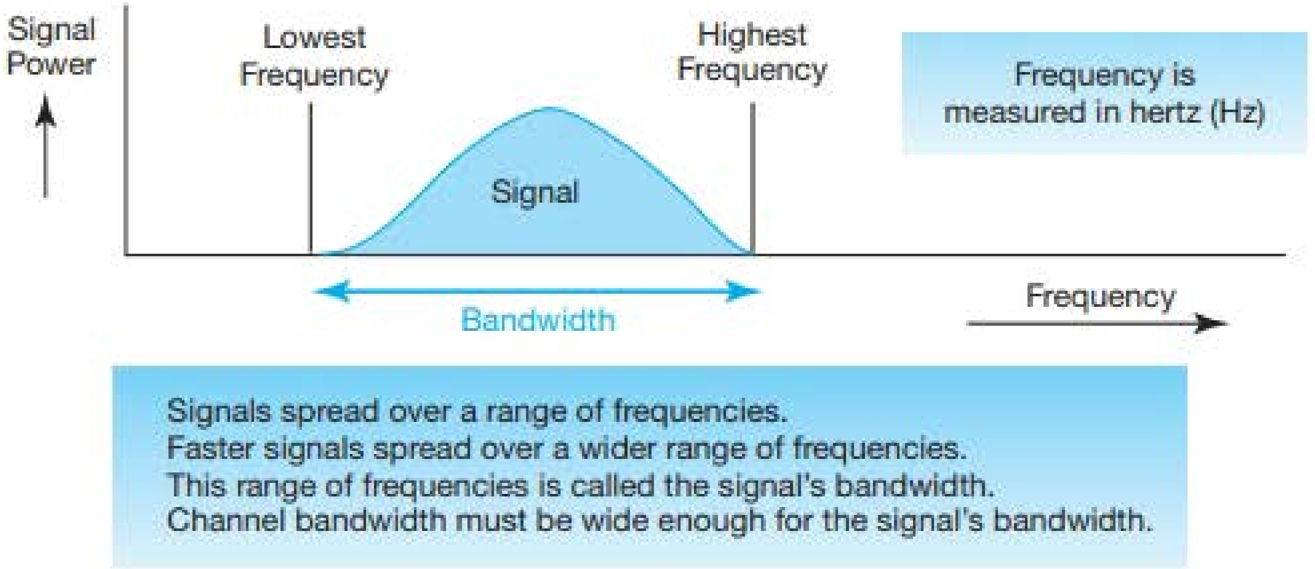

Signals spread over a range of frequencies.
Faster signals spread over a wider range of frequencies.
This range of frequencies is called the signal's bandwidth.
Channel bandwidth must be wide enough for the signal's bandwidth.

FIGURE 6-9 Signal Bandwidth

How wide must the channel bandwidth be? The channel bandwidth must be wide enough for a signal's bandwidth. Claude Shannon discovered a remarkable thing about signal transmission. A signal carrying X bits per second only needs half the bandwidth of a signal carrying 2X bits per second.[2] Looked at the other way, if you want to transmit twice as many bits per second, you need to double your bandwidth. More generally, if you want to be able to transmit N times as fast, you need N times as much channel bandwidth. High bandwidth brings high radio transmission speed.

To transmit N times as fast, you need N times as much channel bandwidth.

Radio channels with large bandwidths are called **broadband channels**. They can carry data very quickly. Although the term *broadband* technically refers only to the width of a channel in radio, **broadband** has come to mean "fast," whether or not radio is used.

Transmission systems that are very fast are usually called broadband systems even when they do not use radio channels.

Test Your Understanding

10. a) Does a signal travel at a single frequency, or does it spread over a range of frequencies? b) If the lowest frequency in a channel is 1.22 MHz and the highest frequency is 1.25 MHz, what is the channel bandwidth? c) If you want to transmit seven times as fast, how much wider must the channel be? d) Why is large channel bandwidth desirable? e) What do we call a system whose channels are wide? f) What other types of system do we call broadband?

[2] Speaking more precisely, Shannon also found that the signal-to-noise ratio (the ratio of signal power to noise power) also affects propagation speed. However, engineers find it far easier to increase speed by increasing bandwidth than by increasing the signal-to-noise ratio. Increasing signal power is usually limited by regulations, and reducing noise power is technically very difficult.

Signal Bandwidth

Figure 6-3 shows a wave operating at a single frequency.

However, most signals are spread over a range of frequencies (see Figure 6-9).

The range between the highest and lowest frequencies is the signal's bandwidth.

As transmission speed increases, the signal bandwidth increases.

Channel Bandwidth

Channel bandwidth is the highest frequency in a channel minus the lowest frequency.

An 87.5 MHz to 88.1 MHz channel has a bandwidth of 0.2 MHz (200 kHz).

Channel Bandwidth and Propagation Speeds

The maximum possible transmission speed increases with bandwidth.

Doubling the bandwidth doubles the maximum possible transmission speed.

Multiplying the bandwidth by X multiplies the maximum possible speed by X.

Higher-speed signals need wider channel bandwidths.

Channel bandwidth must be sufficient for the signal's bandwidth.

Broadband Channels

Broadband means wide channel bandwidth and therefore high speed.

Today, "broadband" has come to mean "fast," whether or not radio transmission in channels is used.

FIGURE 6-10 Channel Bandwidth and Transmission Speed (Study Figure)

Licensed and Unlicensed Service Bands

If two nearby transceivers send at the same frequency, their signals will interfere with each other. To prevent chaos, governments regulate radio transmission. The International Telecommunications Union, which is a branch of the United Nations, creates worldwide rules that define service bands and specify how individual radio service bands are to be used. Individual countries enforce these rules but are given discretion over how to implement controls.

Licensed Service Bands In **licensed service bands**, transceivers must have a government license to operate. They also need a license change if they move. Commercial television bands are licensed bands, as are AM and FM radio bands. Government agencies control who may have licenses in these bands. By doing so, the government limits interference to an acceptable level. In some licensed service bands, the rules allow mobile hosts to move about but central transceivers are regulated. This is the case for mobile telephones.

Unlicensed Service Bands However, for companies that have wireless access points and mobile computers, even the requirement to license central antennas (in this situation, access points) is an impossible burden. Consequently, the International Telecommunications Union created a few **unlicensed service bands**. In these bands, a company can add or drop access points any time it chooses. It can also have as many wireless hosts as it wishes. All 802.11 Wi-Fi networks operate in these unlicensed radio bands.

Licensed Service Bands

If two nearby radio hosts transmit in the same channel, their signals will interfere.

Most service bands are licensed bands, in which hosts need a license to transmit.

The government limits licenses to reduce interference.

Television bands, AM service bands, etc. are licensed.

In cellular telephone bands, which are licensed, only the central antennas are licensed, not the mobile phones.

Unlicensed Service Bands

Some bands are set aside as unlicensed bands.

Hosts do not need to be licensed to be turned on or moved.

802.11 Wi-Fi operates in unlicensed service bands.

This allows access points and hosts to be moved freely.

However, there is no legal recourse against interference from other nearby users.

Your only recourse is to negotiate.

At the same time, you may not cause unreasonable interference by transmitting at illegally high power.

FIGURE 6-11 Licensed and Unlicensed Radio Service Bands (Study Figure)

The downside of unlicensed service bands is that companies must tolerate interference from others. If your neighbor sets up a wireless LAN next door to yours, you have no recourse but to negotiate with him or her over such matters as which channels each of you will use. At the same time, the law prohibits unreasonable interference by using illegally high transmission power.

Test Your Understanding

11. a) Do WLANs today use licensed or unlicensed service bands? b) What is the advantage of using unlicensed service bands? c) What is the downside?

Channel Use and Co-Channel Interference

Figure 6-12 illustrates two important points about how Wi-Fi uses its channels. The first is that an access point normally uses a single channel (although some can operate on more than one). Access Point A is transmitting on "Channel 1," and so is adjacent Access Point B. They will interfere. This is called **co-channel interference** because they are using the same channel.

What about Access Point A and Access Point D? They are adjacent, but they are operating on different channels (1 and 6). Therefore, they will not interfere with one another.

When co-channel interference occurs, it does not stop transmissions, but it does slow them down. One hotel decided to be "consistent" and put all access points on the same channel. Service was terrible.

To reduce co-channel interference, network administrators try to set adjacent access points on different channels. However, there are a limited number of channels

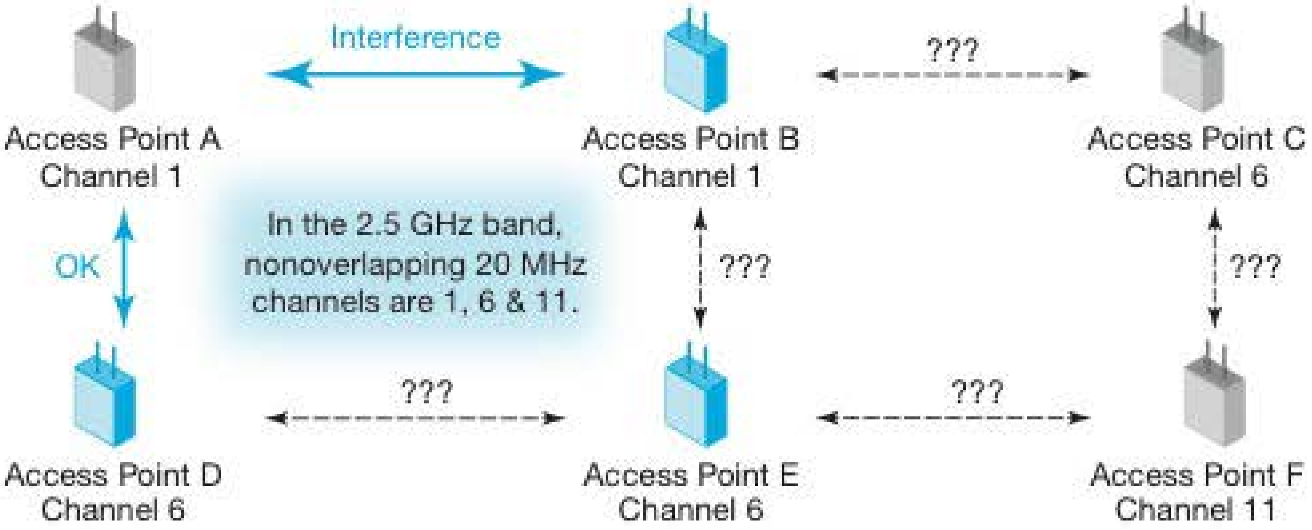

FIGURE 6-12 Channels and Co-Channel Interference in Wi-Fi

in the service bands that Wi-Fi uses, so if there are many nearby access points, some of them will inevitably suffer co-channel interference. Reducing co-channel interference is an important goal in design.

Test Your Understanding

12. In Figure 6-12, there are question marks between several pairs of routers. For each of these pairs, list their channels of operation and whether they will interfere.

The 2.4 GHz and 5 GHz Unlicensed Service Bands

802.11 Wi-Fi WLANs today use two unlicensed service bands. One is the 2.4 GHz unlicensed band. The other is the 5 GHz unlicensed band.

The 2.4 GHz Unlicensed Service Band

The **2.4 GHz unlicensed service band** is the same in most countries in the world. Unfortunately, it only has 83.5 MHz of

The 2.4 GHz Unlicensed Service Band

2.400 GHz to 2.835 GHz for the entire unlicensed service band.

This is small total bandwidth (435 MHz).

There can only be three nonoverlapping 20 MHz channels.

Difficult to put nearby access points on different channels.

If not, there will be co-channel interference.

The 5 GHz Unlicensed Service Band

Slightly shorter propagation distance because of higher absorption at higher frequencies.

Deader dead zones because of higher frequencies.

More bandwidth than the 2.4 GHz band.

Usually allows nearby access points to operate on nonoverlapping channels.

With increasingly wider channels, the ease of channel selection is declining.

FIGURE 6-13 The 2.4 GHz and 5 GHz Unlicensed Service Bands (Study Figure)

total service band bandwidth. Traditionally, each 802.11 channel was 20 MHz wide, although 40 MHz bandwidth channels were introduced in 802.11n. Due to the way channels are allocated, there are only three possible nonoverlapping 20 MHz 802.11 channels. These are centered at Channels 1, 6, and 11.[3] In addition, there can only be a single 40 MHz channel, and if an 802.11n station finds itself in a crowded area, it will drop back from 40 MHz to 20 MHz to reduce interference. This will, of course, cut transmission speed in half.

The 5 GHz Service Band Wi-Fi also operates in the **5 GHz unlicensed service band**. The big advantage of the 5 GHz band is that it is far wider than the 2.4 GHz band. In contrast to the 2.4 GHz band's mere three 20 MHz channels, the 5 GHz band provides between 11 and 24 nonoverlapping 20 MHz channels today, depending on the country. This number of channels in the 5 GHz band is going down as channels become wider to provide higher speed per channel. The 5 GHz unlicensed band will soon be as crowded as the 2.4 GHz unlicensed radio band.

Adding to the attractiveness of the 5 GHz unlicensed band, regulators in several countries have been expanding it to add more total bandwidth and therefore more channels. The United States added more bandwidth in 2003. In 2013, the Federal Communications Commission announced that it would add a further 35%. In contrast, the 2.4 GHz band has no expansion potential because it is bordered by services that cannot be moved.

Test Your Understanding

13. a) In what two service bands does 802.11 operate? b) How many 20 MHz nonoverlapping channels does the 2.4 GHz band support? c) Why is this a problem? d) Why are companies moving rapidly into the 5 GHz band? e) If you triple channel bandwidth, what happens to the number of channels in a service band? (The answer is not directly in the text.)

SPREAD SPECTRUM TRANSMISSION

At the frequencies used by WLANs, there are numerous propagation problems. To address the worst of these problems, multipath interference, regulators mandate the use of a form of transmission called spread spectrum transmission (Figure 6-14). **Spread spectrum transmission** uses far wider channel bandwidth than the transmission speed requires. However, there is no increase in total energy. The signal is simply spread out. Consequently, there is no increase in speed when these wider channels are used.

[3] Channel numbers were defined for the 2.4 GHz band when channels were much narrower. A 20 MHz 802.11 channel overlaps several initially defined channels. Channels 1, 6, and 11 operate in the 2.402 GHz to 2.422 GHz, 2.427 GHz to 2. 447 GHz, and 2.452 GHz to 2.472 GHz frequency ranges, respectively.

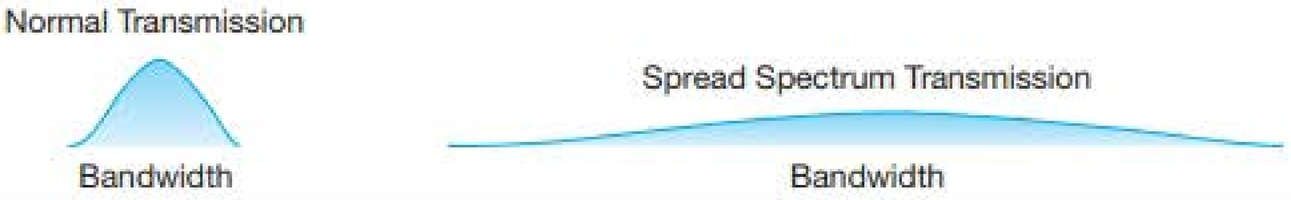

In normal transmission, channel bandwidth is selected to meet the speed requirements of the signal.

In spread spectrum transmission, the signal is spread over a much wider bandwidth.

In spread spectrum transmission, there is no more energy; it is merely spread out.

So there is no increase in transmission speed with spread spectrum transmission.

Goal: Reduce propagation effects at specific frequencies, mainly multipath interference.

Done to improve transmission reliability, not to increase speed.

Not done for security as in military spread spectrum transmission.

FIGURE 6-14 Spread Spectrum Transmission

Spread spectrum transmission uses far wider channel bandwidth than the transmission speed requires.

It is required by regulators to reduce multipath interference problems at Wi-Fi frequencies.

Spread spectrum channels are much wider than normal channels, but they do not transmit signals faster.

Normal versus Spread Spectrum Transmission

Spread spectrum transmission transmits signals redundantly across its broad channel bandwidth, so that if there are transmission problems at some frequencies, the signal will still get through.[4]

In wireless LANs, spread spectrum transmission is used to reduce propagation problems, not to provide security or higher transmission speed.

In commercial spread spectrum transmission, security is *not* a benefit. The military uses spread spectrum transmission for security, but it does so by keeping certain parameters of its spread spectrum transmission secret. Commercial spread spectrum transmission must make these parameters publicly known to allow parties to communicate easily.

Test Your Understanding

14. a) In the 2.4 GHz and 5 GHz service bands, what type of transmission method is required by regulators? b) What is the benefit of spread spectrum transmission for business communication? c) Is spread spectrum transmission done for security reasons in commercial WLANs? d) Does spread spectrum transmission increase transmission speed thanks to its wider channels?

[4] Spread spectrum transmission was invented by Hollywood Actress Hedy Lamar and composer George Antheil during World War II. Their idea was to transmit RADAR over a very wide range of frequencies so that German interference, which was limited to narrow frequency ranges, would not prevent most of the signal from getting through. Their invention was overlooked at the time.

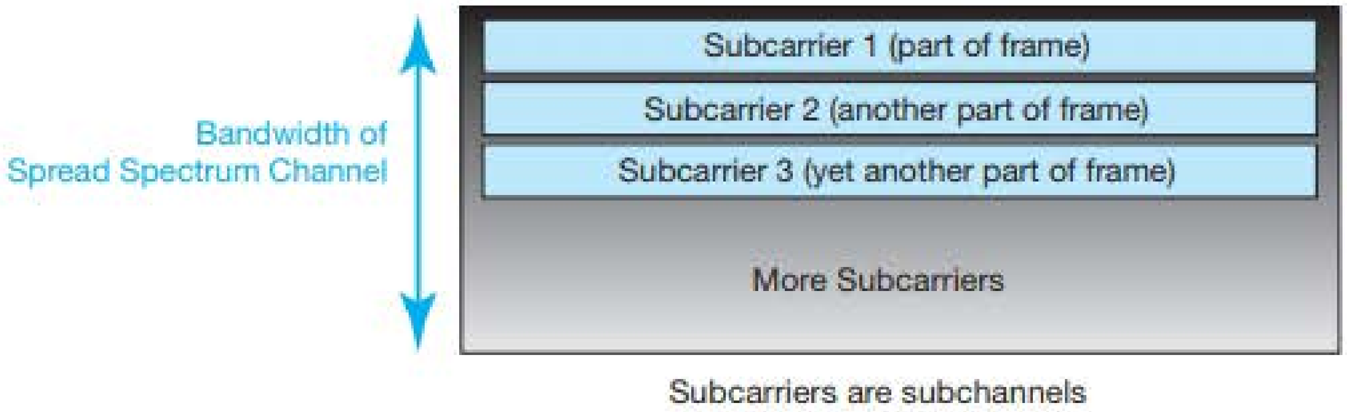
Bandwidth of
Spread Spectrum Channel
Subcarrier 1 (part of frame)
Subcarrier 2 (another part of frame)
Subcarrier 3 (yet another part of frame)
More Subcarriers
Subcarriers are subchannels

FIGURE 6-15 Orthogonal Frequency Division Multiplexing (OFDM)

Orthogonal Frequency Division Multiplexing (OFDM) Spread Spectrum Transmission

There are several spread spectrum transmission methods. The 802.11 Working Group's current standards almost exclusively use **orthogonal frequency division multiplexing (OFDM)**, which Figure 6-15 illustrates.

In OFDM, each broadband channel is divided into many smaller subchannels called subcarriers. OFDM transmits part of a frame in each subcarrier. OFDM sends data redundantly across the subcarriers, so if there is impairment in one or even a few subcarriers, all of the frame will usually still get through.

Why use subcarriers instead of simply spreading the signal over the entire channel? The problem is that sending data over a very wide channel reliably is technically difficult. It is much easier to send many slow signals in many small subcarriers.

Test Your Understanding

15. a) What spread spectrum transmission method dominates today? b) Why does it divide the channel into subcarriers?

802.11 WLAN OPERATION

From 802.11 to 802.3

As Figure 6-16 shows, when a wireless host wishes to send a frame to a server, it transmits the frame to a wireless access point.

- When the wireless host transmits, it puts the packet into an 802.11 frame.[5]
- The frame arrives at the access point. Of course, an 802.11 frame cannot travel over the 802.3 LAN. Wi-Fi has an entirely different frame organization, and Ethernet switches have no idea how to handle 802.11 frames. The access point cannot simply pass the frame on.
- To address this problem, the access point removes the packet from the 802.11 frame and places the packet in an 802.3 Ethernet frame.

[5] 802.11 frames are much more complex than 802.3 Ethernet frames. Much of this complexity is needed to counter wireless propagation problems.

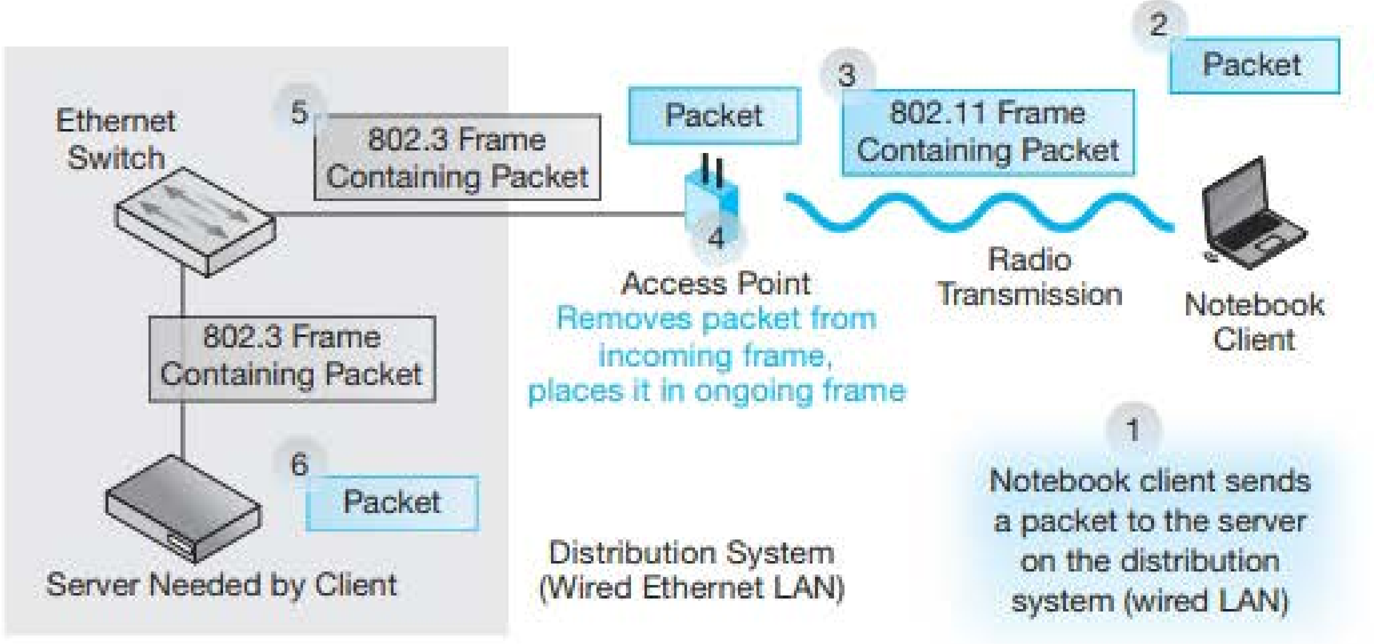
Ethernet Switch
5
802.3 Frame Containing Packet
Packet
3
802.11 Frame Containing Packet
2
Packet
4
Access Point
Removes packet from incoming frame, places it in ongoing frame
Radio Transmission
Notebook Client
802.3 Frame Containing Packet
6
Packet
Server Needed by Client
Distribution System (Wired Ethernet LAN)
1
Notebook client sends a packet to the server on the distribution system (wired LAN)

FIGURE 6-16 Packet and Frame Transmission

- The access point then sends this 802.3 frame to the Ethernet network, which delivers the 802.3 frame to the server.
- Later, when the server replies, the wireless access point receives the 802.3 frame, removes the packet from the Ethernet frame, and forwards the packet to the wireless host in a Wi-Fi frame.[6]

The packet goes all the way from the wireless host to a server. The 802.11 frame travels only between the wireless host and the wireless access point. The 802.3 frame travels only between the wireless access point and the server.

Test Your Understanding

16. a) Why must an access point remove an arriving packet from the frame and place the packet in a different frame when it sends the packet back out? b) Describes the steps that occur when the server transmits a packet back to the wireless client.

Wireless Networks with Multiple Access Points

Access points have limited signal range. To serve a large building or other physical areas, a company must install many access points. The user connects to the nearest access point. To do this, the user must know its **service set ID (SSID)**, which is its name.[7] This is not a

[6] This sounds like what a router does. However, a router can connect any two single networks. Access points are limited to connecting 802.3 and 802.11 networks.

[7] The first author once gave his access point the SSID EvilHacker. He changed it when his neighbors expressed nervousness about seeing it on their list of available access points. On the positive side, there were no attempts by outside hosts to connect to his access point.

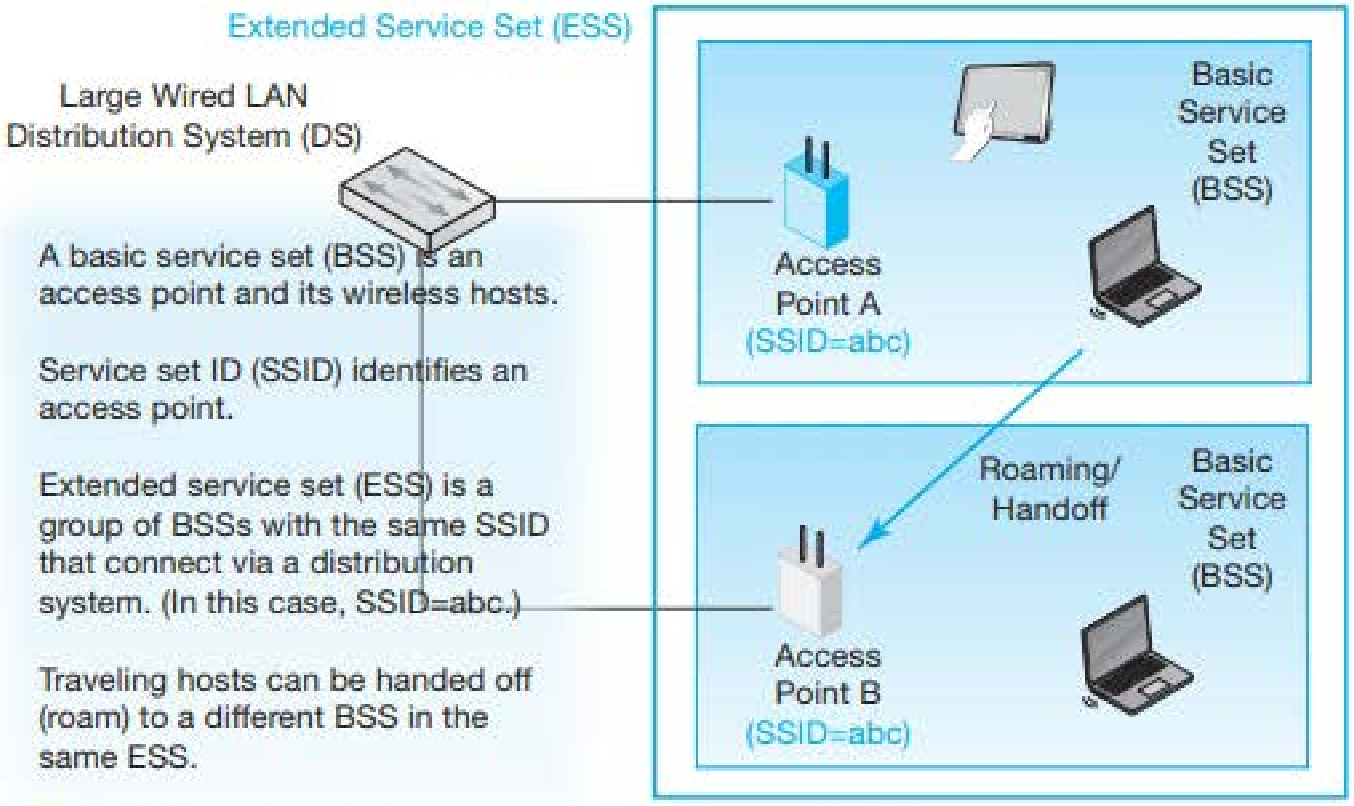

A basic service set (BSS) is an access point and its wireless hosts.

Service set ID (SSID) identifies an access point.

Extended service set (ESS) is a group of BSSs with the same SSID that connect via a distribution system. (In this case, SSID=abc.)

Traveling hosts can be handed off (roam) to a different BSS in the same ESS.

FIGURE 6-17 Wi-Fi Wireless LAN with Multiple Access Points

problem, because all Wi-Fi devices show you the available SSIDs of nearby access points. You just pick the one you want to connect to.

Companies with multiple access points would like their access points to work together. For example, if you connect to the access point in a classroom and then go to the cafeteria after class, you would like to keep your connection without having to connect again in the cafeteria. This is called **roaming**, and it is part of the 802.11 standard. As you pass through several access points on the way to the cafeteria, the one you are leaving and the one you are entering can automatically pass you from the former to the latter.

How do access points know that they are on the same network? The simple answer is that they all have the same SSID. In Figure 6-17, the two access points shown have the unimaginative SSID *abc*. The access point in a single network also needs to transmit messages back and forth to do roaming and other things. They normally do this through the company's switched Ethernet network. In 802.11 jargon, this is called the **distribution system**.

Test Your Understanding

17. a) What is a roaming in 802.11? b) What characteristics do all access points in a corporate network share? c) Over what transmission system do access points communicate with each other to accomplish roaming? c) Distinguish between a BSS, and ESS, and an SSID.

Media Access Control

The access point and all of the wireless hosts it serves transmit and receive in a single channel. Figure 6-18 shows that if two devices transmit in the same channel at the same time, their signals will interfere with each other. This is called a

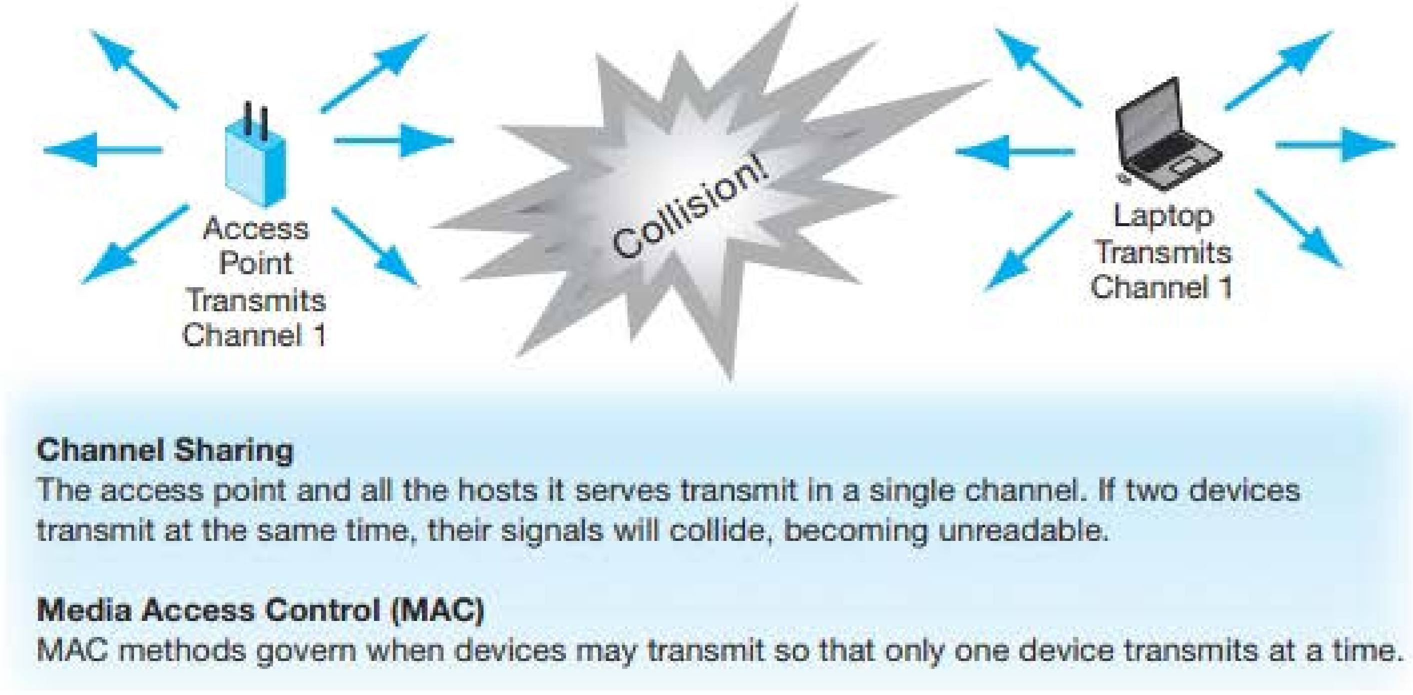

Channel Sharing
The access point and all the hosts it serves transmit in a single channel. If two devices transmit at the same time, their signals will collide, becoming unreadable.

Media Access Control (MAC)
MAC methods govern when devices may transmit so that only one device transmits at a time.

FIGURE 6-18 Hosts and Access Points Transmit on a Single Channel

collision. It makes both signals unreadable. When a wireless host or the access point transmits, all other devices must wait. As the number of hosts served by an access point increases, individual throughput falls because of this waiting. The box "Media Access Control (MAC)" discusses how **media access control (MAC)** methods govern when hosts and access points may transmit so that collisions are avoided.[8]

Media access control (MAC) methods govern when hosts and access points may transmit so that collisions can be avoided.

The access point and all of the wireless hosts it serves transmit and receive in a single channel. When a wireless host or the access point transmits, all other devices must wait.

Test Your Understanding

18. All wireless hosts and the access point that serves them transmit on the same channel. a) What problem does this cause? b) How does media access control (MAC) address this problem? c) Does media access control apply to wireless hosts, access points, or both? d) Can a wireless access point and one of the wireless clients in its BSS transmit simultaneously?

[8] Yes, this is where the term *MAC address* comes from. Conceptually, Media Access Control is a sublayer of the data link layer. It applies to Ethernet, Wi-Fi, and other 802.11 standards. Addresses are defined at this layer so that all 802.11 standards use EUI-48 addresses.

IN MORE DEPTH

Media Access Control (MAC)

The 802.11 standard has two mechanisms for media access control. The first, CSMA/CA+ACK, is mandatory. Access points and wireless hosts must support it. The second, RTS/CTS, is optional.[9]

CSMA/CA+ACK Media Access Control

The mandatory method is Carrier Sense Multiple Access with Collision Avoidance and Acknowledgment, which is mercifully shortened to **CSMA/CA+ACK**.

Carrier sense (CS) means to listen to (sense) traffic (the carrier, in radio parlance). Multiple access (MA) means that this method uses listening to control how multiple hosts can access the network to transmit. Quite simply, if another device is transmitting, the wireless host or access point does not transmit.

Collision avoidance (CA) means that the method attempts to avoid two devices transmitting at the same time. One issues that if one device has been sending for some time, two or more others may be waiting to send. If they both send as soon as the current sender stops, they will both transmit at the same time. (We have all been in conversations like this.) This will cause a collision. Collision avoidance adds a random delay time to decide which device may transmit first. This works, but it is inefficient because it adds dead time when no one is transmitting. If nobody has been transmitting for a long time, this random delay step is skipped because the likelihood of a collision is small.

CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance)

- Sender listens for traffic
- Carrier is the signal; sensing is listening
  1. If there is traffic, waits
  2. If there is no traffic:
     - 2a. If there has been no traffic for less than the critical time value, waits a random amount of time, then returns to Step 1.
     - 2b. If there has been no traffic for more than the critical value for time, sends without waiting.
- These steps avoid the collision that would result if hosts could transmit as soon as one host finishes transmitting.

ACK (Acknowledgment)

- Receiver immediately sends back an acknowledgment
- If sender does not receive the acknowledgment, retransmits using CSMA
- CSMA/CA plus ACK is a reliable protocol

FIGURE 6-19 CSMA/CA+ACK Media Access Control (Study Figure)

[9] Actually, if you have even a single host with older 802.11b equipment connected to an access point, RTS/CTS becomes mandatory. However, 802.11b wireless hosts are almost never encountered anymore.

(continued)

ACK means that if the receiver receives a message correctly, it immediately sends an acknowledgment to the sender, not waiting at all. This is another reason to require stations to delay before sending when a sender stops transmitting.

If the sender does not receive an ACK, it retransmits the frame. Sending acknowledgments and doing retransmissions makes 802.11 Wi-Fi transmission reliable because it provides both error detection and error correction. CSMA/CA+ACK is the only reliable transmission method we will see in this book other than TCP. Most early wired DLL protocols were reliable because transmission then was unreliable, even in wired networks. Under these circumstances, error correction at the data link layer made sense. This is no longer true today generally. Today, wired transmission protocols such as Ethernet are unreliable. Doing error correction is simply not worth the effort at each hop between switches when transmission errors are rare. We have seen that wireless transmission, however, is encumbered with propagation problems, and lost or damaged frames are far too common. It makes sense under these conditions to make 802.11 (and many other wireless protocols) reliable.

Thanks to CSMA/CA+ACK, 802.11 is a reliable protocol.

CSMA/CA+ACK works well, but it is inefficient. Waiting before transmission wastes valuable time. Sending ACKs and doing retransmissions also is time consuming. Overall, an 802.11 LAN has throughput substantially lower than rated speeds.

Test Your Understanding

19. a) What does CS mean? (Do not just spell out the abbreviation.) b) How is carrier sensing used in multiple access? c) Why is CA desirable? d) Does a frame's receiver transmit an ACK immediately or after a random delay? e) Is CSMA/CA+ACK reliable or unreliable? f) Why was 802.11 made reliable? g) Is CSMA/CA+ACK efficient?

Request to Send/Clear to Send (RTS/CTS)

Although CSMA/CA+ACK is mandatory, there is another control mechanism called request to send/clear to send (**RTS/CTS**). Figure 6-20 illustrates RTS/CTS. As noted earlier, the RTS/CTS protocol is usually optional. Avoiding RTS/CTS whenever possible is wise because RTS/CTS is much less efficient, and therefore slower, than CSMA/CA+ACK.

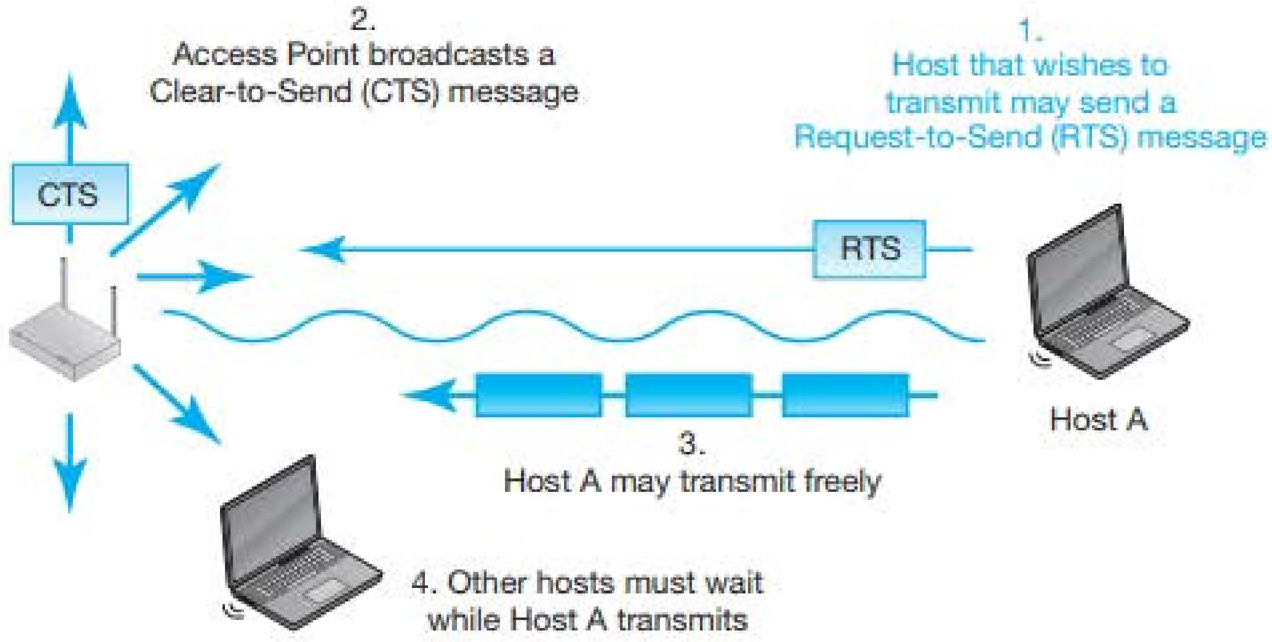
2.
Access Point broadcasts a
Clear-to-Send (CTS) message
1.
Host that wishes to
transmit may send a
Request-to-Send (RTS) message
CTS
RTS
Host A
3.
Host A may transmit freely
4. Other hosts must wait
while Host A transmits

FIGURE 6-20 Request to Send/Clear to Send Media Access Control

- When a host wishes to send, the host may send a request-to-send (RTS) message to the wireless access point. This message asks the access point for permission to send messages for a short period of time on an exclusive basis. It is like someone asking for recognition by a human meeting director so that they may take the floor.
- If the access point responds by broadcasting a clear-to-send (CTS) message, then other hosts must wait. The host sending the RTS may then transmit, ignoring CSMA/CA.

RTS/CTS makes sense primarily when two wireless clients can both hear the access point but cannot hear each other. With CSMA/CA+ACK, the two stations may transmit at the same time. RTS/CTS eliminates this.

Test Your Understanding

20. a) Describe RTS/CTS. b) Is CSMA/CA+ACK required or optional? c) Is RTS/CTS required or optional? d) Which is more efficient, RTS/CTS or CSMA/CA+ACK? e) When does it make sense to use RTS/CTS?

802.11 TRANSMISSION STANDARDS

The 802.11 Working Group has created several WLAN transmission standards since 1997. We will look at the two important standards today, 802.11n and 802.11ac.[10]

Channel Bandwidth and Service Band Bandwidth

One major difference between 802.11n and 802.11ac is channel bandwidth. Recall that wider channel bandwidth means faster speed. As Figure 6-21 shows, the total bandwidth in the 5 GHz service band is about 665 MHz in the United States. The normal bandwidth of 802.11n is 40 MHz. Therefore, there can be about ten 40 MHz 802.11n channels in the 5 GHz unlicensed service band. That is a good number of channels, but speed is limited.

	802.11n 40 Hz Channels	802.11ac 80 MHz Channels	802.11ac 160 Hz Channels
Total Service Band Bandwidth	665 MHz	665 MHz	665 MHz
Channel Bandwidth	40 MHz	80 MHz	160 MHz
Total Service Band Bandwidth Divided by Channel Bandwidth	16.6	8.3	4.2
Actual Number of Channels	12	6	2

FIGURE 6-21 Number of Channels in the 5 GHz Unlicensed Radio Band

[10] Sometimes access points also have to deal with stations communicating with the older 802.11g standard in the 2.4 GHz band and the even older 802.11a standard in the 5 GHz band.

The 802.11ac standard, in contrast, has either 80 MHz or 160 MHz channels. That means more speed—twice or four times the speed of 802.11n's 40 MHz channels. (Other things make 802.11ac transmission even faster.) However, with a total bandwidth of about 665 MHz, there can only be six 80 MHz channels and only two 160 GHz channels, due to the way that channels are allocated.

As channel bandwidth increases, the number of channels in a service band decreases proportionately.

Test Your Understanding

21. If you triple channel bandwidth in a service band, what happens to the number of channels in a service band?

Speed and Market Status

Figure 6-22 compares the **802.11n** and **802.11ac** standards. 802.11n products deliver speeds of 100 to 600 Mbps. The newer 802.11ac standard delivers far higher rated speeds of 433 Mbps to 6.9 Gbps. 802.11n still dominates the installed base today, but 802.11ac dominates sales and will soon supplant 802.11n as the dominant Wi-Fi technology.

Test Your Understanding

22. a) Compare the rated speeds of 802.11n and 802.11ac. b) Compare the market status of 802.11n and 802.11ac. c) If you need an access point providing 3 Gbps service, what choice do you have?

Characteristic	802.11n Dual Band	802.11ac
Rated Speed	100 Mbps to 600 Mbps.	433 Mbps to 6.9 Gbps.
Status	Widely used	Widely used and dominates sales
Unlicensed Band(s)	2.4 GHz and 5 GHz	5 GHz
Channel bandwidth	40 MHz, but will drop back to 20 MHz if there is interference with older 20 MHz standards	80 MHz or 160 MHz
Number of Non-Overlapping Channels (varies by country)	3 in 2.4 GHz band; 12 in the United States in 5 GHz band	6 at 80 MHz channel bandwidth and 2 at 160 MHz channel bandwidth in the United States in the 5 GHz band
Maximum MIMO spatial streams	4	8
Multi-User MIMO / Beamforming?	No	Yes

FIGURE 6-22 Characteristics of Major 802.11 Wi-Fi Standards

- **The rated speed of the access point**
    - **The actual throughput (aggregate) of the access point, which is lower**
    - **The number of users transmitting simultaneously determines individual throughput**
- **Your distance from the access point will also affect how fast the access point transmits to you**
    - **As you travel away, the access point will transmit more slowly to be more easily understood**

FIGURE 6-23 Your Individual Throughput Will Vary. A Lot. (Study Figure)

Your Service Speed Will Vary. A Lot.

The *rated speed* of a network is the speed that it called for in the standard. This is the speed that is advertised on the box. In reality, *throughput*—the speed your network actually provides—is always lower, often substantially lower. Let's say that the rated speed of your access point is 600 Mbps. The throughput might be 500 Mbps—or a lot lower.

In addition, the access point and the wireless clients all transmit in a single channel. As we saw earlier, when the access point or a wireless host is transmitting, the others must wait. The 500 Mbps for an access point, then, is aggregate throughput. Suppose that an access point serves 50 devices. At a particular moment in time, 5 of them wish to transmit. These 5 would share the aggregate throughput. On average, each would get 100 Mbps of *individual throughput*. If 20 wished to transmit, each would receive only 25 Mbps.

Things get worse for an individual wireless host as it gets farther from the access point. The signal degrades with distance, creating more transmission errors. Standards compensate for this by transmitting more slowly to a host if errors become substantial. This reduces errors, but it also reduces individual throughput.

In general, given the uncertainties involved in rated speed versus throughput, individual throughput changing as the number of stations needing to transmit changes, and individual throughput changing as a function of distance, it is impossible to know how much transmission speed you will get as a wireless client. It is always good to install a program that measures your actual upload and download speed when you are being served by an access point.

Test Your Understanding

23. a) You are using an access point with a rated speed of 4 Gbps. Why will you experience much less speed? b) What will happen to your speed as you move away from the access point?

Multiple Input/Multiple Output (MIMO)

Increasing bandwidth is the easiest way to boost transmission speed, but there is a more elegant way to increase speed without increasing bandwidth. Figure 6-24 notes that standards beyond 802.11g use a technique called **multiple input/multiple output (MIMO)** to double, triple, or quadruple transmission speed (or even increase it more) without increasing channel bandwidth.

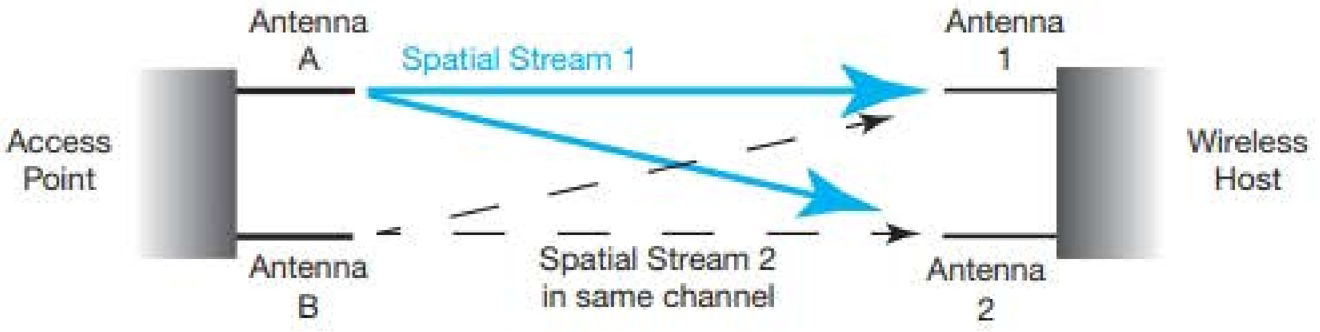

Two spatial streams are sent in the same channel, but from different sending antennas. The two signals arrive at slightly different times at the two receiving antennas. This allows the receiver to distinguish between the two signals.

FIGURE 6-24 Multiple Input/Multiple Output (MIMO) Operation

The key to higher throughput in MIMO is that the host or access point sends two or more spatial streams (radio signals) in the same channel between two or more different antennas on access points and wireless hosts. Earlier, we said that that was impossible. Actually, it used to be impossible, but newer technology has made this possible.

In the figure, there are two spatial streams. Each carries different information. As we saw earlier in this chapter, two signals in the same channel should interfere with each other. However, the two spatial streams sent by different antennas will arrive at the two receiving antennas with slightly different time lags. Using detection and separation methods based on differences in arrival times for the two spatial streams, the receiver can separate the two spatial streams in the same channel and so can read them individually.

Even with only two spatial streams using two antennas each on the sender and receiver, MIMO can roughly double throughput. Using more antennas and therefore more spatial streams can increase throughput even more. MIMO is not limited to two spatial streams.

The 802.11n standard introduced MIMO to Wi-Fi. With two spatial streams, the rated speed in 802.11n with 40 MHz channels is 300 Mbps. Three spatial streams raise the rated speed to 450 Mbps, and four raise it to 600 Mbps. The 802.11n standard requires access points to support four spatial streams, although wireless hosts are only required to support two spatial streams. Typical speeds in 802.11n products today have rated speeds of about 300 Mbps.

The 802.11ac standard, in addition to doubling or quadrupling channel bandwidth compared to 802.11n, doubles the number of possible spatial streams to eight. The standard offers 16 possible combinations of bandwidth (80 MHz or 160 MHz) and number of spatial streams (1 to 8). This creates a large number of possible rated speeds: 433 Mbps to 6.9 Mbps. Products today typically provide rated speeds of about 1.5 Gbps, but speed is increasing rapidly.

Another benefit of MIMO, beyond greater transmission speed, is greater transmission range. Greater propagation distances may permit fewer access points to be installed, and this will lower equipment and installation cost.

Test Your Understanding

24. a) How does MIMO use spatial streams to increase transmission speed? b) What is the main benefit of MIMO? c) What is its other benefit?

Beamforming and Multiuser MIMO

Today, jet fighters use phased array radar systems that are flat dishes with many tiny antennas spread over the surface. Controlling the relative phases of the signals from these antennas can focus the radar beam in a particular direction very rapidly. The antennas on advanced MIMO systems can do the same, focusing the radio power instead of broadcasting it isotropically (in all directions equally). Figure 6-25 illustrates this focusing, which is called **beamforming**.

Obviously, beamforming means that when the access point transmits to (or receives from) a wireless device the signal will be stronger. The radio can either operate at lower power or send the signal farther.

Beamforming also allows **multiuser MIMO (MU-MIMO)**, in which the access point focuses on two wireless devices at the same time. With focused transmissions, it can communicate with two or more devices simultaneously. This eliminates the time a device may have to wait before transmitting in order to avoid collisions.

Test Your Understanding

25. a) What is beamforming? b) What benefits can it bring? c) Distinguish between MIMO and multiuser MIMO (MU-MIMO).

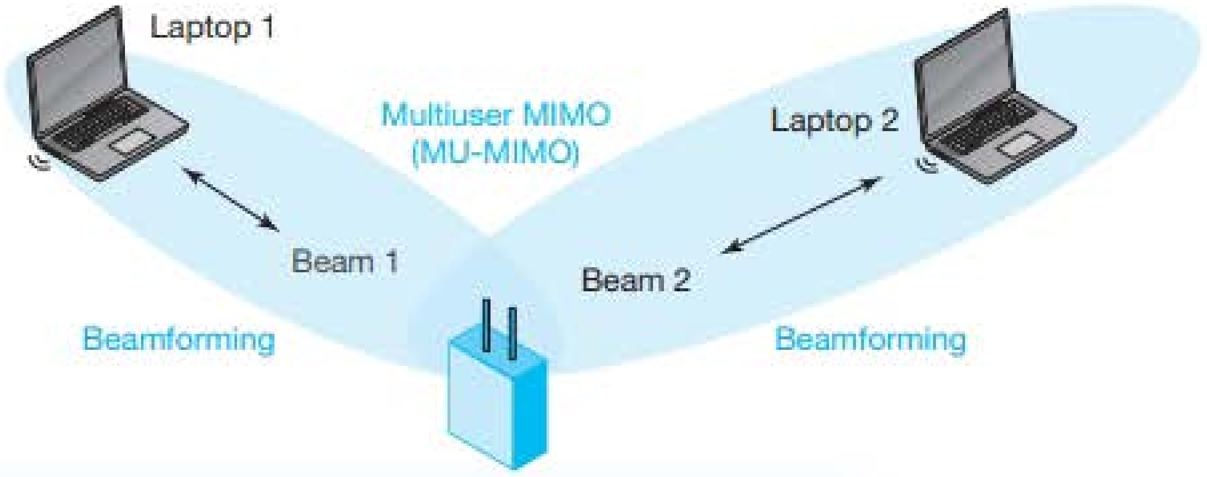

Beamforming can direct signal energy toward individual devices.
This sends stronger signals, bringing longer range.
It can also allow an access point to communicate with multiple devices in a single channel through multiuser MIMO (MU-MIMO).

FIGURE 6-25 Beamforming and Multiuser MIMO

IN MORE DEPTH

802.11/WI-FI NOTES

Backward Compatibility

You have several 802.11n devices in your home—two notebooks, two laptop computers, a voice-activated home controller, and an access point. You are thinking of upgrading your access point to 802.11ac. Will your devices work with the new access point? They will. Wi-Fi devices have **backward compatibility**. This means that new devices will always work with existing devices (although perhaps not with truly ancient devices). This is good, because otherwise you would have to throw out all your 802.11n devices, or at least add an external USB device to implement 802.11ac (see Figure 6-26).

This does not mean that your devices will magically operate at 802.11ac speeds. Their radios can only give 802.11n speeds. This is where backward compatibility comes in. The new 802.11ac access point retains the ability to transmit 802.11n signals, including those in the 2.4 GHz channel. All the devices, then, will transmit using 802.11n. If you want 802.11ac speeds, you will have to buy new client devices or buy an external 802.11ac device. The good news is that you are free to upgrade them individually or not at all, waiting until you buy new 802.11ac-compatible devices.

Test Your Understanding

26. a) You are considering a laptop computer that uses 802.11ay. (802.11ay is discussed in the next subsection.) Will your existing 802.11ac access point be able to communicate with the new device? b) What standard will they use in the communication if communication is possible? c) What principle does this communication exemplify?

Profile Waves for Wi-Fi Devices

As noted earlier, the Wi-Fi Alliance tests for interoperability between 802.11 devices. Until an 802.11 device pass certification, it cannot display the Wi-Fi logo on its box.

However, 802.11 standards have many options, some of which would be impossible or at least very expensive to implement when a standard is first released. The Wi-Fi Alliance addresses this challenge by releasing a series of **profile waves** over time, each specifying certain things that must be included.

The Wave 1 profile for a new standard always gives a good increase in performance compared to the previous standard. The Wave 2 profile gives still better performance by implementing more advanced options. Waves may continue beyond Wave 2, although by the time there is a need for a third wave, a better standard is often available. It is not unusual for profile waves to stop below the theoretical best speeds of the standard.

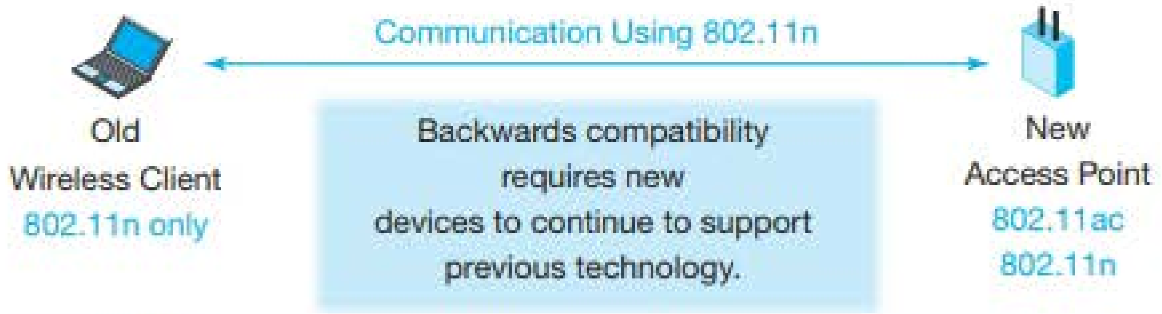

FIGURE 6-26 Backward Compatibility

Standards Have Many Options

- Some may be impossible or too expensive to implement initially
- The Wi-Fi Alliance defines profile waves by doing compatibility testing
- Devices must be tested for compatibility with a particular profile wave
- Only then are they certified as Wi-Fi compliant

Wave Profile Progression

- Wave 1 profiles are usually good improvements over past standards
- Wave 2 profiles provide more speed and other features
- Confusingly, wave profiles themselves have options
- This gives sometimes unwelcome variability for performance between device pairs

802.11ac

- Wave 1 profile gives a data stream of up to 1.3 Gbps
- Wave 2 profile gives a data stream of 2.5 Gbps, plus MU-MIMO

FIGURE 6-27 Profile Waves for Wi-Fi Devices (Study Figure)

To give an example, 802.11ac Wave 1 products were limited to 80 MHz channels and up to three MIMO spatial streams. At the high end, this gives a data rate of 1.3 Gbps. A cellular telephone with only a single antenna will probably receive a throughput around 250 Mbps. A wireless computer with three antennas will probably see about 750 Mbps.

The Wave 2 802.11ac profile brings 160 MHz channels (if there is room for them) and a fourth antenna to give four MIMO spatial streams. This and other improvements will bump the data stream to about 2.5 Gbps.

This is roughly a doubling in potential speed, but newer waves also introduce features beyond speed. For example, Wave 2 adds MU-MIMO capability, allowing separate beamformed transmission with more than one device simultaneously.

Test Your Understanding

27. a) Why does the Wi-Fi Alliance release compatibility testing profiles in waves instead of combining the entire standard's features initially? b) When someone says that an access point is a Wave 1 802.11ac device, what improvements do you expect to receive with a Wave 2 802.11ac device?

Coming Attractions

The 802.11 Working Group produces new standards constantly. Many are minor standards or are important management standards that you will learn if you go into networking as a career. However, three standards under development are worth knowing about broadly because they have the potential to make 802.11ac a quaint memory.

802.11ax The **802.11ax** standard now under development could be thought of as a supercharged 802.11ac. It uses the same 5 GHz unlicensed band. It seems like a minor upgrade; it will only raise the maximum speed from 7 Gbps to 10 Gbps.

However, 802.11ax addresses a problem that is becoming more important than speed—increasing **density** (the number of hosts per access point). The 802.1ax standard promises to serve four times as many hosts per access point as 802.1ac. It will do so by being more spectrum

(continued)

802.11ax

- In the unlicensed 5 GHz Band
- A little faster than 802.11ac
- But can serve many more stations per access point
- High-density operation is becoming important as the number of devices in an area grows

The 60 GHz Unlicensed Band

- Very high attenuation so short range, strong dead zones, and difficulty penetrating walls
- 14 gigahertz of total bandwidth

802.11ad in the 60 GHz Band

- Up to 7 GHz of rated speeds
- Today's products offer only about half that

802.11ay in the 60 GHz Band

- More sophisticated version of 802.11ad
- Should bump basic speed to 20 to 30 Gbps over longer distances than 802.11ad
- *May* be able to penetrate walls

FIGURE 6-28 Coming Attractions (Study Figure)

efficient, sending more bits per hertz of bandwidth. For example, it will transmit with 1,048 states per clock cycle, and it will introduce a much more efficient media access control mechanism to control when hosts transmit.

The 60 GHz Unlicensed Band: 802.11ad and 802.11ay In corporations, the 5 GHz band is close to being as saturated like the 2.4 GHz band in corporations. The enormous 160 MHz channels of 802.11ac at its most aggressive chew up massive amount of the 5 GHz band.[11]

802.11ad in the 60 GHz Band A new higher-frequency 60 GHz unlicensed band has been approved, and products have begun appearing to exploit it in wireless LANs. The actual range of frequencies varies in different parts of the world, but it is usually very wide. In the United States, the Federal Communication Commission has allocated the frequencies between 57 GHz and 71 GHz. This is 14 GHz in total. The base channel bandwidth is 2.16 GHz, which is wider than the entire 5 GHz band.

The first 802.11 standard for 60 GHz, **802.11ad**, can provide 7 Gbps of speed. The 802.11ac standard's maximum rate is this fast, but 802.11ac products do not reach it. Today, Wave 2 802.11ac speeds are only 3.2 Gbps. The 802.11ad standard is attractive for high-end residential use. It can support wireless communication between a laptop and a television for streaming 4K video. It also can make wireless connections to replace USB cords and do so at ultrahigh speeds. In residences, it is a high-speed cable replacement technology.

High Absorptive Attenuation Although the 60 GHz band has a great deal of capacity, it also has serious propagation problems. In Figure 6-5, we saw that absorptive attenuation

[11] The 5 GHz band does not extend to 6 GHz, and it has sections within its range that have not been approved for unlicensed use.

increases as frequency increases. Consequently, the 60 GHz band has much higher attenuation than the 5 GHz band.[12] Its maximum propagation distance is very short.

Shadow Zones and Clear Lines of Sight Recall the other problem with increasing frequencies: objects block waves far more as frequency increases. Tests have shown that 802.11ad signals cannot go through building walls. In fact, wooden doors stop them almost completely. The 802.11ad standard requires a **clear line of sight** with no obstacles between the access point and a host. In office areas, this is difficult. Together, high absorptive attenuation and strong shadow zones limit 802.11ad to a single room, and not a room with significant obstacles.

802.11ay The existing 802.11ad standard is not very sophisticated. It does not offer MIMO or other advances that 802.11ac has offered for some time. A next-generation 60 GHz standard promises to add MIMO and much more. This **802.11ay** standard is still under development, but it should bump basic speed to 20 to 30 Gbps over substantially longer distances than 802.11ad, and by bonding several channels together, it will be able to provide much higher speeds. Using MU-MIMO, it also can direct energy with beamforming to give much better range. This and other improvements should even allow it to penetrate walls and other obstacles, at least to some extent. Of course, we will have to wait and see.

Test Your Understanding

28. a) What is the main promise of 802.11ax over 802.11ac? b) Why is the 60 GHz unlicensed band attractive? c) What problems does it pose for Wi-Fi? d) How is 802.11ay likely to be better than 802.11ad?

END-OF-CHAPTER QUESTIONS

Thought Questions

6-1. Why might a company decided to use 80 MHz channels in 802.11ac instead of 160 MHz channels?

6-2. a) What do physical layer standards govern? b) What do data link layer standards govern? For the following lettered question parts, say whether the concept is a Layer 1 concern or a Layer 2 concern. Explain your reasoning. c) Multipath interference. d) Media Access Control. e) MIMO. f) Converting between 802.11 frames and 802.3 frames. g) Wireless propagation problems. h) Roaming. i) 802.11ac.

6-3. You can transmit 1.54 Gbps in a channel you use frequently. You want to transmit at 4.32 Gbps. How much wider must your channel be than its current bandwidth?

[12] Right around 60 GHz, there is an even more severe absorption problem. Radio waves at that frequency cause oxygen atoms in the air to vibrate, leaching energy from the signal. Around 60 GHz, attenuation can be up to 100 times higher than it is at neighboring frequencies. The FCC originally allocate the range from 57 GHz to 64 GHz, much of which is heavily affected by oxygen absorption attenuation. In 2016, it added the range from 64 GHz to 71 GHz, which is well past the oxygen absorption peak.

Troubleshooting Question

6-4. You have been using your phone and your school's Wi-Fi network to access hosts on the Internet. Suddenly, you cannot reach Internet hosts. Create a two-column table. a) In the first column, create a list of possible causes. b) In the second column, describe how you would test each one. (You may not be able to test them all.)

Hands On

Chapter 6a has a set of hands-on exercises that will help you make the things you have learned in this chapter more concrete.

Perspective Questions

6-5. What was the most surprising thing you learned in this chapter?

6-6. What was the most difficult part of this chapter for you?

Chapter **6a**

Hands-On: Using Xirrus Wi-Fi Inspector

LEARNING OBJECTIVES

by the end of this chapter, you should be able to:

- Use Xirrus Wi-Fi Inspector with some facility.
- Interpret output from Wi-Fi Inspector in specific situations.
- Do a site survey.

INTRODUCTION

Wi-Fi analysis programs listen to nearby access points (and sometimes wireless hosts) to determine such things as how strong their signals are, what types of security they use, what their SSIDs and BSSIDs are, and sometimes the directions of the individual access points.

There are many Wi-Fi analysis programs for mobile devices. Many have "stumbler" in their names in homage to one of the first examples, NetStumbler. This chapter looks at *Wi-Fi Inspector* from Xirrus, which runs on Microsoft Windows and is available as a free download from Xirrus. A comparable Windows Widget that always remains on the desktop is also available from Xirrus.

THE FOUR WINDOWS

Figure 6a-1 shows the ribbon menu and four tiled windows that appear when you bring up Wi-Fi Inspector. This view shows all information in a single window. This is the default. It is also what you see if you click on Show All in the Layout ribbon.

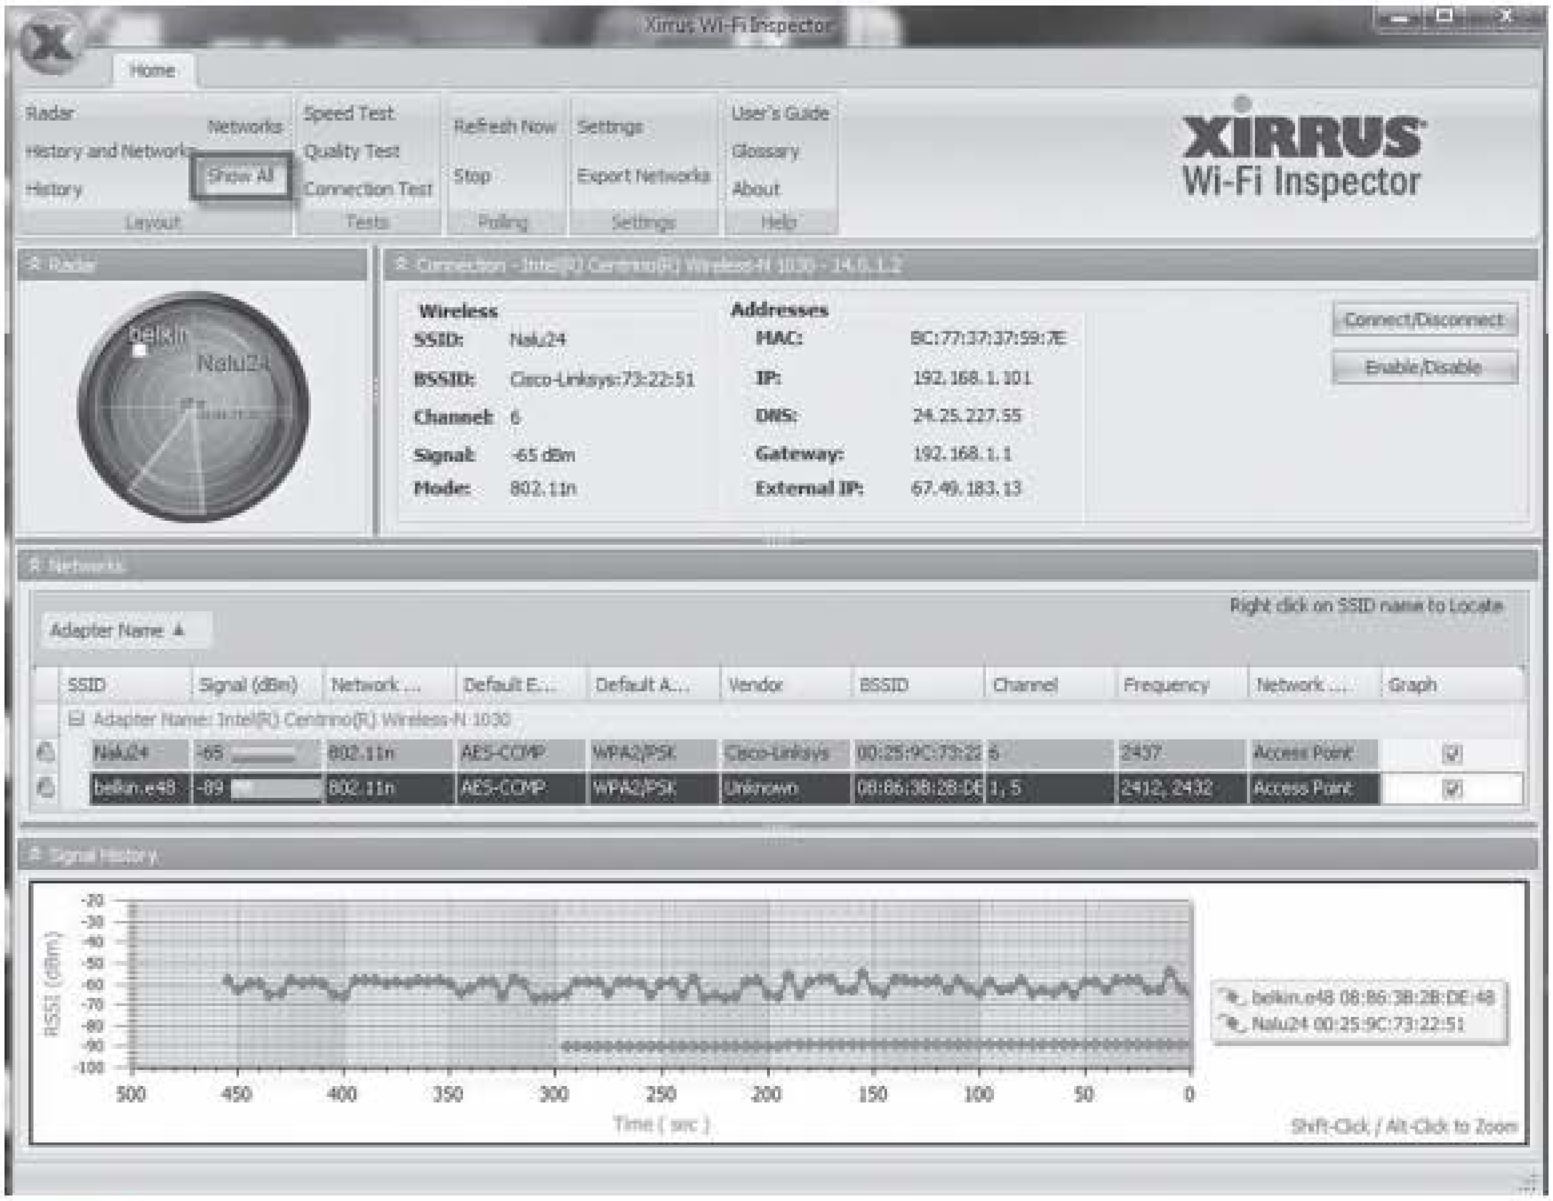

Connection - Intel(R) Centrino(R) Wireless-N 1030 - 14.0.1.2

Wireless

SSID:	Nalu24
BSSID:	Cisco-Linksys:73:22:51
Channel:	6
Signal:	-65 dBm
Mode:	802.11n

Addresses

MAC:	BC:77:37:37:59:7E
IP:	192.168.1.101
DNS:	24.25.227.55
Gateway:	192.168.1.1
External IP:	67.49.183.13

Adapter Name ▲

Right click on SSID name to Locate

SSID	Signal (dBm)	Network ...	Default E...	Default A...	Vendor	BSSID	Channel	Frequency	Network ...	Graph
Adapter Name: Intel(R) Centrino(R) Wireless-N 1030										
Nalu24	-65	802.11n	AES-CCMP	WPA2/PSK	Cisco-Linksys	00:25:9C:73:22	6	2437	Access Point	☑
belkin.e48	-89	802.11n	AES-CCMP	WPA2/PSK	Unknown	08:86:3B:2B:DE	1, 5	2412, 2432	Access Point	☑

FIGURE 6a-1 Four Windows in Wi-Fi Inspector

The Radar Window (Read the Fine Print)

The most obvious window is the radar window, which shows all access points in the vicinity. The access points are spread out across the two-dimensional picture.

Relative Direction (Meaningless) It appears that the radar window shows the relative directions of the access points, much as an air traffic radar display shows the directions of nearby aircraft. Actually, it does not. The access points are merely spread out for readability. Direction is meaningless. In this sense, the radar window is misleading. However, it looks cool.

Distance From the Center (Signal Strength) What does distance from the center mean? It looks like it means physical distance, as it would on a physical radar screen. Rather, it means *signal strength*. Access points that are shown closest to the center are the *strongest*, and access points that are the farthest from the center are the *weakest*.

Measuring Signal Strength Signal strength gives the RSSI (relative signal strength indicator) for the access point. Smaller negative numbers are better. For example, –60 dBm is a very strong signal, while –87 dBm is a very weak signal. In

Figure 6a-1, Nalu24 has a signal strength of –65, which is quite good. Belkin has a signal strength of –89, which is terrible.

For signal strength, smaller negative numbers are better. (It's a double negative.)

Expanding the Radar Window The radar window in its normal small form can only display four access points. Under the Layout section of the menu, selecting Radar in the Layout Group will maximize the radar window. This allows up to ten access point names to be seen. By the way, "network" and "SSID" are synonyms.

Figure 6a-2 shows the expanded radar window. There are only two nearby access points, so there is no need for a large radar window. However, it certainly is easier to read the relative indicated signal strength.

Connection Window

The connection window (in the upper right in Figure 6a-1) shows information about the access point to which the computer running Wi-Fi Inspector is currently connected (Nalu24). It shows the SSID (the network name, in this case, Nalu24), the BSSID (the

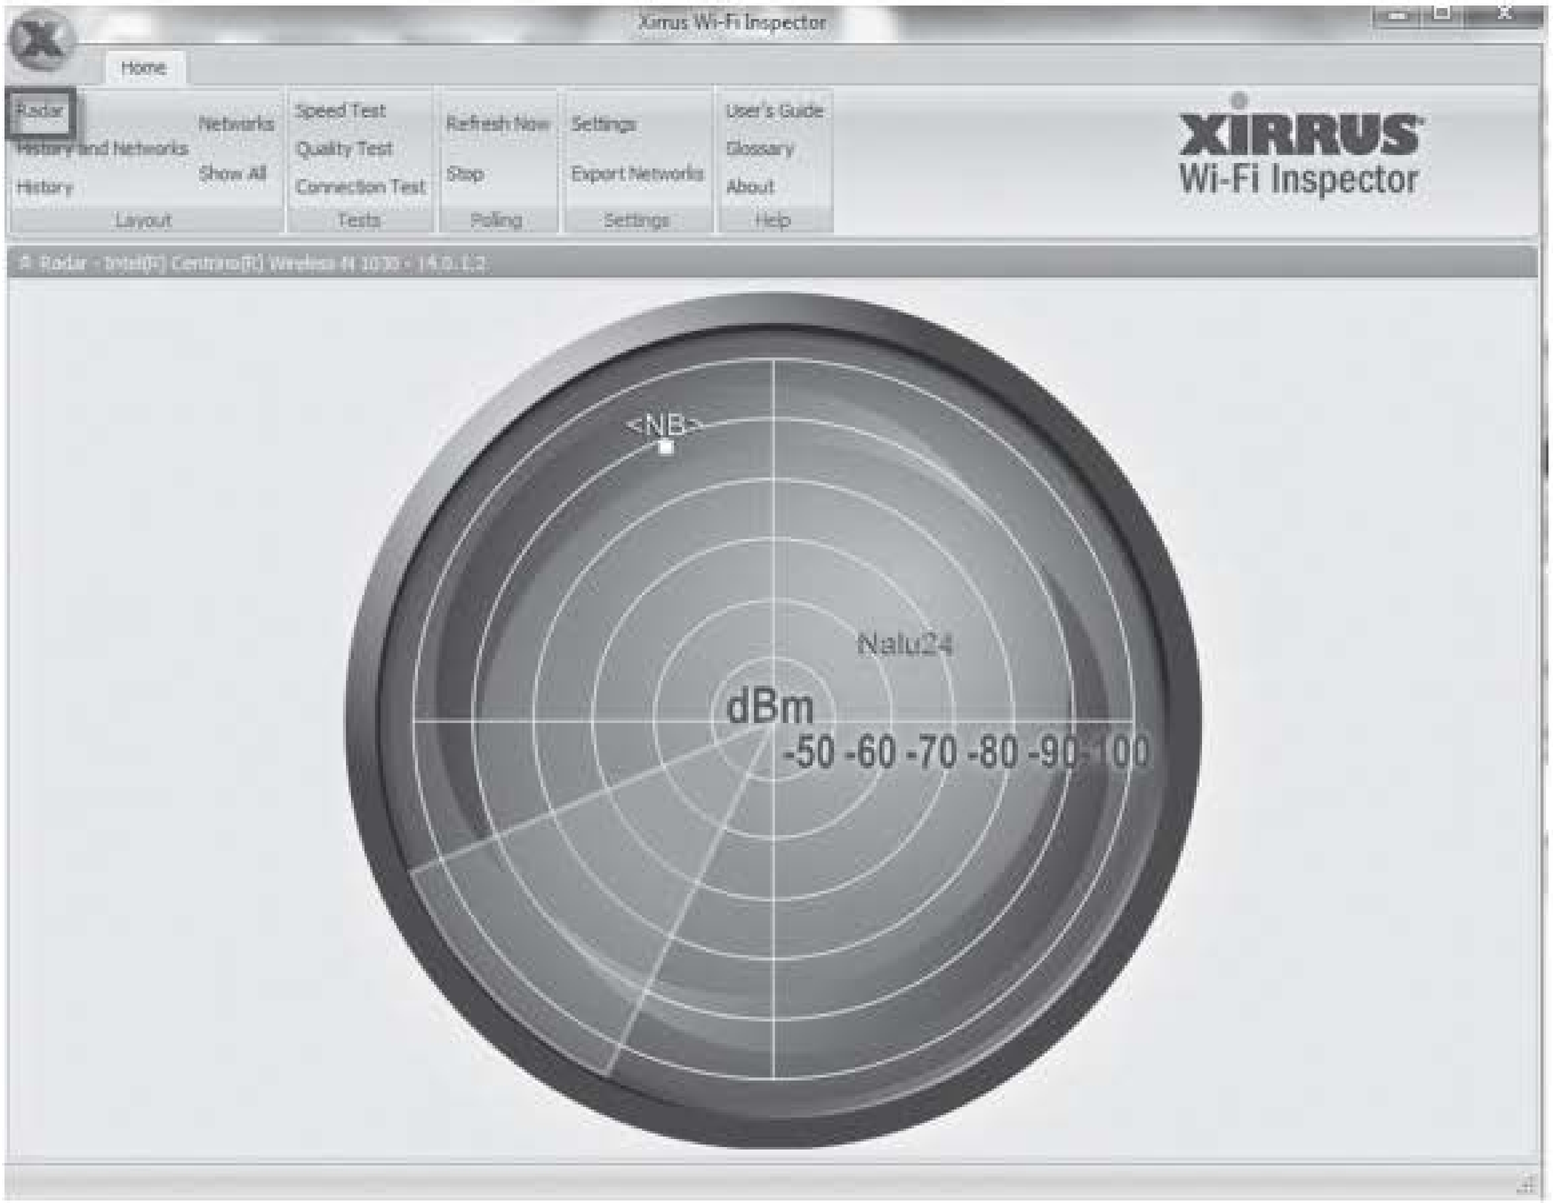

FIGURE 6a-2 Expanded Radar Window

access point's MAC address, in this case, Cisco-Linksys:73:22:51[1]), the channel (6), the signal strength (–65 dBm), and the network mode (802.11n).

In the middle is information about the user's PC. It shows the user's MAC address and configuration information, including the user's IP address, the IP address of the destination server, the IP address of the default gateway (router), and the network's external IP address given to it by the ISP. (This is a home network.) This information does not tell the user about nearby access points, but it can be very useful in assessing connection problems.

On the right is a Connect/Disconnect button. Clicking this button shows a list of potential networks and allows the user's computer to disconnect from the current access point and pick another to connect to. The user can also turn off the computer's wireless adapter.

The Networks Window

The networks window shows detailed information about each of the nearby access points. This is what the user goes to when he or she wants detailed information. The row for the access point to which the user is currently connected is shown in orange highlighted. Wi-Fi Inspector updates the information in the networks window frequently. As Figure 6a-3 shows, the information in this window is detailed.

- *SSID*. The network name.
- Signal Level in either dBm or percentage. Remember that smaller negative dBm numbers indicate higher strength. Next to the number is a colored bar.
  - Green is for signals of –70 dBm and above (–60 dBm, etc.).
  - Yellow is for signals between –71 dBm and –80 dBm.
  - Orange is for signals between –81 dBm and –90 dBm.
  - Red is for –91 dBm and below.

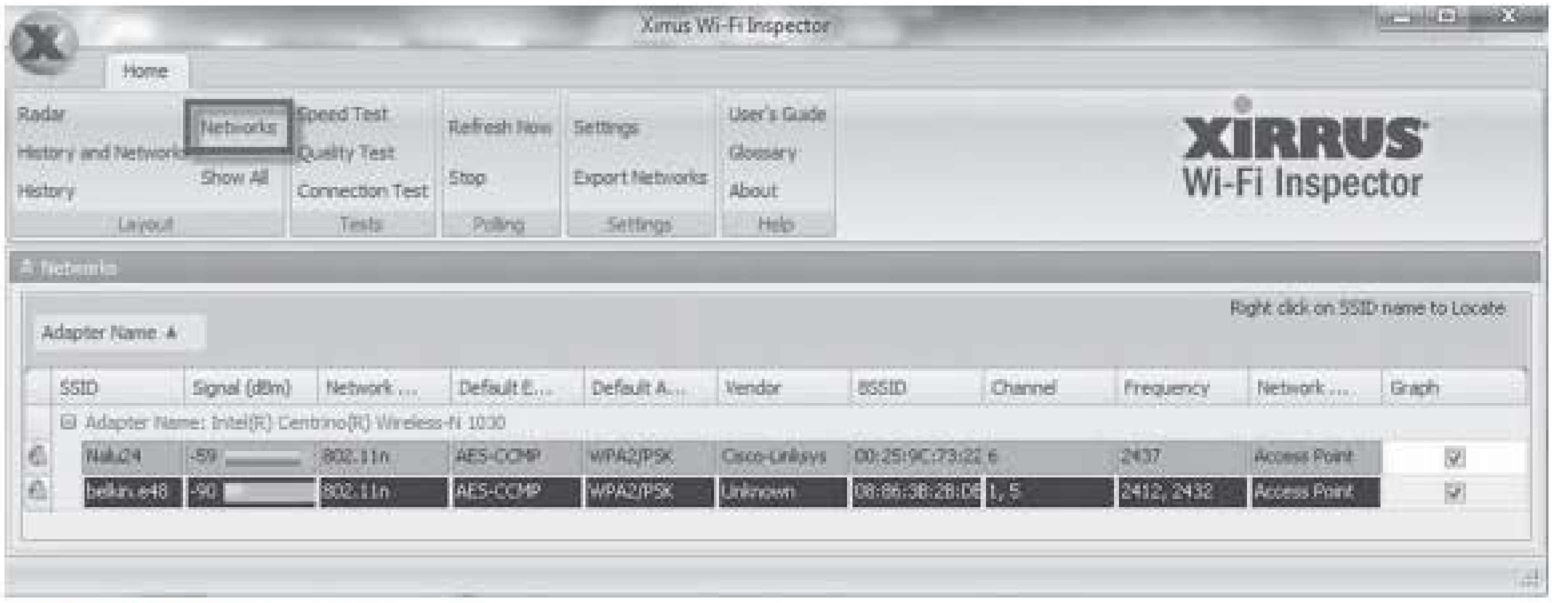

Right click on SSID name to Locate

⊟ Adapter Name: Intel(R) Centrino(R) Wireless-N 1030

SSID	Signal (dBm)	Network ...	Default E...	Default A...	Vendor	BSSID	Channel	Frequency	Network ...	Graph
Naku24	-59	802.11n	AES-CCMP	WPA2/PSK	Cisco-Linksys	00:25:9C:73:22	6	2437	Access Point	☑
belkin.e48	-90	802.11n	AES-CCMP	WPA2/PSK	Unknown	08:86:3B:28:DE	1, 5	2412, 2432	Access Point	☑

FIGURE 6a-3 Networks Window

[1] The first two octets in a MAC address identify the company making the network adapter in the access point. Wi-Fi Inspector converts this information into a humanly readable name.

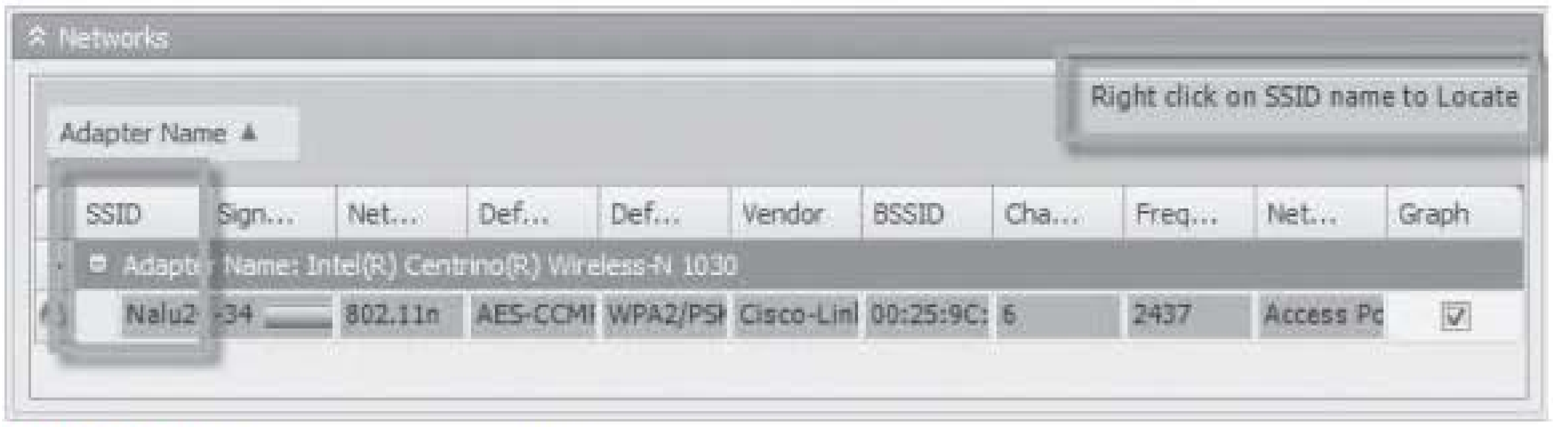

Right click on SSID name to Locate

SSID	Sign...	Net...	Def...	Def...	Vendor	BSSID	Cha...	Freq...	Net...	Graph
Adapter Name: Intel(R) Centrino(R) Wireless-N 1030										
Nalu2	-34	802.11n	AES-CCMI	WPA2/PSI	Cisco-Linl	00:25:9C:	6	2437	Access Pc	☑

FIGURE 6a-4 Locating an Access Point

- *Network Mode*. 802.11g, 802.11n, etc.
- *Default Encryption*. None, WEP, TKIP (in WPA), or AES (802.11i).
- *Default Authentication*. Open (none), WPA/PSK, WPA2/PSK, WPA/802.1X, or WPA2/802.1X.
- *Vendor*. The name of the device manufacturer.
- *BSSID*. The access point's MAC address.
- *Channel*. The channel number.
- *Frequency*. The center frequency of the channel.
- *Network Type*. Access point or ad hoc (no access point).
- *Graph*. This is a checkbox that tells Wi-Fi Inspector to graph the signal level over time (checked) or not to do so (unchecked). In the figure, both are checked, so both will be graphed.

In the figure, the access points are listed in terms of declining signal strength. However, the networks table can be sorted by any column heading. The user merely clicks on the column heading.

Figure 6a-4 zooms in on the networks window. In the upper right, there are instructions to "Right click on SSID name to Locate." In the section on the radar window, we saw that the window does not give the physical locations of access points. The Locate function under networks addresses this lack of physical location in a limited but interesting way. If you right click on an SSID name such as Nalu24, your computer begins beeping. If you are far away, it will beep slowly. As you approach it, the beeping speed will be increased. Essentially, you are using the network analysis version of a Geiger counter.

Signal History

In the networks window, we saw that the user can check or uncheck whether graphing should be done. The Signal History window shows these graphs. The graphs in Figure 6a-5 show that the signal strength for Nalu24 is uniformly excellent and that the signal strength for Belkin in is uniformly poor. Major fluctuations would indicate serious problems.

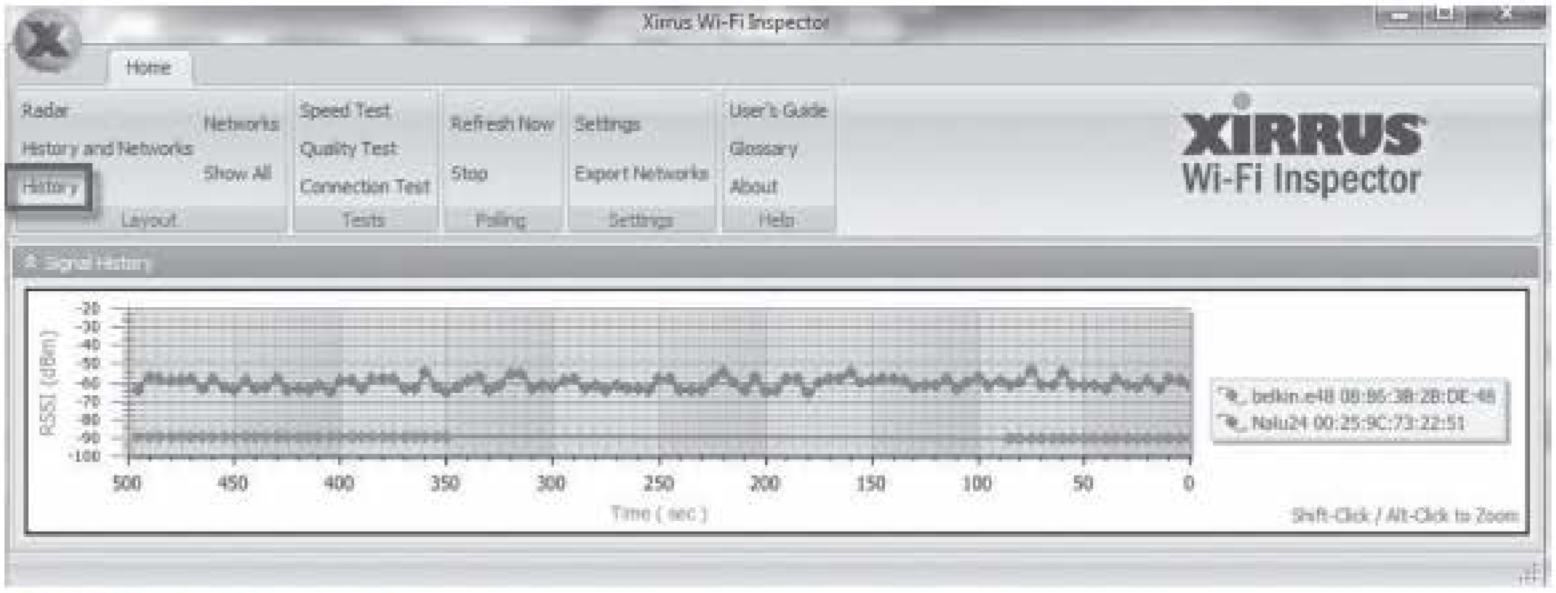

FIGURE 6a-5 Signal History

Other Groups on the Ribbon

The Layout group on the ribbon is the most-used feature of the Xirrus Wi-Fi Inspector.

Help Group The Help group provides a user's guide to explain the program's detailed functionality. There is also a helpful glossary of terms.

Settings Group The Settings group allows the user to adjust many settings, for example, expressing RSSI in percentage terms instead of in terms of dBm.

Tests Group The windows in Wi-Fi Inspector provide information visually. The Tests group allows the user to conduct more detailed tests. These tests are good for troubleshooting.

TESTS

As just noted, the Tests group actively tests the quality of your service. The Tests group performs three important tests.

Connection Test

The connection test shows how well you are connected to the outside world and to critical internal devices. Figure 6a-6 shows the results of a connection test. It shows that Wi-Fi Inspector uses ping to test latency to your DNS server, default gateway (router), and a host on the Internet (Internet Reachable). It also does a DNS lookup, in this case for www.google.com.

The test shows that the user has low latency for the default router and an Internet host. It also shows that the DNS lookup was successful. In color, these are shown in green, with the word *Pass*. However, there is relatively high latency to the user's DNS server (152 ms). This is indicated by a yellow bar with the text *Warning: high latency*. However, the latency is not very high. This connection looks good.

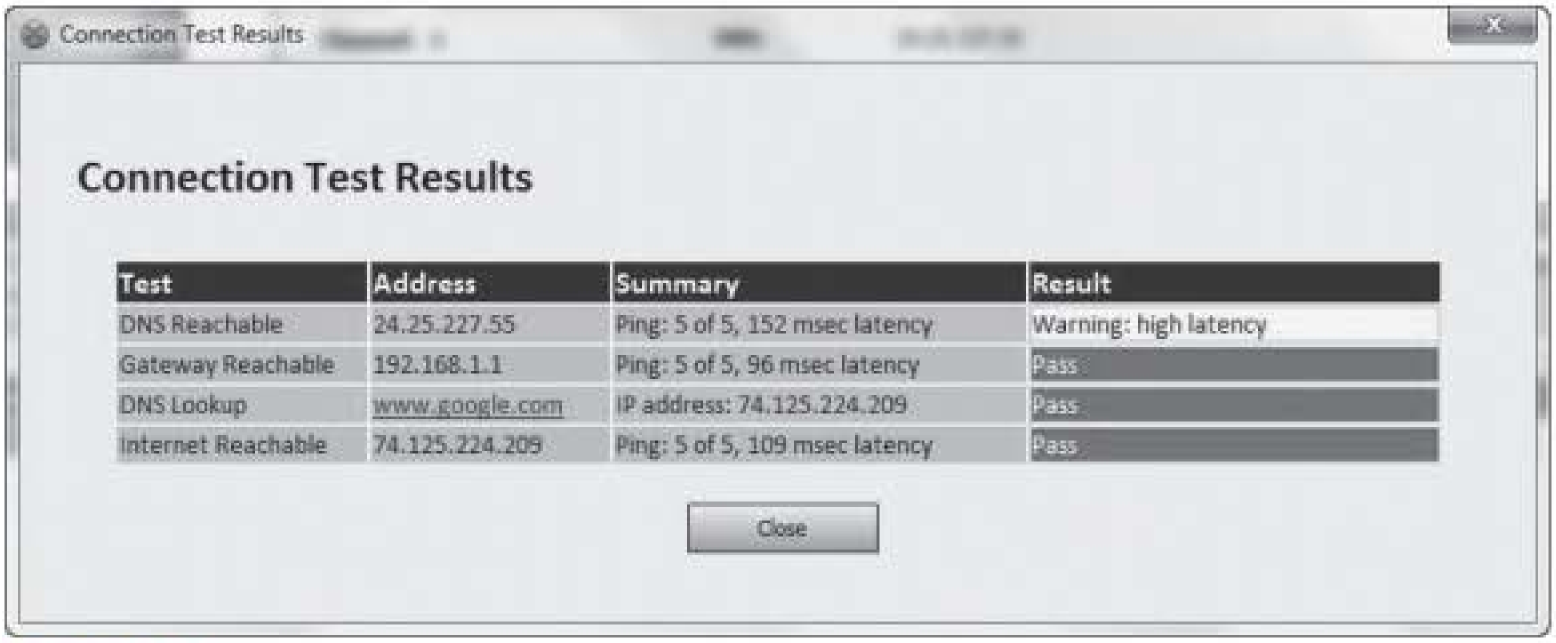

Connection Test Results

Test	Address	Summary	Result
DNS Reachable	24.25.227.55	Ping: 5 of 5, 152 msec latency	Warning: high latency
Gateway Reachable	192.168.1.1	Ping: 5 of 5, 96 msec latency	Pass
DNS Lookup	www.google.com	IP address: 74.125.224.209	Pass
Internet Reachable	74.125.224.209	Ping: 5 of 5, 109 msec latency	Pass

FIGURE 6a-6 Connection Test

Speed Test

The speed test takes the user to speedtest.net. Figure 6a-7 shows a test in which there was a download speed of 14 Mbps and an upload speed of just under 1 Mbps. These are reasonable numbers.

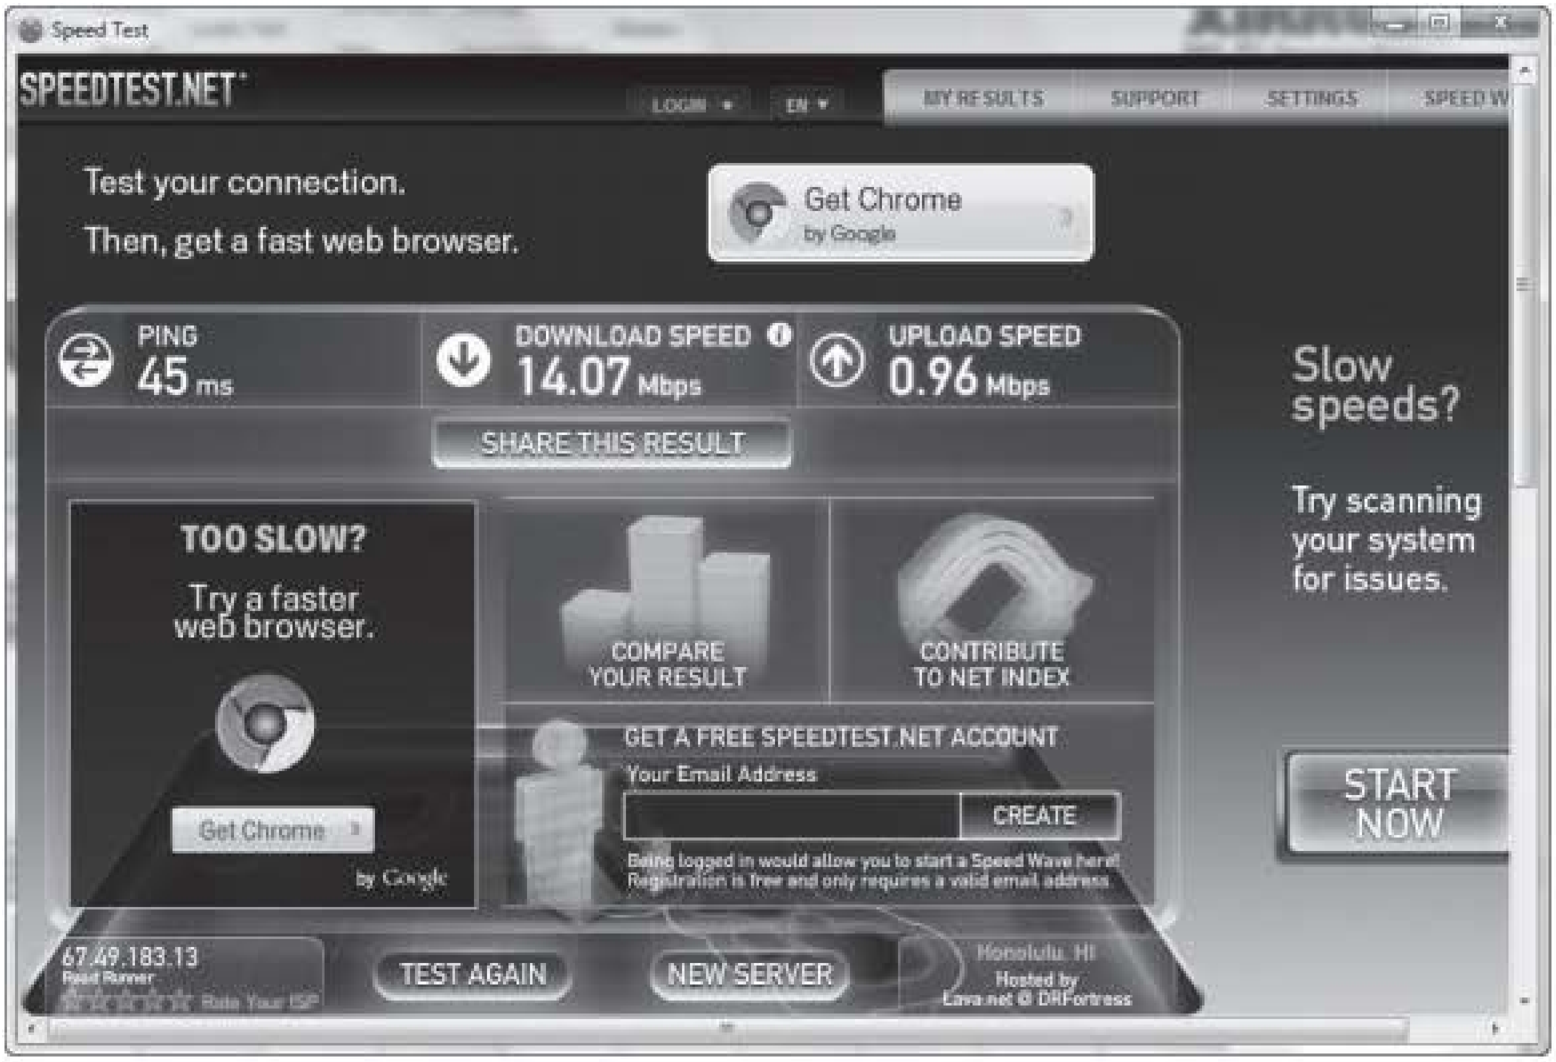

Test your connection.
Then, get a fast web browser.

PING 45 ms

DOWNLOAD SPEED 14.07 Mbps

UPLOAD SPEED 0.96 Mbps

Try a faster web browser.

GET A FREE SPEEDTEST.NET ACCOUNT
Your Email Address

Being logged in would allow you to start a Speed Wave here!
Registration is free and only requires a valid email address

Slow speeds?

Try scanning your system for issues.

67.49.183.13
Road Runner

Honolulu, HI
Hosted by
Lava.net @ DRFortress

FIGURE 6a-7 Speed Test in Wi-Fi Inspector

Quality Test

Figure 6a-8 shows results from the quality test, which takes you to pingtest.net. The results give the user's quality level a B. However, the box on the left notes that the connection should be fine for anything but gaming.

- The ping (latency) averaged 84 ms, which is a little high for games. The server is less than 50 miles away. Connecting to a more distant server would increase latency.
- Jitter, which is variation in latency from packet to packet is 24 ms. This can affect voice and video, for which jitter can result in jittery voice or video. Again, the number is fairly good.
- There was zero packet loss. The connection appears to be reliable.
- There is a MOS score of 4.33. This is a traditional subjective indicator of voice call quality. A MOS score of 5 indicates toll-call quality on the telephone system. A MOS of 4.33 is quite good.

One caveat is that pingtest.net is a bit "grabby." It tries to sell you its tools and is slightly aggressive. In addition, the site uses Java, which you may have to download. You may also have to give a firewall exception to this Java program.

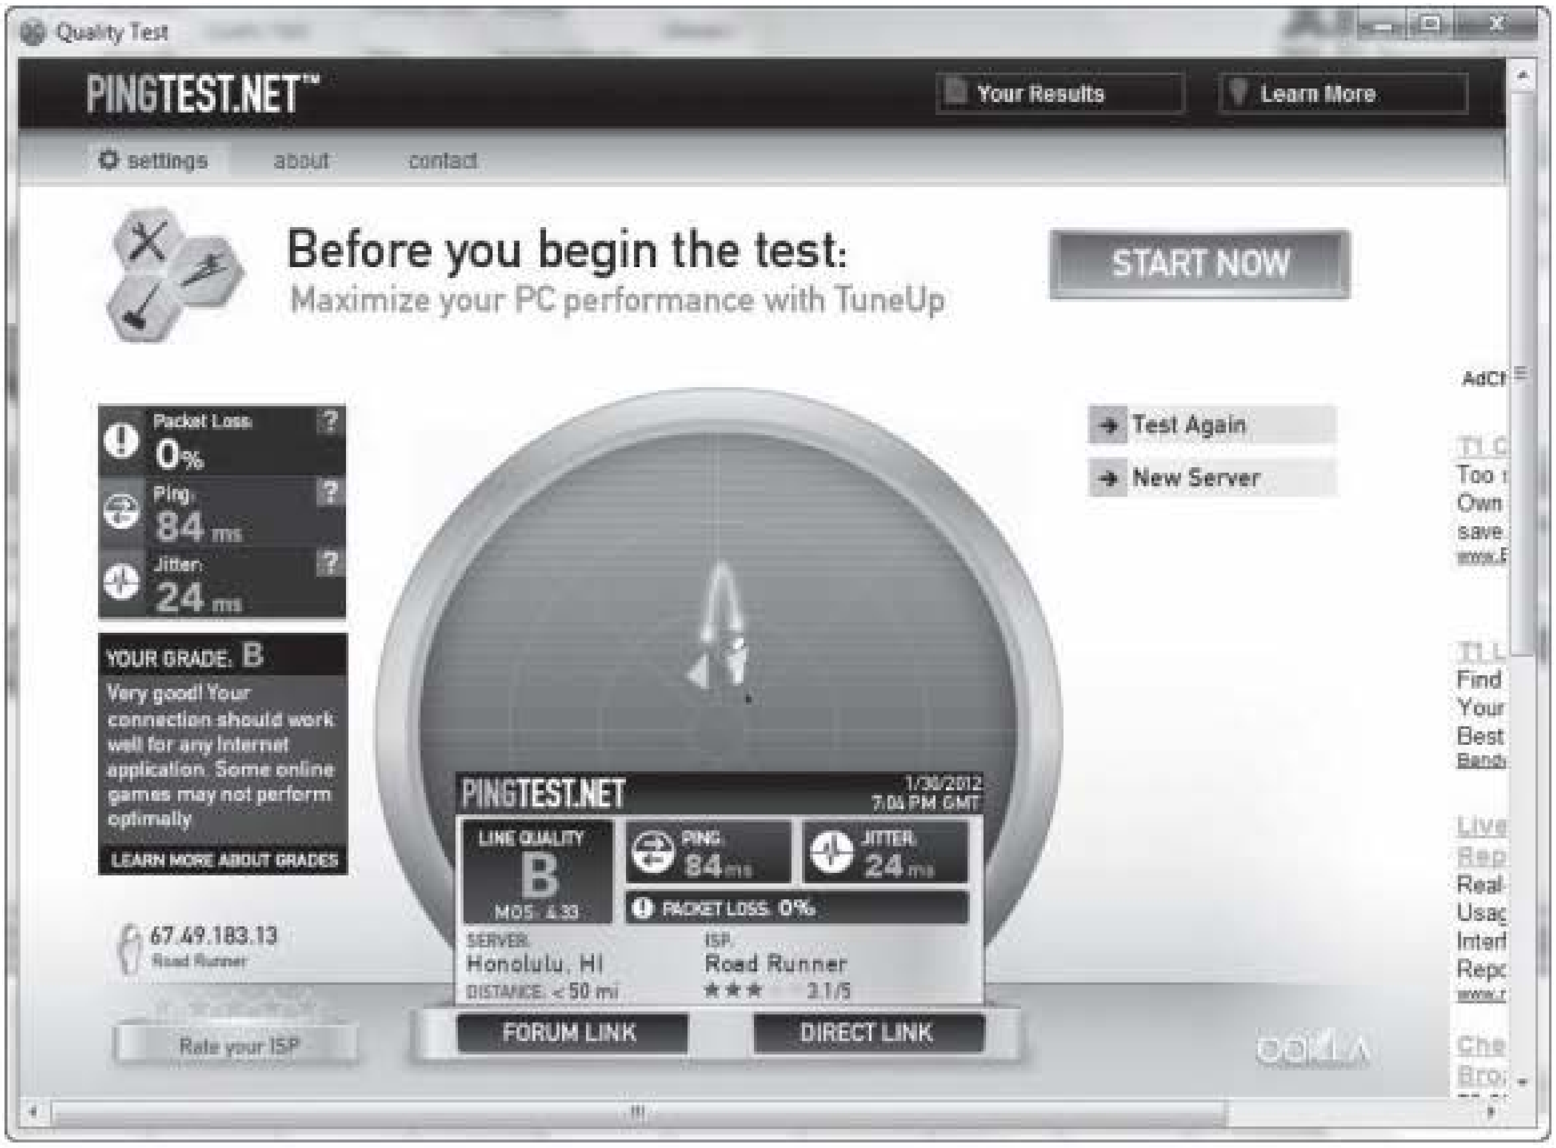

Packet Loss: 0%
Ping: 84 ms
Jitter: 24 ms

YOUR GRADE: B

Very good! Your connection should work well for any Internet application. Some online games may not perform optimally

67.49.183.13
Road Runner

PINGTEST.NET 1/30/2012 7:04 PM GMT

LINE QUALITY B MOS: 4.33
PING: 84 ms
JITTER: 24 ms
PACKET LOSS: 0%
SERVER: Honolulu, HI
DISTANCE: < 50 mi
ISP: Road Runner ★★★ 3.1/5

FIGURE 6a-8 Quality Test

HANDS-ON EXERCISES

Questions

1. Why is the radar window's image of a radar scope misleading?
2. How would you locate an access point despite the limitations of the radar window? This will take one to four paragraphs.
3. There is a value of –44 dBm for signal strength. How good is this?
4. How can you sort the networks window?
5. What information does the Connection Test give you?
6. What information does the Speed Test give you?
7. What information does the Quality Test give you?

Activity

Select a building. Go to at least ten locations. At each location, record the information in the networks window. Also, do a connection and speed test. Write a brief report describing what you learned about Wi-Fi service in the building, referring to the data you collected.

This page intentionally left blank

Chapter **7**

Wireless LANs II

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Explain 802.11i Wi-Fi security.
- Explain why 802.11i security is not enough for WLANs.
- Discuss 802.11 WLAN management.
- Work with decibel representations of power ratios (if you read the box on decibels).
- Compare peer-to-peer local wireless technologies that will be important for the Internet of Things, including Bluetooth.

CHILD'S PLAY[1]

A young girl sat at a computer and connected to the local Wi-Fi network. Like many public wireless networks, this one was "open," meaning that it offered no security. Although a regular computer user, Betsy Davies was only seven years old and not a computer genius. She did not even know how to do what she intended

[1] Nicola, "Hidemyass! Experiment: 7-Year-Old Girl Hacks Public Wi-Fi Network in Less Than 11 Minutes," HideMyAss.com, January 22, 2015, http://blog.hidemyass.com/2015/01/22/hidemyass-experiment-7-year-old-girl-hacks-public-wi-fi-in-less-than-11-minutes/; Ben Rossi, "How a 7-Year-Old Girl Hacked a Public Wi-Fi Network in 10 Minutes," www.informationage.com, January 21, 2015, http://www.information-age.com/how-7-year-old-girl-hacked-public-wi-fi-network-10-minutes-123458891/; Victoria Woollaston, "Hacking Wi-Fi Is Child's Play! 7-Year-Old Shows How Easy It Is to Break Into a Public Network in Less Than 11 MINUTES," DailyMail.com, http://www.dailymail.co.uk/sciencetech/.

to do that day—hack the connection of a nearby computer. This would allow her to eavesdrop on all the traffic sent between the victim and the access point. To learn how, she read a brief tutorial. She then quickly hacked her target connection. It took her 10 minutes and 54 seconds, including reading the tutorial. It had literally been child's play.

This incident did not take place in a coffee shop or other public hot spot. It was an experiment, done with the permission of her parents. However, there was nothing remarkable about the experimental situation. Ms. Bailey demonstrated how easy it is to hack a connection in the many public networks that many people frequently use. Later in this chapter, we will see the type of attack she used. It was a man-in-the middle attack using an evil twin access point. The experiment was conducted by a vendor that offered a virtual private network (VPN) service, HideMyAss.com. We will see later in this chapter how this would have prevented the hack.

In one survey, 59% of people in Britain used unsecure Wi-Fi hotspots in 2015. One in five did so weekly. Among unsecure hot spot users, 19% did online banking, and 31% sent e-mails and documents. In the United States, 87% of people surveyed had used a public hot spot.[2] More than 60% believed that they were protected while using a public access point.[3] Seventeen percent believed that the Wi-Fi supplier protected them; the same percent believed the website did. The simple reality is that your signal spreads out like a sphere when you transmit, reaching everyone nearby. Without encryption and other protections, everything you send is electronically visible to everyone nearby.

You do not have to be in an unsecure Wi-Fi hot spot to have your connection to the Internet hacked. You can be sitting in your office at work. Figure 7-1 illustrates a typical organizational site. It has a border firewall that scrutinizes traffic going into

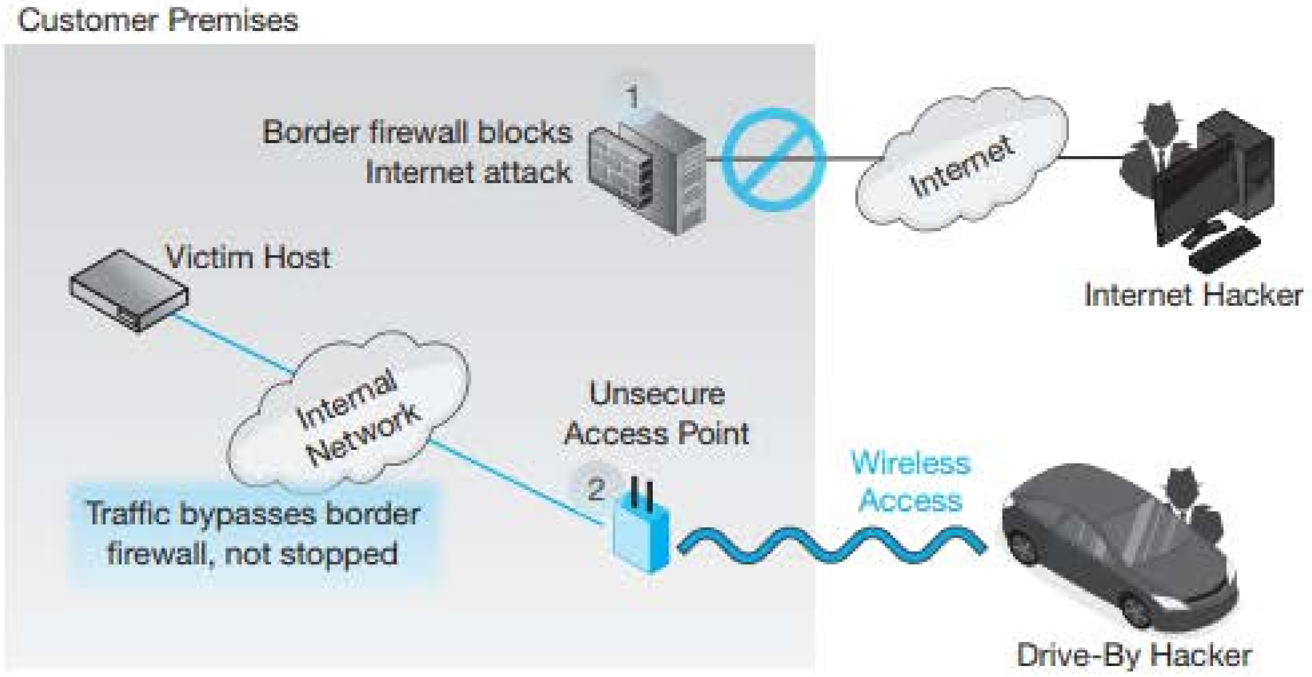

FIGURE 7-1 Drive-By Hacking

[2] Michael Covington, "Free Wi-Fi and the Dangers of Mobile Man-in-the-Middle Attacks," betanews.com. 2015, https://betanews.com/2016/10/08/free-wi-fi-mobile-man-in-the-middle-attacks/.

[3] Ibid.

and out of the site. Within the site, clients connect to the internal network through Wi-Fi access points. Their communication is not filtered by the border firewall because they are treated as being inside its protection.

The figure also shows a **drive-by hacker** located outside the corporate premises. He or she connects to an unsecure access point within the site.[4] If the attempt is successful, then the attacker can communicate with any hosts within the site—without going through the border firewall. The attacker can send attack packets to any host and will be able to intercept at least some conversations within the customer premises.

Companies may mistakenly believe that someone outside their walls will be too far away to communicate with internal access points. However, drive-by hackers use highly directional antennas that allow them to send very strong signals and to receive signals that would be too weak to hear with normal Wi-Fi equipment. Many use Pringles cans.

Test Your Understanding

1. a) Do public hot spots protect your transmissions? b) What type of attack did Ms. Davies use? c) How long did it take her to hack the connection, including reading the tutorial? d) How can a drive-by hacker defeat a site's border firewall?

802.11i WLAN SECURITY

802.11i

Realizing the danger of drive-by hackers, the 802.11 Working Group created the **802.11i** standard. Figure 7-2 shows that 802.11i provides cryptographic protection

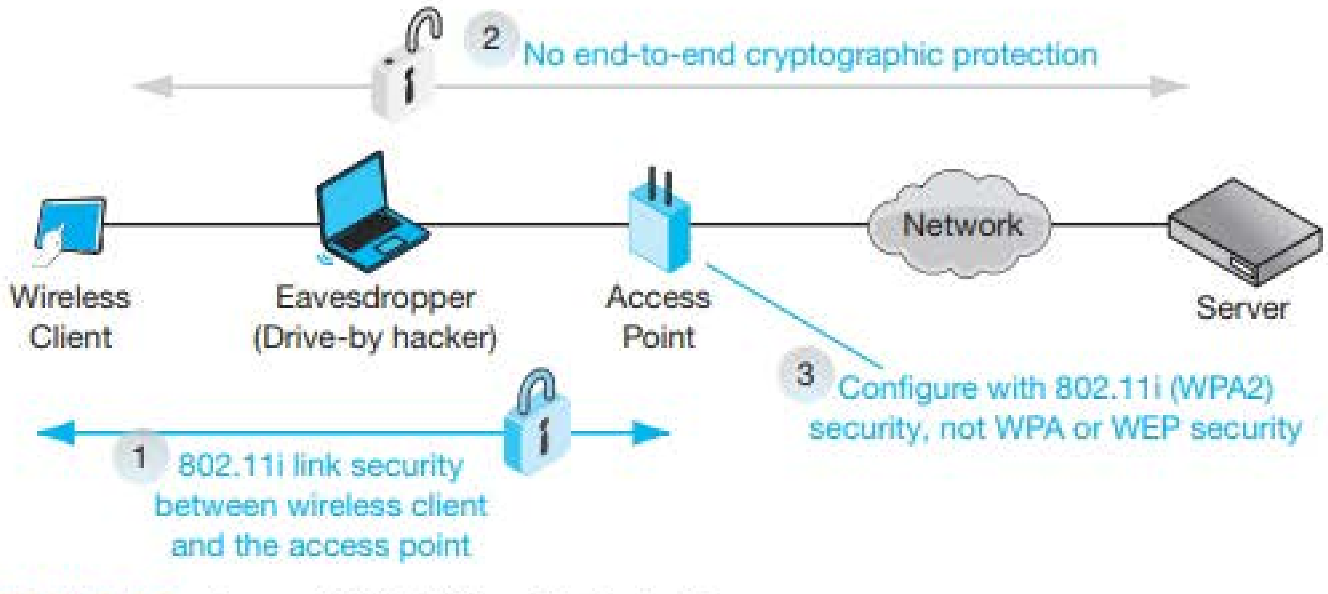
2 No end-to-end cryptographic protection
Wireless Client
Eavesdropper (Drive-by hacker)
Access Point
Network
Server
3 Configure with 802.11i (WPA2) security, not WPA or WEP security
1 802.11i link security between wireless client and the access point

FIGURE 7-2 Scope of 802.11i Security Protection

[4] Merely collecting wireless transmissions to determine such things as SSID, signal strength, and channel is not illegal. This practice, although called war driving, is built into every Wi-Fi program. It cannot be illegal because you need this information to connect to an access point. Of course, a subsequent attempt to connect to an access point without authorization *is* illegal.

between the wireless access point and the wireless host. This protection includes initial authentication plus message-by-message confidentiality, integrity, and authentication (CIA). A drive-by hacker cannot read traffic (confidentiality), modify traffic (integrity), or connect to the access point to send traffic (authentication).[5] Hot spot access points should also secure local communication with 802.11i security. Unfortunately, this security is not mandatory. In fact, because it involves authentication, many hot spot owners avoid it because this makes the access point harder to use.

Note in the figure that 802.11i protection only provides **link security** on the link between the wireless client and the wireless access point. It does not provide **end-to-end security** all the way between the wireless client and the server on the wired LAN (or a server on the Internet). The 802.11i standard has a very limited objective—to protect wireless transmission between the access point and the wireless client host.

The protection provided by 802.11i only extends between the wireless access point and the wireless client host.

Although its physical scope is limited, 802.11i protects transmissions within its scope very well. For example, the standard uses the Advanced Encryption Standard (AES) for confidentiality. It also uses strong standards for all other aspects of cryptology.

Historically, the 802.11i standard was the third standard created to protect communication between wireless clients and access points in 802.11 WLANs. The original standard was **wired equivalent privacy (WEP)**. The 802.11 Working Group created WEP as part of the original 802.11 standard in 1997. WEP was deeply flawed. As a stopgap measure, the Wi-Fi Alliance created an interim security standard based on an early draft of 802.11i but using much weaker standards for cryptographic protections. The Wi-Fi Alliance called their interim standard **Wireless Protected Access (WPA)**.

Today, there is no reason to use WPA because 802.11i is superior, and using WEP is malpractice at best. However, many wireless access points and wireless routers continue to offer WEP and WPA. To add to the confusion, the Wi-Fi Alliance calls the 802.11i standard **WPA2**, and many wireless access points and wireless routers still use this terminology. All access points and wireless clients today support WPA2 at no extra cost. The only choice today should be to use 802.11i/WPA2.

The choice today should be to use 802.11i/WPA2.

[5] Some people recommend further security protections, such as turning off the periodic broadcasting of the access point's SSID. Users need to know this SSID to use an access point. However, the SSID is transmitted in the clear (without encryption) in *every frame header*. Hacker software reads it effortlessly. Another common recommendation is to accept only computers whose wireless network interface cards have preapproved EUI-48 addresses. Again, however, the EUI-48 address is also transmitted in the clear in every packet, and attackers can easily read and spoof one of these addresses. Overall, these measures take a great deal of work, and they are easily pushed aside by readily available hacking software. They might make sense if you are only concerned about a home network and unsophisticated but nosy neighbors, but turning on 802.11i protection is easier, and it provides security automatically without additional rabbit's-foot gambits.

Test Your Understanding

2. a) What cryptographic protections does 802.11i provide? b) How is this protection limited? c) Distinguish between link security and end-to-end security. d) What does the Wi-Fi Alliance call 802.11i? e) When offered the choice when you are configuring a wireless access point, which WLAN security standard should you choose?

802.11i Stages

The 802.11i standard provides a broad spectrum of security protections. At the beginning of a session between a client and an access point, the two parties exchange information. This normally includes **initial authentication**, which is distinct from ongoing message-by-message authentication that takes place after the initial handshaking stage. In initial authentication, the wireless client is the supplicant. It must prove its identity to the access point before the access point will allow the client to connect.

When the 802.11 Working Group created the 802.11i standard, it realized that different initial authentication methods would be needed in homes and large enterprises. These two initial authentication methods are shown in Figure 7-3. Note that whatever initial authentication mode is used, ongoing communication has the same very strong protections, with message-by-message confidentiality, integrity, and authentication. These ongoing protections are extremely strong.

Figure 7-4 shows that these two initial authentication modes are designed for very different environments. **802.1X initial authentication mode** was created for corporations with many access points. It is extremely strong but complex to implement. Using 802.1X for initial authentication would be overkill in residences. To use it, you would have to have a separate authentication server in addition to the other devices in your home! The Wi-Fi Alliance has rightly dubbed this **enterprise mode**.

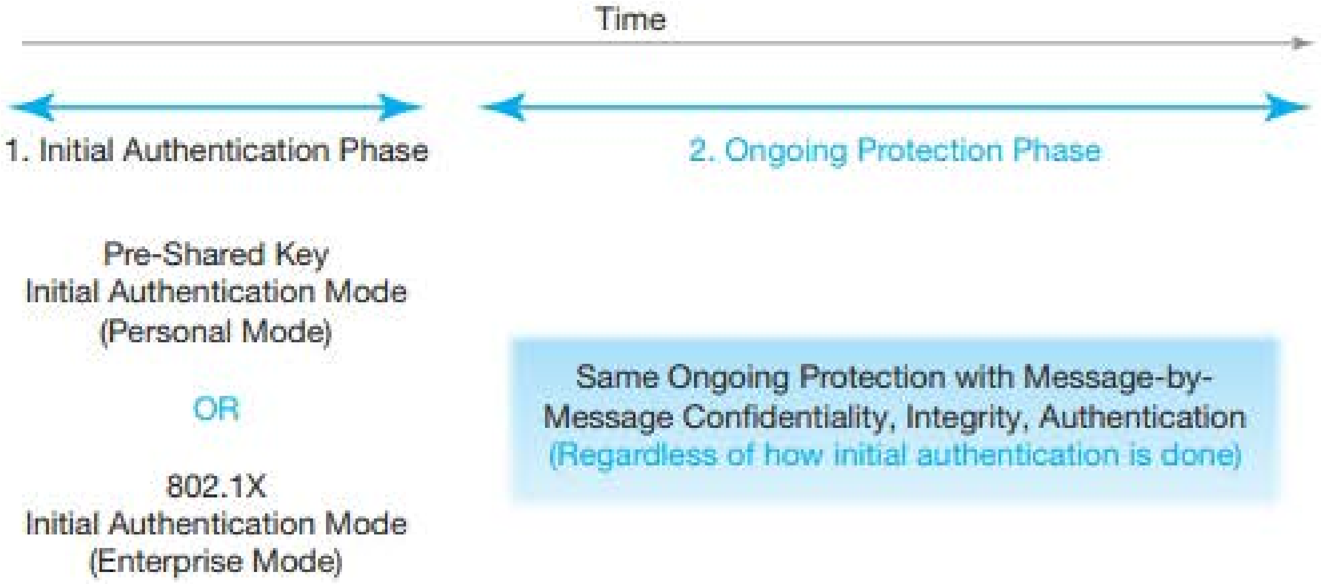
Time
1. Initial Authentication Phase
2. Ongoing Protection Phase
Pre-Shared Key
Initial Authentication Mode
(Personal Mode)
OR
802.1X
Initial Authentication Mode
(Enterprise Mode)
Same Ongoing Protection with Message-by-Message Confidentiality, Integrity, Authentication
(Regardless of how initial authentication is done)

FIGURE 7-3 Phases in 802.11i Cryptographic Security Between the Wireless Client and the Access Point

Mode of 802.11i Initial Authentication	Pre-Shared Key Mode	802.1X Mode
Environment	Home, business with single access point	Companies with multiple access points
Uses a central 802.1X authentication server	No	Yes
Authentication Basis	Knowledge of pre-shared key	Credentials on the 802.1X authentication server
Technical Security	Technologically strong, but weak human security can compromise the technological security	Technically extremely strong but can be defeated by rogue access points and evil twin attacks
Operational Threats	Mismanaging the pre-shared key	Rogue access points, evil twin attacks

FIGURE 7-4 802.11i Modes of Initial Authentication

The 802.11 Working Group created the simpler **Pre-Shared Key (PSK) initial authentication** mode for homes with a single access router. PSK mode is also attractive for small businesses with a single access point. PSK initial authentication mode is a bit weaker than 802.1X initial authentication mode, but it is still strong if implemented properly. The Wi-Fi Alliance calls this **personal mode**.

Test Your Understanding

3. a) For what use scenario was 802.11i PSK mode created? b) For what use scenario was 802.11i's 802.1X mode created? c) Does the choice of initial authentication mode change how other phases of 802.11i work?

Pre-Shared Key (PSK) Initial Authentication Mode in 802.11i

Pre-Shared Session Keys Figure 7-5 shows that the access points and wireless hosts need to know the same **pre-shared key (PSK)** for initial authentication.[6] Demonstrating to the access point that the client knows the PSK authenticates the client to the access point. As "pre-shared" suggests, all hosts on the single access point have the same pre-shared key to authenticate themselves. In fact, *anyone* who knows the PSK can authenticate himself or herself to the access point.

Unshared Pairwise Session Keys After authentication using the pre-shared key, the wireless access point gives each authenticated device a new unshared **pairwise session**

[6] The figure shows the PSK being sent by Host X to the wireless access router/access point. In fact, when a host tries to connect to the access point, the access point sends a challenge message, which the host *encrypts* with the PSK. If the encrypted challenge message can be decrypted back to the challenge message via the PSK, the access point knows that the host knows the PSK and should be accepted.

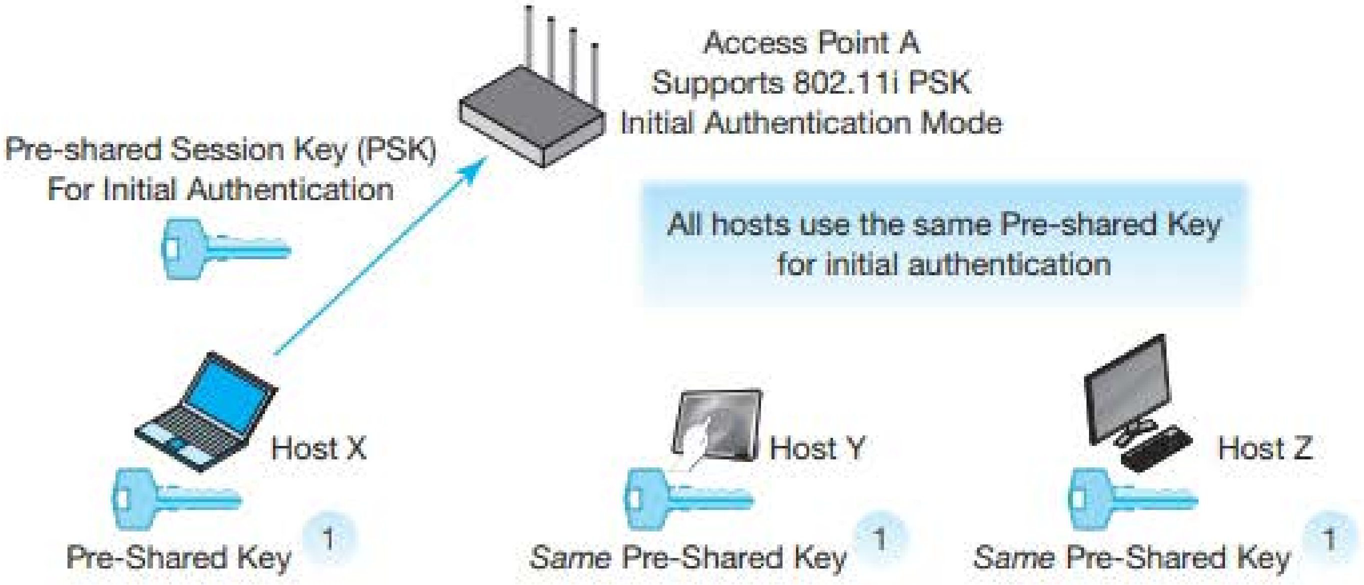
Access Point A
Supports 802.11i PSK
Initial Authentication Mode
Pre-shared Session Key (PSK)
For Initial Authentication
All hosts use the same Pre-shared Key
for initial authentication
Host X
Host Y
Host Z
Pre-Shared Key
1
Same Pre-Shared Key
1
Same Pre-Shared Key
1

FIGURE 7-5 802.11i Pre-Shared Key (PSK) Initial Authentication Mode

key to use while communicating with the access point subsequently. Figure 7-6 shows this second key. It is a **session key** because it will only be used for a single communication session. The next time a client authenticates itself, it will receive a different session key. It is a **pairwise** key in the sense that each client will have a different session key to use with the access point. Each client will use its own pairwise session key to encrypt frames sent to the access point. Other clients, not knowing the unshared pairwise session keys of others, will not be able to read these frames.

Security Threats in 802.11i PSK Mode Although 802.11i PSK mode is technically strong, it faces some threats involving how the PSK is managed. Operational (human) security must be equal to the technical security if a residence or small business is to be safe.

One operational security threat is that someone who is not authorized to use the network will learn the pre-shared key. In a home or very small business, there is the danger that someone, rationalizing that "everyone knows" the pre-shared key, will

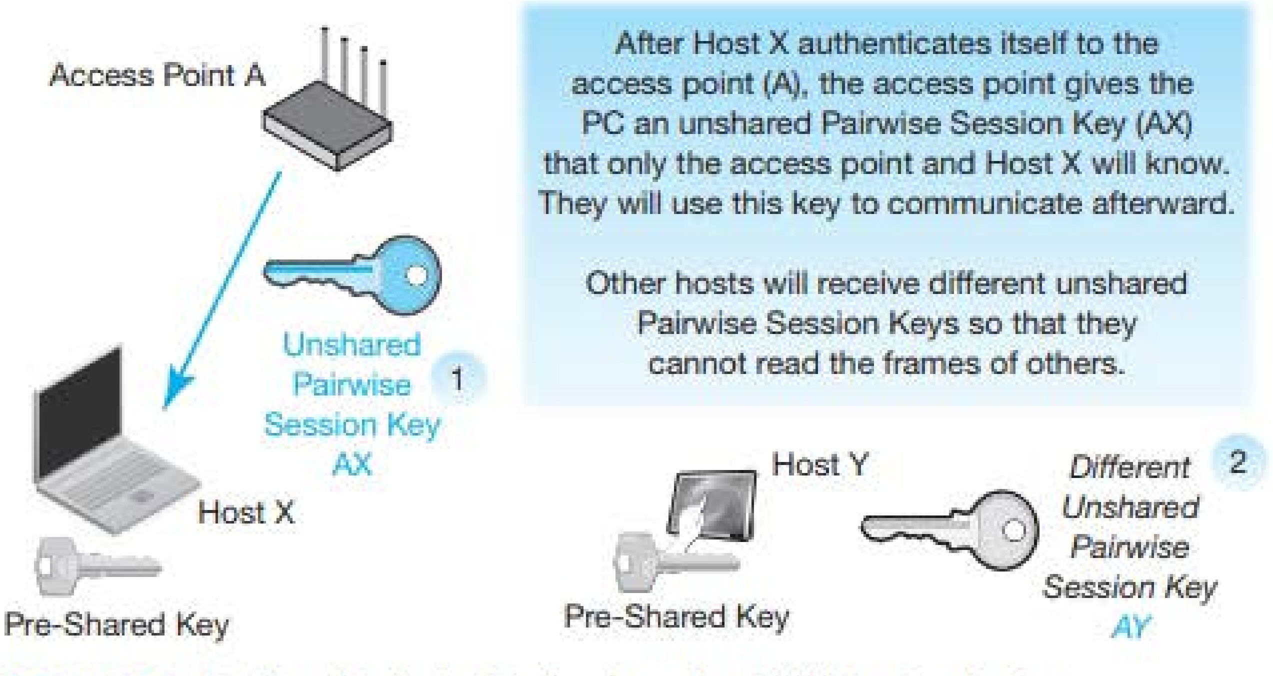
Access Point A
After Host X authenticates itself to the
access point (A), the access point gives the
PC an unshared Pairwise Session Key (AX)
that only the access point and Host X will know.
They will use this key to communicate afterward.
Other hosts will receive different unshared
Pairwise Session Keys so that they
cannot read the frames of others.
Unshared
Pairwise
Session Key
AX
1
Host X
Pre-Shared Key
Host Y
Pre-Shared Key
Different
Unshared
Pairwise
Session Key
AY
2

FIGURE 7-6 Unshared Pairwise Session Key after Initial Authentication

give it to an unauthorized person. Some PSK mode access points at least have a guest account to provide temporary access to outsiders as appropriate.

If a person leaves a company that uses 802.11i PSK mode, it is important to change the pre-shared key. There is no automated way to do this. It must be changed on every device that will use the access point. Given the fact that work is involved, it is all too easy to delay this.

Another danger is that the household or small business will select a weak passphrase. To create the pre-shared key, the household or company creates a long **passphrase**, which is much longer than a password. The client or access point enters this passphrase; the system then automatically generates the 64-bit PSK. The passphrase must be at least 20 characters long to generate a strong pre-shared key. If short passphrases are used, 802.11i in PSK mode can be cracked in seconds.

In 802.11i pre-shared key mode, the passphrase must be at least 20 characters long to generate a strong pre-shared key.

Test Your Understanding

4. a) For what use scenario was 802.11i PSK mode created? b) What must a user know to authenticate his or her device to the access point? c) In what ways is the pairwise session key the user receives after authentication different from the PSK? d) What three operational security threats must PSK users consider? e) Why is this risk probably acceptable for the PSK use scenario? (The answer is not in the text.) f) How long must passphrases be to generate strong pre-shared keys?

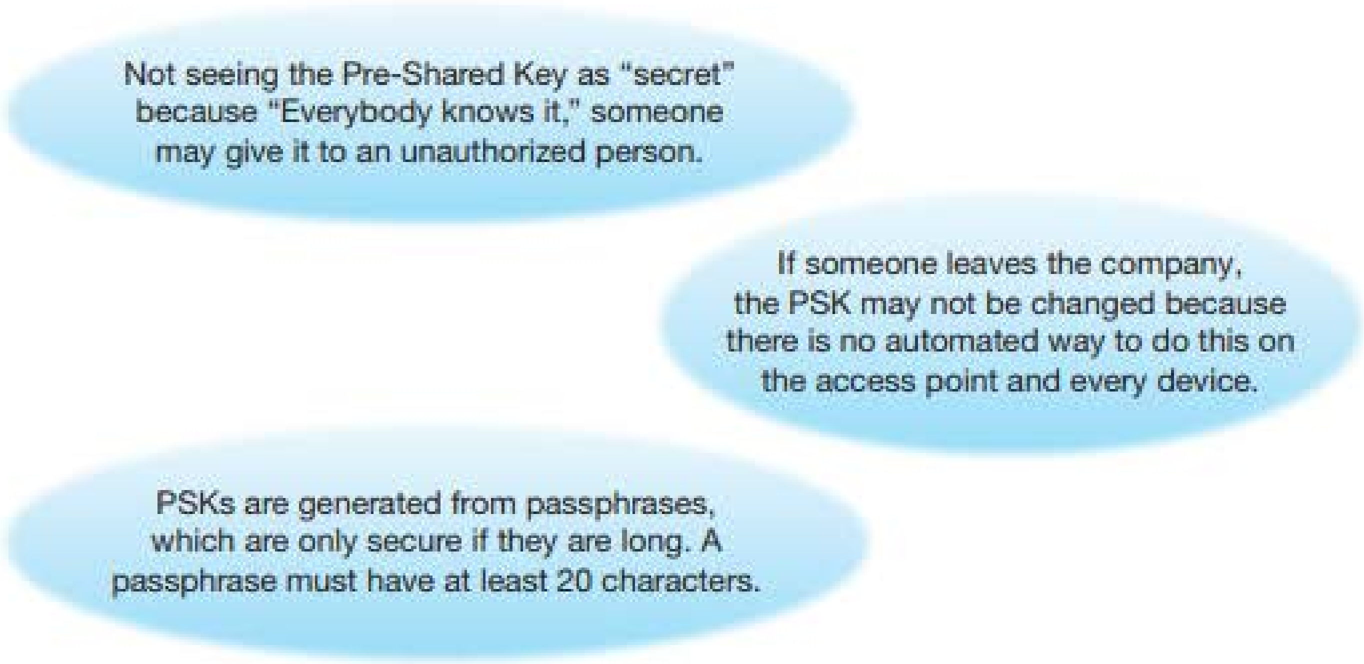
Not seeing the Pre-Shared Key as "secret" because "Everybody knows it," someone may give it to an unauthorized person.
If someone leaves the company, the PSK may not be changed because there is no automated way to do this on the access point and every device.
PSKs are generated from passphrases, which are only secure if they are long. A passphrase must have at least 20 characters.

FIGURE 7-7 Operational Security Threats in Pre-Shared Key Mode (Study Figure)

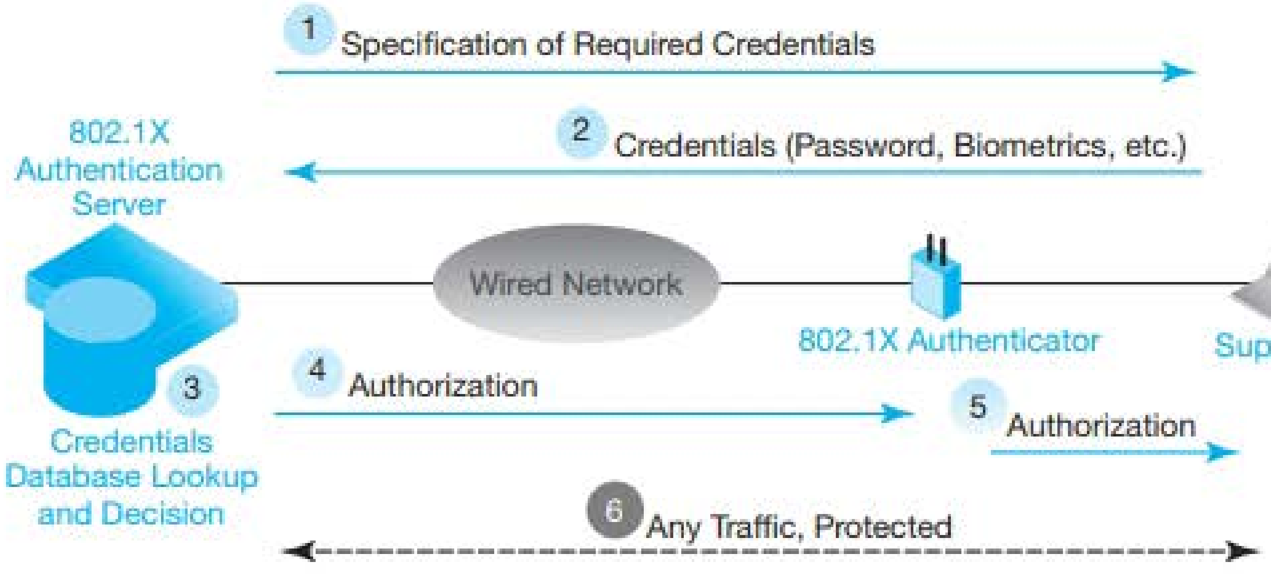
1 Specification of Required Credentials
802.1X Authentication Server
2 Credentials (Password, Biometrics, etc.)
Wired Network
802.1X Authenticator
Supplicant
3
Credentials Database Lookup and Decision
4 Authorization
5 Authorization
6 Any Traffic, Protected

FIGURE 7-8 802.1X Initial Authentication Mode

802.1X Initial Authentication Mode Operation

Again, 802.11i with PSK mode for initial authorization is for homes and for small businesses with a single access point. Large firms with many access points must use a different 802.11i initial authentication mode, **802.1X mode**. (The Wi-Fi Alliance appropriately calls this mode **enterprise mode**.)

The Elements of 802.1X Initial Authentication Mode Figure 7-8 shows that there are three devices involved in 802.1X initial authentication in 802.11i. The wireless client is called the supplicant, of course. However, there is no single verifier. Instead, the verification function is distributed over two devices. The first is the access point, which is the **802.1X authenticator**. The second is a central **802.1X authentication server**. The access point/authenticator is mostly a pass-through device during initial authentication. The real work of authentication is done by the 802.1X authentication server. It has the database on credentials, and its job is to do the heavy work of checking supplicant authentication credentials.

The 802.1X Authentication Process Figure 7-8 also shows the four steps in the 802.1X initial authentication process in 802.1i.[7]

- ***Requirements for Credentials.*** When the supplicant first contacts the access point, the access point authenticator notifies the 802.1X central authentication server. The server sends requirements for credentials to the supplicant. The access point authenticator passes it on to the supplicant.
- ***Provide Credentials.*** The supplicant sends the required credentials to the access point, which passes it on to the authentication server.

[7] In many apartment buildings, a person reaching the front door must buzz your apartment and ask you to open the building's outer door. The buzzer is the authenticator. It buzzes your apartment. You are then the authentication server. You talk to the person and decide whether to let the person in. If you decide to let the person in, you send a signal, and the door opens for the visitor.

- ***Credentials Check***. The 802.1X authentication server receives the credentials and checks them against its authentication database. For example, if the credentials are a username and password, the central authentication checks to see if the password matches the password in the credentials database for the username.
- ***Authorization Message to the Authenticator***. If the authentication succeeds, the 802.1X authentication server sends back an authorization message. This message is not sent to the supplicant like the earlier messages. Instead, it is sent to the access point/802.1X authenticator itself. It tells the access point to accept a connection from the authenticated user.
- ***Authorization Message to the Client***. When it receives the authorization message, the access point authenticator authorizes the connection to the client supplicant and sends an authorization message to the client.
- The client may now send packets to any host on the network.

Test Your Understanding

5. a) Contrast the use scenarios for initial authentication in PSK mode and 802.1X mode. b) Which initial authentication mode or modes of 802.11i authentication use(s) a central authentication server? c) What does the Wi-Fi Alliance call this 802.11i initial authentication mode? d) In 802.1X operation, what device acts as the authenticator in Wi-Fi? e) In 802.1X, which is the verifier?

6. a) What initial authentication mode does 802.11i use? (This is a trick question.) b) Which initial authentication mode is used for message-by-message encryption, authentication, and message integrity? (Another trick question!)

BEYOND 802.11I SECURITY

Again, the 802.11i standard protects communication between the wireless access point and wireless clients. This greatly reduces risks. However, two types of attack can succeed even if a company implements 802.11i security well. These are attacks on rogue access points and evil twin attacks.

Rogue Access Points

The first threat that can defeat 802.11i security is the creation of rogue access points. A **rogue access point** is an unauthorized access point set up within a firm by an employee or department. Rogue access points are dangerous because they are typically configured with no security or poor security. Figure 7-9 shows that even if a firm carefully applies 802.11i to every one of its authorized access points, the presence of a single unsecure rogue access point will give a drive-by hacker access to the firm's

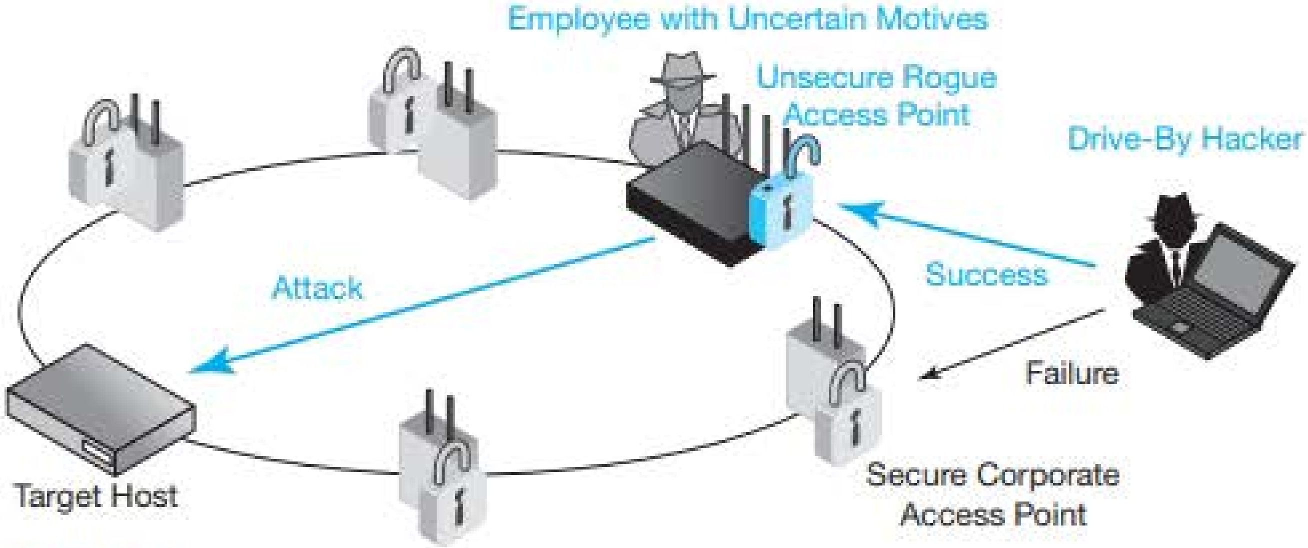
Employee with Uncertain Motives
Unsecure Rogue Access Point
Drive-By Hacker
Attack
Success
Failure
Target Host
Secure Corporate Access Point

FIGURE 7-9 Rogue Access Point

internal network. In other words, a single rogue access point destroys the security that the firm has so laboriously created with 802.11i. In the terminology of the Appendix, this is a weakest link problem. The least secure access point determines the strength of the entire network.

A rogue access point is an unauthorized access point set up within a firm by an employee or department.

The employees who set up rogue access points may not have malicious intent. In many cases, they set up their own access points because they are getting poor Wi-Fi service. However, even nonmalicious employees who set up unauthorized access points can ruin wireless security.

Test Your Understanding

7. a) Who creates a rogue access point? b) Why can they defeat 802.11i security? c) Do employees who set up rogue access points have malicious motives?

Evil Twin Access Points and Virtual Private Networks (VPNs)

The second type of attack that 802.11i will not stop is the ominous-sounding evil twin access point attack. An **evil twin attack** is a man-in-the-middle attack in which the evil twin intercepts traffic passing between a wireless host and a legitimate access point.

An evil twin attack is a man-in-the-middle attack in which the evil twin intercepts traffic passing between a wireless host and a legitimate access point.

Evil Twin Access Points Figure 7-10 illustrates an evil twin access point attack. Normally, the wireless client shown in the figure will associate with its legitimate access point. The two will establish an 802.11i connection between them, communicating via a common encryption key.

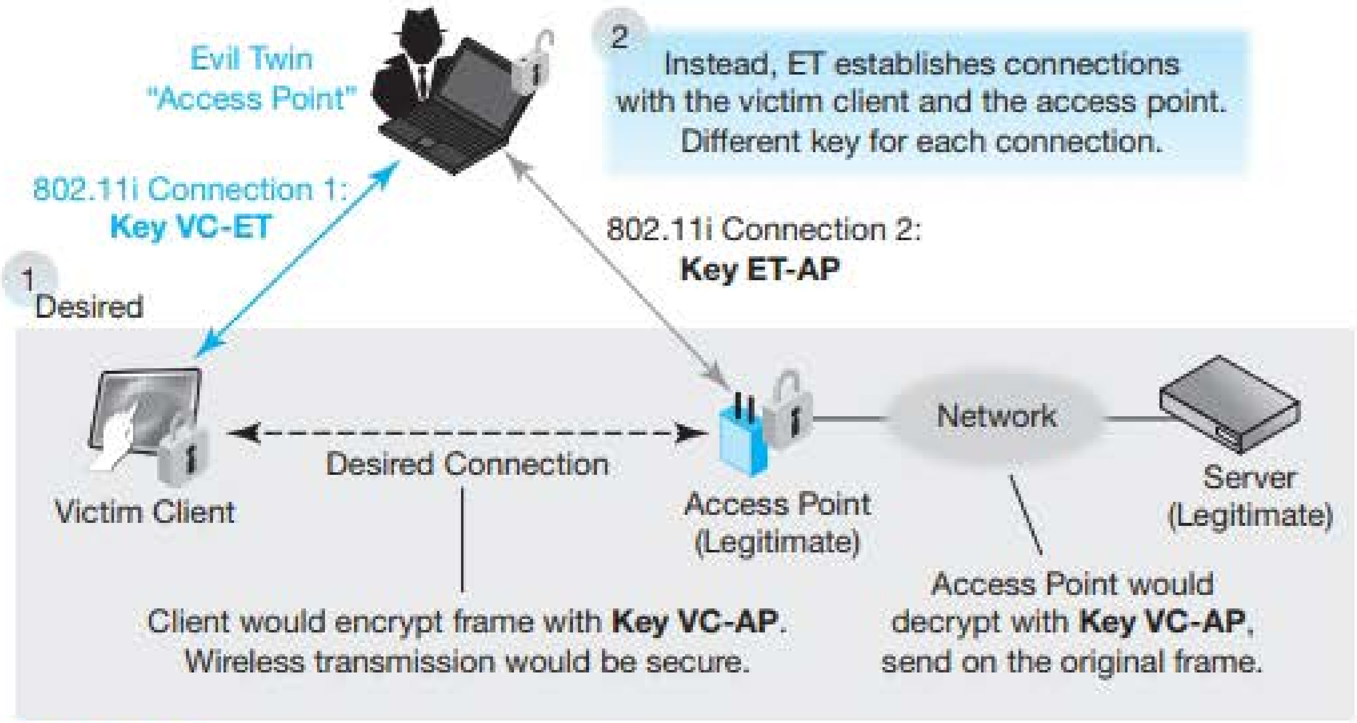
Evil Twin
"Access Point"
2
Instead, ET establishes connections with the victim client and the access point. Different key for each connection.
802.11i Connection 1:
Key VC-ET
802.11i Connection 2:
Key ET-AP
1
Desired
Desired Connection
Victim Client
Access Point
(Legitimate)
Network
Server
(Legitimate)
Client would encrypt frame with Key VC-AP.
Wireless transmission would be secure.
Access Point would decrypt with Key VC-AP, send on the original frame.

FIGURE 7-10 Desired Operation and Evil Twin Connection

An evil twin access point (usually a notebook computer) has software to impersonate a real access point. The evil twin operates at very high power. If the wireless host is configured to choose the highest-power access point it can reach, it will associate with the evil twin access point instead of with the legitimate access point. The evil twin will establish a secure 802.11i connection with the wireless victim client. This is *Security Connection 1*. It will use *Key Client-ET (VC-ET)* for encryption.

An evil twin access point is a notebook computer configured to act like a real access point.

Next, the evil twin associates with the legitimate access point using 802.11i, creating *Security Connection 2*. This connection will use *Key ET-AP* for encryption. The evil twin now has two symmetric session keys—one that it shares with the victim client and one that it shares with the legitimate access point.

Normal Operation Figure 7-11 shows what happens when an evil twin operates normally.

- When the host transmits a frame, the host first encrypts it with key Client-ET. It then transmits the encrypted frame to the evil twin.
- The evil twin decrypts the received frame with key Client-ET. It then reads the message in the clear. Its eavesdropping task is done.
- To continue the deception, the evil twin reencrypts the frame, this time with Key ET-AP. Then it sends the encrypted frame to the legitimate access point, which decrypts it and passes it on.

A man–in-the-middle attack is difficult to detect because it is transparent to both the wireless client and the access point. Both operate as usual. Neither can tell that it is dealing with an impostor.

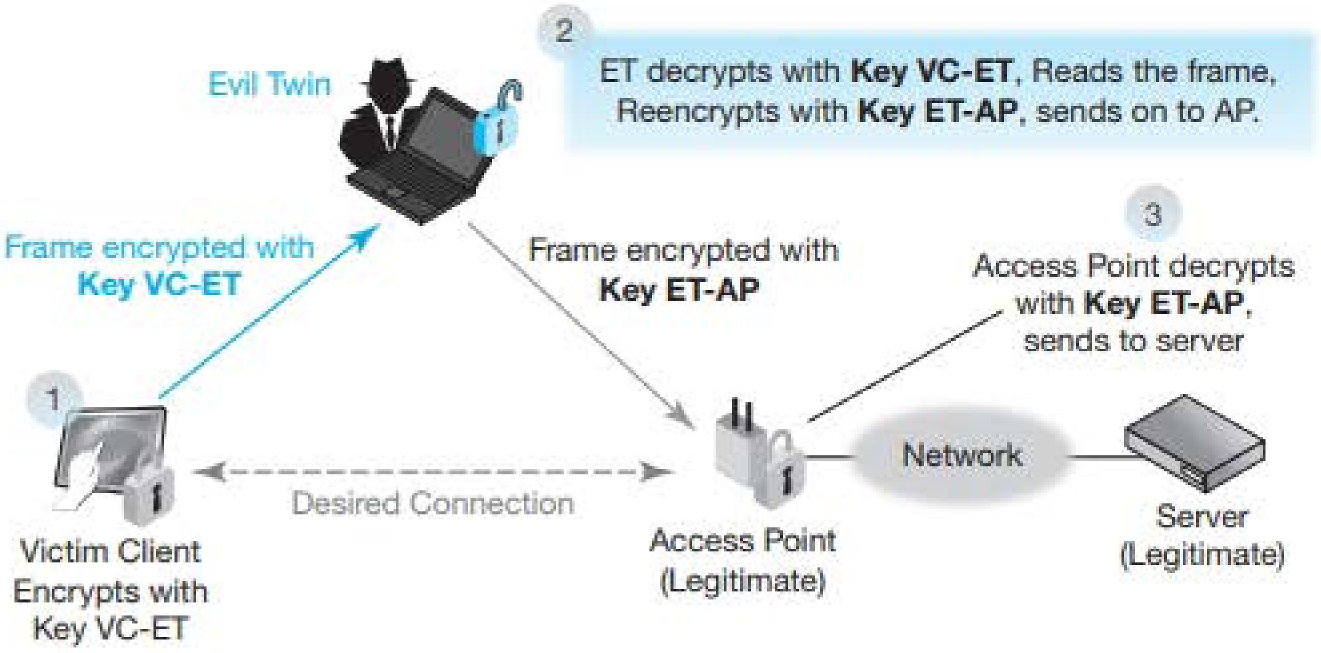
Evil Twin
2
ET decrypts with Key VC-ET, Reads the frame, Reencrypts with Key ET-AP, sends on to AP.
Frame encrypted with Key VC-ET
Frame encrypted with Key ET-AP
3
Access Point decrypts with Key ET-AP, sends to server
1
Desired Connection
Network
Victim Client Encrypts with Key VC-ET
Access Point (Legitimate)
Server (Legitimate)

FIGURE 7-11 Operation with Evil Twin Connections

A man–in-the-middle attack is difficult to detect because it is transparent to both the wireless client and the access point. Both operate as usual. Neither can tell that it is dealing with an impostor.

Using a VPN to Defeat Evil Twins If a client cannot detect that it is being deceived by an evil twin access point attack, how can it protect itself? The answer is that it can take a simple precaution. As Figure 7-12 shows, a client can implement a virtual private network (VPN) between itself and the server it wishes to communicate with. We saw VPNs in Chapter 4. A VPN is simply an encrypted path through an untrusted network. Because the transmission is encrypted, others cannot read it. It is as if the transmission was traveling over its own private network.

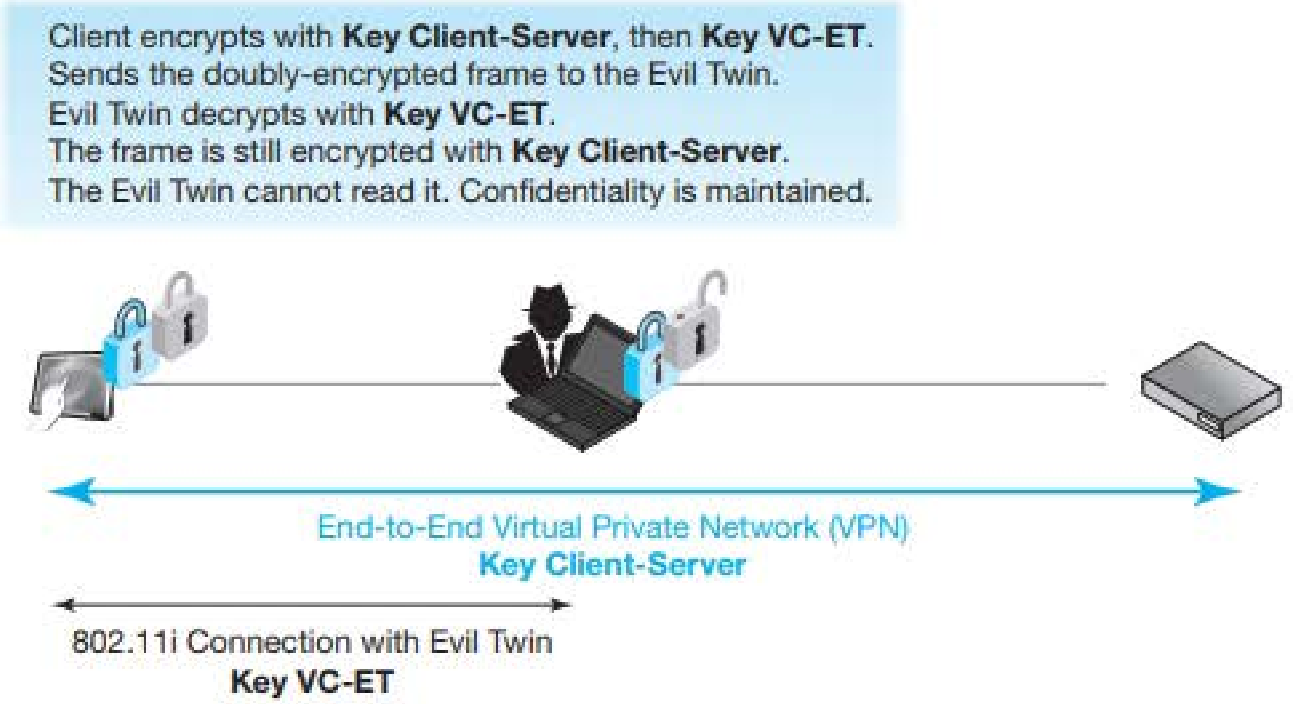
Client encrypts with Key Client-Server, then Key VC-ET.
Sends the doubly-encrypted frame to the Evil Twin.
Evil Twin decrypts with Key VC-ET.
The frame is still encrypted with Key Client-Server.
The Evil Twin cannot read it. Confidentiality is maintained.
End-to-End Virtual Private Network (VPN)
Key Client-Server
802.11i Connection with Evil Twin
Key VC-ET

FIGURE 7-12 Defeating an Evil Twin Attack by Using a Virtual Private Network (VPN)

The evil twin still intercepts traffic. Now, however, intercepting the traffic does it no good. Consider what happens when the client transmits a frame.

- The client first encrypts a frame it with the VPN key, *Key Client-Server*, which it shares with the server. It then encrypts the frame again, this time with the key it shares with the evil twin (*Victim Client-ET*). Now it sends the doubly encrypted frame to the evil twin.
- The evil twin decrypts the frame with the Victim Client-ET key. However, the frame is still encrypted with the VPN key. The ET cannot read the message.

Test Your Understanding

8. a) What kind of physical device is an evil twin access point? b) What does the evil twin do after initial association when the victim client transmits? c) Distinguish between evil twin access points and rogue access points. (The answer is not explicitly in the text.) d) How are VPNs able to defeat evil twin attacks? Explain in detail. e) How can you tell if your client computer has succumbed to an evil twin attack? f) Why is this important?

802.11 WI-FI WIRELESS LAN MANAGEMENT

Until recently, the term WLAN management was almost an oxymoron. Large WLANs were like airports without control towers. Companies knew that they needed tools to centralize WLAN management. Vendors began to provide these tools.

Access Point Placement

The first management issue is where to place access points throughout a building or site. If access points are placed poorly, there will be overloaded access points, dead spots, and crippling interference between access points.

Initial Planning The first step in placing access points is to determine how far signals should travel. This determines the radius of service around each access point.

- If the radius is too great, many hosts will be far from their access points. Hosts far from the access point must drop down to lower transmission speeds, and their frames will take longer to send and receive. This will reduce the access point's effective capacity. Also, a large circle may contain too many users to handle.
- If the radius is too small, however, the firm will need many more access points to cover the area to be served. Having access points too close together will also increase co-channel interference if it is present.

Once an appropriate radius is selected (say, 10 meters), the company gets out its architecture drawings and begins to lay out 10-meter circles that cover all points in the building, as shown in Figure 7-13. Where there are thick walls, filing cabinets, or other obstructions, shorter propagation distances must be used. When this is done, it will be clear that access points often cannot be placed precisely in the middle of the circle, so other adjustments must be made.

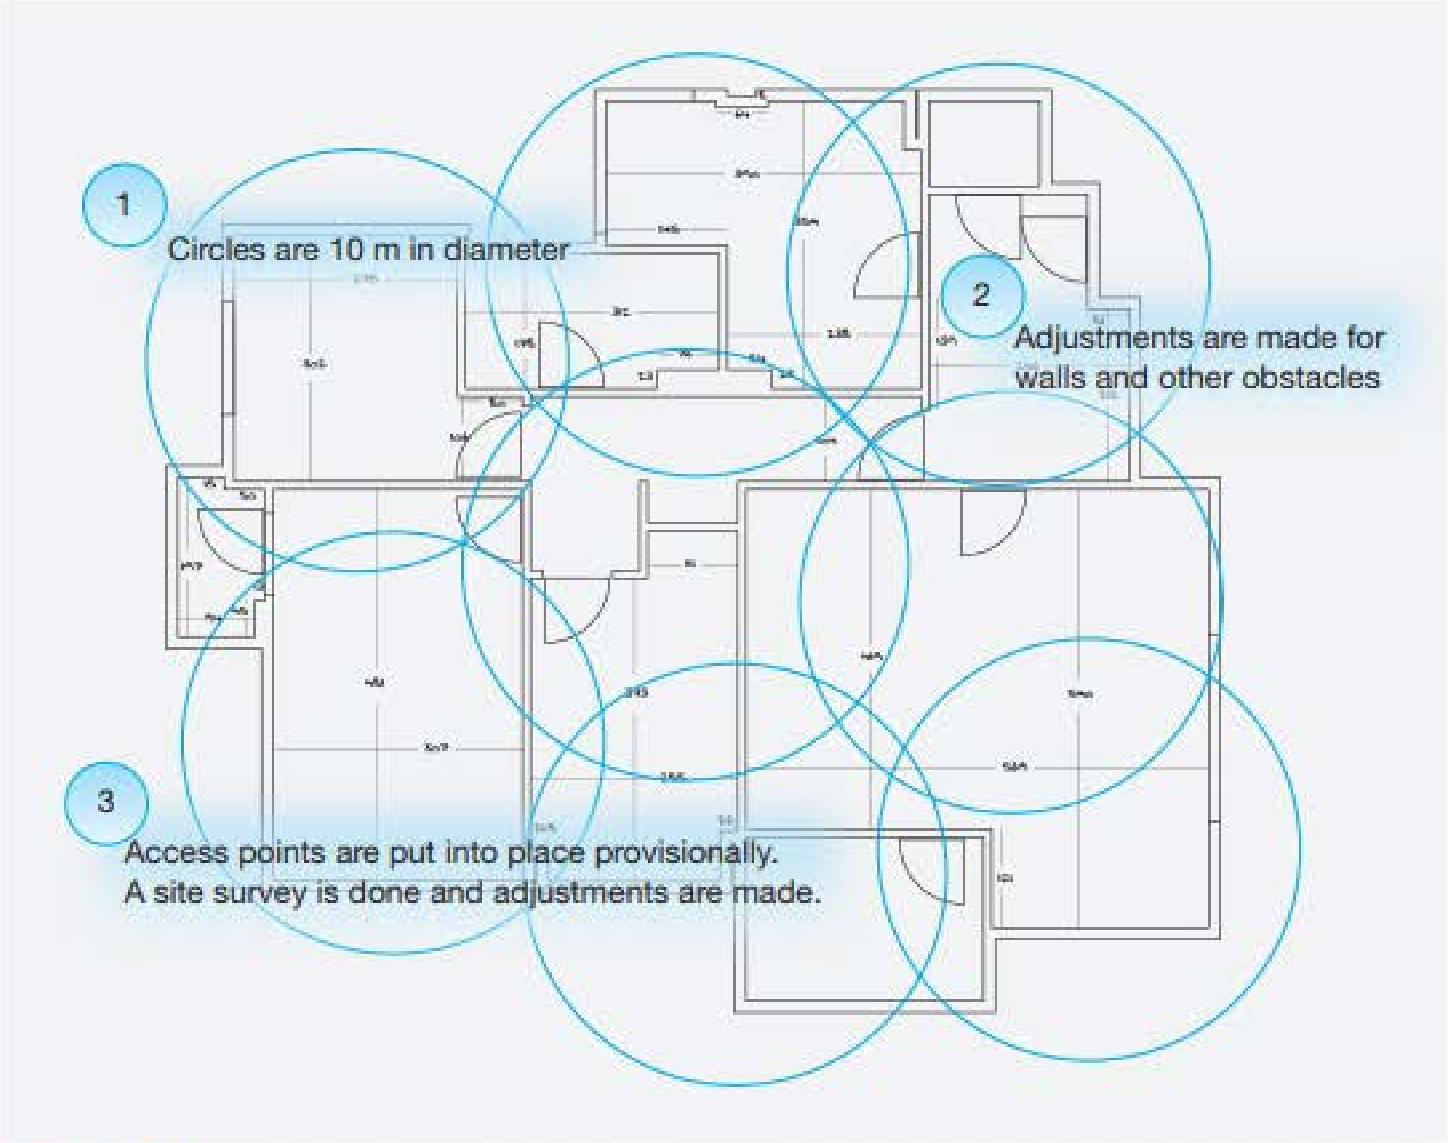
1
Circles are 10 m in diameter
2
Adjustments are made for walls and other obstacles
3
Access points are put into place provisionally.
A site survey is done and adjustments are made.

FIGURE 7-13 Access Point Placement in a Building

Of course, in a multistory building, planning must be done in three dimensions. The "circles" are now bubbles with radiuses of 10 meters. Again, the goal is to provide coverage to all points within the building while reducing overlap as much as possible.

Finally, planners assign channels to access point positions. They attempt to minimize co-channel interference while doing so.

Installation and Initial Site Surveys Next, the access points are installed provisionally in the planned locations. However, the implementation work has just begun. When each access point is installed, an initial **site survey** must be done of the area to discover any dead spots or other problems. This requires **signal analysis software**, which can run on a notebook computer or even a smartphone.

When areas with poor signal strength are found, surrounding access points must be moved appropriately, or their signal strengths must be adjusted until all areas have good signal strength. Users should now have good service.

Ongoing Site Surveys Although the initial site survey should result in good service, conditions will change with time. More people may be given desks in an access point's service areas, signal obstructions may be put up for business purposes, and other changes may occur. Site surveys must be done frequently and routinely; they also may be done in response to specific reports of problems.

Test Your Understanding

9. a) Describe the process by which access point locations are determined. b) When must firms do site surveys to give users good service?

Centralized Management

Large organizations have hundreds or thousands of access points. Traveling to each one for manual configuration and troubleshooting would be extremely expensive. To keep management labor costs under control, organizations must be able to manage access points remotely. The Simple Network Management Protocol, which we saw in Chapter 3, makes this possible. Figure 7-14 shows that the management console constantly requests data from the individual access points. This data includes signal strengths, indications of interference, error rates, configuration settings, power levels, channels, security settings of nearby access points, and other diagnostic information.

If the administrator detects a problem in the network when reading the data, he or she can send SNMP Set commands to access points to increase power, decrease power, switch channels, or make other changes.

The figure also shows a wireless access point initiating an SNMP trap command. A trap might indicate an abnormal error rate, the detection of a rogue access point, or disassociate messages that break connections. The last category, **disassociate messages**, may indicate that an attacker is committing a denial-of-service attack by sending disassociate messages to wireless clients, telling them to stop using the access point. This knocks them off the network.

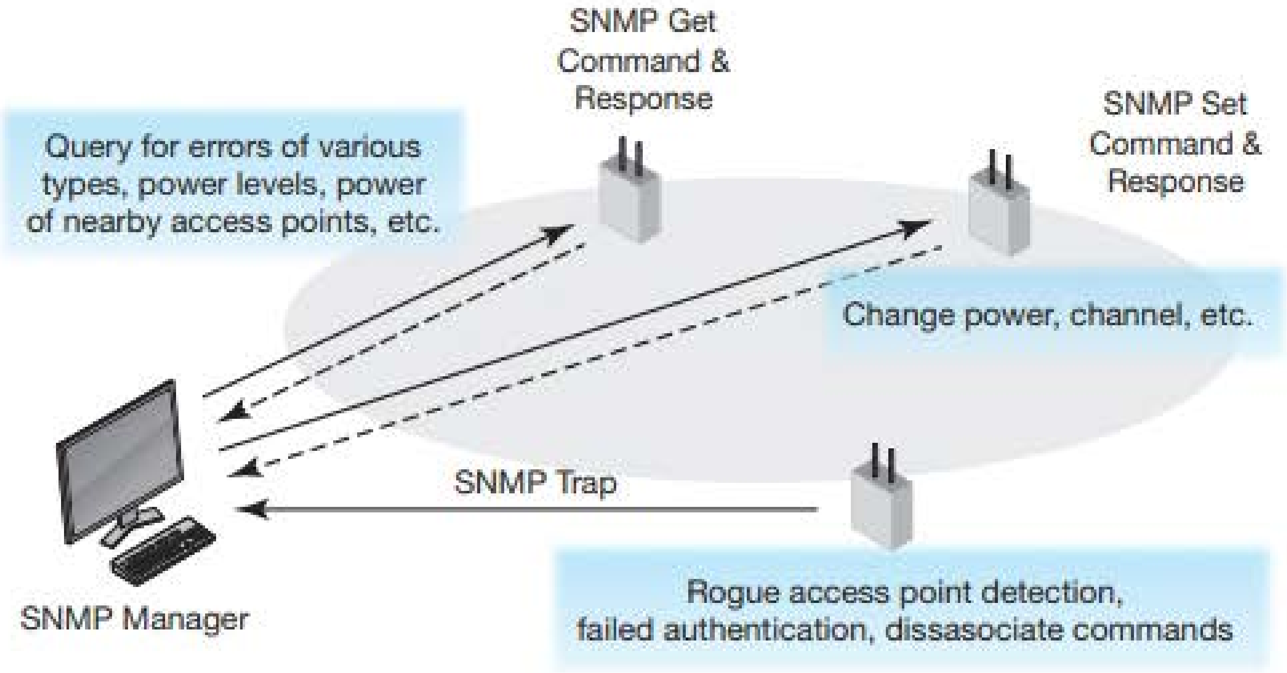
SNMP Get
Command &
Response
Query for errors of various
types, power levels, power
of nearby access points, etc.
SNMP Set
Command &
Response
Change power, channel, etc.
SNMP Trap
SNMP Manager
Rogue access point detection,
failed authentication, dissasociate commands

FIGURE 7-14 Remote Access Point Management

Centralized network management software and hardware on the management console and switches or access points is expensive. However, it greatly reduces management labor, so there should be considerable net savings from its use.

In addition, centralized WLAN management's wireless intrusion detection functionality is the only real way to manage WLAN security. Manual detection of threats would be far too slow and require prohibitive amounts of labor.

Test Your Understanding

10. a) How might a security administrator use SNMP Get commands to access points? b) How does centralized management provide for the detection of rogue access points? c) Comment on the cost of central access point management.

IN MORE DEPTH

Expressing Power Ratios in Decibels

Signal power is usually measured in **milliwatts (mW)**. Networking professionals often compare two signal strengths. For instance, if signal power is 20 mW at 10 meters and 2 mW at 20 meters, the ratio of the second power to the first is 0.1. To give another example, if a larger antenna doubles a transceiver's transmission power, then the ratio of the final power to the initial power is 2:1. Power ratios are expressed in several ways—as decimal numbers, percentages, or ratios (such as 2:1).

Calculating Decibel Values for Power Ratios

Networking professionals typically express the ratio of two powers in **decibels (dB)**, using Equation 7-1. L_{dB} is the decibel relative value of two power levels, P_1 and P_2. P_1 is the initial power level. P_2 is the final power level. The equation shows that the decibel expressions use a logarithmic scale.

$$L_{dB} = 10 * Log_{10}\left(\frac{P_2}{P_1}\right)$$

(Equation 7-1)

This looks complicated, but it really is not. Figure 7-15 shows how to do decibel calculations in Excel or some other spreadsheet program. In the first example, the initial power is 40 mW and the final power is 10 mW. This gives a power ratio of 0.25. Excel has a LOG10 function, and this is

Data or Formula	Example 1: Attenuation	Example 2: Amplification
Initial Power: P_1 (mW)	40	10
Final Power: P_2 (mW)	10	30
P_2/P_1	0.25	3
LOG10(P_2/P_1)	−0.602059991	0.477121255
L_{db}: 10*LOG10(P_2/P_1)	−6.020599913 (Negative)	4.771212547 (Positive)

FIGURE 7-15 Decibel Calculation for Power Levels

(continued)

applied to the power ratio. The result is -0.602. This logarithm is multiplied by a factor of 10. This gives a value of −6.02 decibels. Whenever the second value is smaller than the initial value, the decibel value is negative.

Whenever the second value is smaller than the initial value, the decibel value is negative.

In the second example, the final power is *larger* than the initial power. For example, the signal may be increased by a larger antenna. The initial power is 10 mW, and the final signal power is 30 mW. This gives a power ratio of 3:1. This time, the decibel value is 4.77 dB, a positive value. Whenever the second value is larger than the initial value, the decibel value is positive.

Whenever the second value is larger than the initial value, the decibel value is positive.

Test Your Understanding

11. a) The power level at 10 meters is 100 mW. At 20 meters, it is 5 mW. How many decibels has it lost? b) Compared to an omnidirectional antenna, a dish antenna quadruples radiated power. How much is this change in decibels? c) Compute the decibel value for a power ratio of 17:1. d) Of 1:33.

Approximating Decibel Values

You do not always have a spreadsheet program with you. Nobody can calculate logarithms in his or her head. However, you can use two approximations to roughly estimate decibel values if you know the power ratio.

First, Figure 7-16 shows that if you double the signal power, this is a gain of approximately 3 dB. If you quadruple the signal power, this is a gain of approximately 6 dB. For each additional doubling, the gain is another approximately 3 dB. This calculation is approximate, but it is close. (The exact value is 3.0103.)

Powers of 2		Powers of 10	
Power Ratio	Approximate dB	Power Ratio	Exact dB
2	3 dB	10	10 dB
4	6 dB	100	20 dB
8	9 dB	1,000	
16		10,000	
32		100,000	
1/2	−3 dB	1/10	−10 dB
1/4		1/100	
1/8		1/1,000	

FIGURE 7-16 Decibel Approximations

Each doubling of power gives a gain of approximately 3 dB.
Each multiplying by 10 in power gives a gain of approximately 10 dB.

What if the power ratio is less than 1? If it is 0.5, then the decibel value is approximately −3 dB. Cutting this power in half gives −6 dB. Every additional halving is another −3 dB. Again, if the power ratio is greater than 1, the decibel value will be positive, and if the power ratio is less than 1, the decibel value will be negative.

If the power ratio is greater than 1, the decibel value will be positive, and if the power ratio is less than 1, the decibel value will be negative.

For positive or negative powers of 10, the situation is similar. A power ratio of 10:1 is exactly 10 dB. (There is no approximation.) A power ratio of 100 is 20 dB. Each further increase by a factor of 10 is another 10 dB. Likewise, a power ratio of 0.1 is −10 dB, and a power ratio of 0.01 is −20 dB.

What if a ratio is *not* a multiple of 2 or 10? What if it is, for example, 3:1? Well, 2:1 is 3 dB; and 4:1 is 6 dB. So the answer is somewhere between 3 dB and 6 dB. That is not very precise, but it can be useful in practical situations. The 2:1 and 10:1 approximation will not always be useful, but they are good tools for networking professionals to have.

Test Your Understanding

12. a) Fill in the missing values in Figure 7-16. Approximate, without using Excel, the decibels for a ratio for b) 8:1. c) 9:1. d) 110:1. e) 1:7. f) 1:90. Use terms like "a little higher than" or "a little lower than."

PEER-TO-PEER PROTOCOLS FOR THE INTERNET OF THINGS (IoT)

In Chapter 1, you learned that the Internet of Things involves hosts talking to other hosts without human involvement. IoT machines simply communicate directly with one another to coordinate their work. Much of their communication will be peer-to-peer, that is between the two devices with no server involved. Normally, more distance and speed are desirable. But "fast and far" is also a recipe for draining batteries rapidly. If communication takes place over short distances and at slower speeds, this "slow and close" communication extends battery life. As we will see later, RFID tags can transmit without have *any* internal power.

"Slow and close" communication extends battery life.

Battery drain is especially important for IoT devices too small to plug into a wall or to use a traditional rechargeable battery. Many use small **coin batteries** like the one in Figure 7-17. These batteries will need to last months or even years in most devices. This is possible only if energy demands are kept very low. Energy restrictions for IoT transmissions that use coin batteries require new standards.

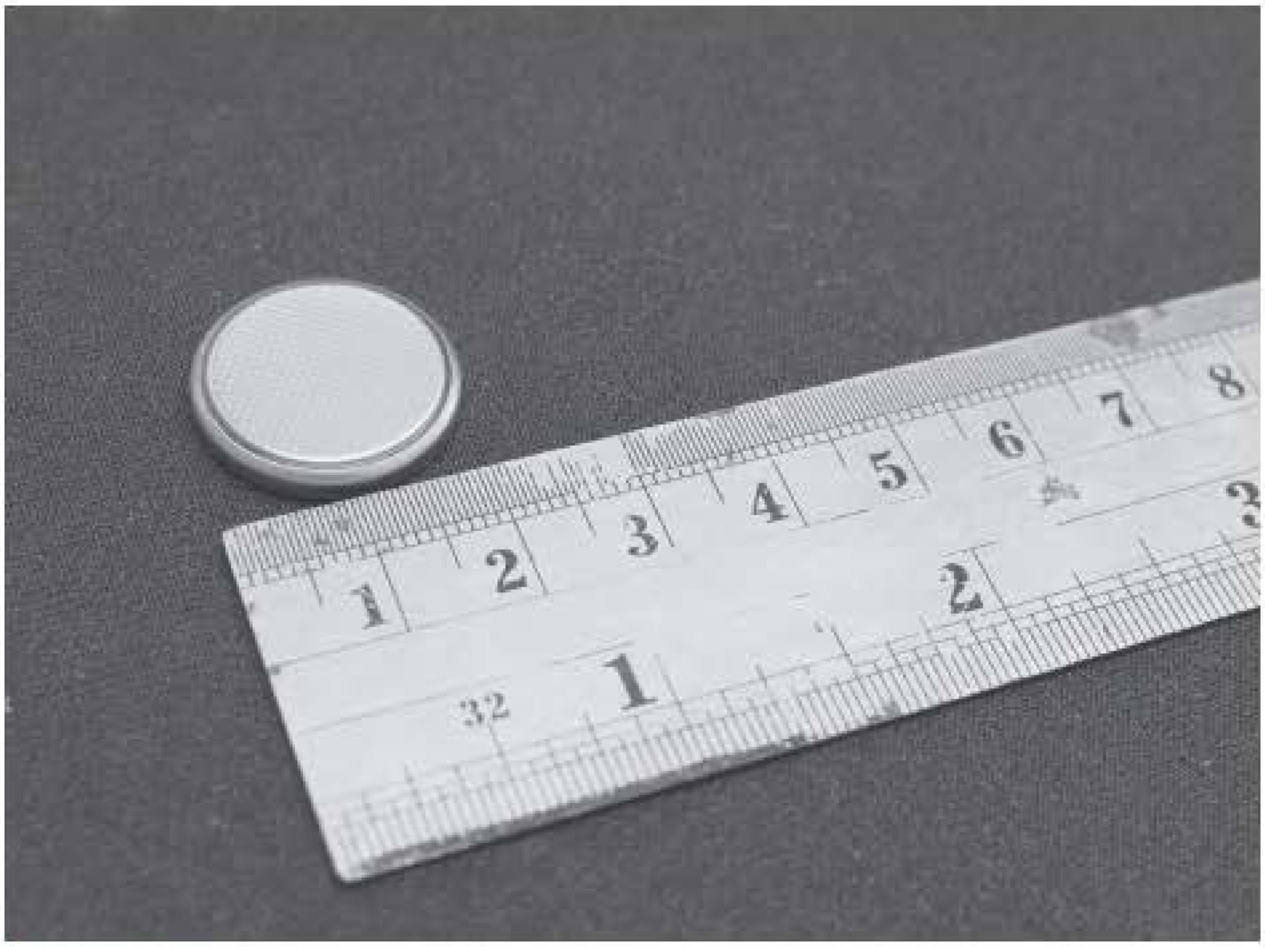

FIGURE 7-17 Coin Battery

Energy Restrictions for IoT Devices that use Coin Batteries Require New Standards Figure 7-18 shows several communication peer-to-peer protocols that promise to be attractive for IoT communication in general. They vary widely in the possible distance between the two devices and in transmission speed. Those at the lowest level use the least energy and are suitable for IoT devices with coin batteries

Test Your Understanding

13. a) Why is low speed and short distance good in the Internet of Things? b) Is there a single dominant IoT communication standard?

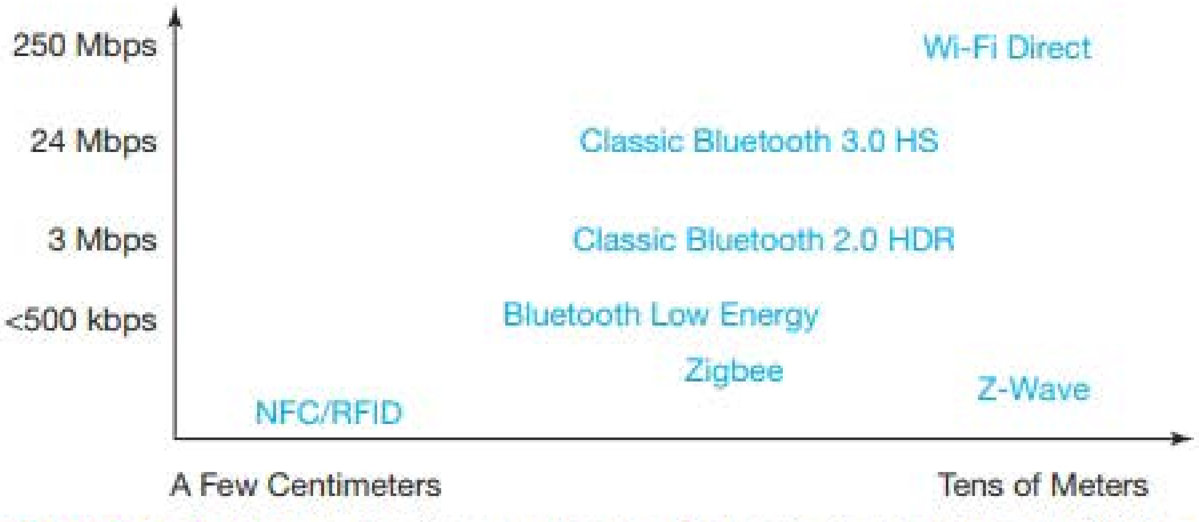

FIGURE 7-18 Peer-to-Peer Communication Protocols for the Internet of Things (IoT)

BLUETOOTH

If you have a wireless headset for your mobile phone or pocket music player, or if you have a hands-free cellular system in your car, you are already using Bluetooth. These are precisely the kinds of short-range moderate-speed applications that Bluetooth was created to handle. **Bluetooth** is a short-range radio technology designed for **personal area networks (PANs)**—small groups of devices in a communication bubble around a person's body or a single desk (Figure 7-19). Bluetooth is essentially a cable replacement technology. In contrast to 802.11, Bluetooth is not standardized by the IEEE. Rather, it is standardized by the **Bluetooth Special Interest Group (SIG)**, which is an industry trade association.

Bluetooth is a short-range radio technology designed for personal area networks (PANs)—small groups of devices around a person's body or a single desk.

Bluetooth is essentially a cable replacement technology.

Classic Bluetooth and Bluetooth Low Energy (LE)

There are two forms of Bluetooth, Classic Bluetooth and Bluetooth Low Energy. Perhaps surprisingly, they are deeply incompatible. Both can use the 2.4 GHz unlicensed radio band, and both use the same radio and antennas, so they do not need entirely separate technology. However, they work very differently. Bluetooth Low Energy is not a mere extension of Classic Bluetooth. Figure 7-20 compares Classic Bluetooth with Bluetooth LE.

Classic Bluetooth and Bluetooth Low Energy are incompatible.

Classic Bluetooth The original version of Bluetooth, **Classic Bluetooth**, has two data rates: an Extended Data Rate speed of 3 Mbps and a High Speed rate of 24 Mbps. The two form a single service with modestly fast normal operation and high burst-speeds operation for occasional file transfers and other actions. Until recently, these were the only types of Bluetooth.

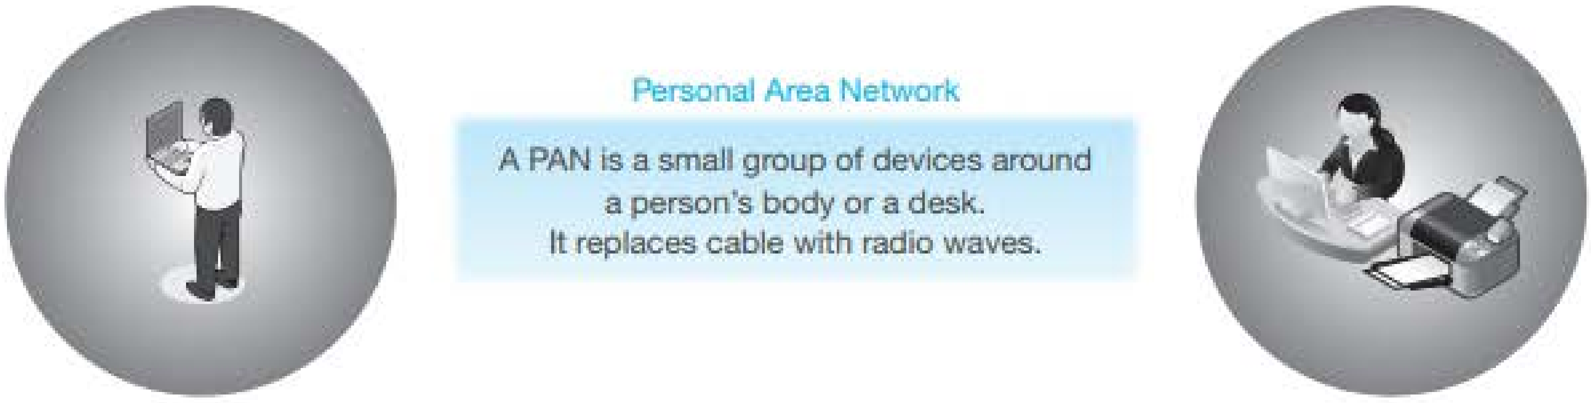
Personal Area Network
A PAN is a small group of devices around
a person's body or a desk.
It replaces cable with radio waves.

FIGURE 7-19 Bluetooth Personal Area Networks (PANs)

Operating Mode	Classic Bluetooth: EDR	Classic Bluetooth: HS	Bluetooth LE
Abbreviation	Enhanced Data Rate	High Speed	Low Energy
Use Case	Headsets, speakers, keyboards, etc. High duty cycle (percentage of time in use)		Fitness trackers Low duty cycle
Principal Benefit	Good performance at modest power	Brief high-speed transfers at modest power	Low cost for very brief, low-speed, and infrequent communication
Speed	Up to 3 Mbps	Up to 24 Mbps	Up to 2 Mbps but usually 125 kbps or 500 kbps.
Power Required	Low (Rechargeable mobile phone battery)		Very Low (coin battery)

FIGURE 7-20 Bluetooth Modes of Operation

Bluetooth Low Energy More recently, the Bluetooth Alliance introduced **Bluetooth Low Energy (Bluetooth LE)**. Compared to Classis Bluetooth, Bluetooth Low Energy has similar range but greatly reduced power consumption. Classic Bluetooth requires wall power or a rechargeable battery. Bluetooth Low Energy was created to work for a new class of small devices, such as light switches, that use a small coin battery that is expected to last for a long time, even years. This requires extremely low energy output.

Bluetooth Low Energy (LE) is for devices with coin batteries.

Dual-Mode and Single-Mode Device Devices with rechargeable batteries, such as mobile phones, usually offer both Classic Bluetooth and Bluetooth LE. They have a rechargeable battery, so they can easily implement both modes. In contrast, small IoT devices typically only have a coin battery. They usually only support Bluetooth LE because even a brief use of traditional Bluetooth would slash battery life.

Test Your Understanding

14. a) What is a PAN? (Do not just spell out the abbreviation.) b) Compare the relative benefits of the two types of Classic Bluetooth. c) Why would you not want to use high-speed Bluetooth all the time? d) What is the benefit of Bluetooth Low Energy? e) What type of battery do very small Bluetooth LE devices require, and why is this important? f) Why do small IoT devices only implement Bluetooth LE?

One-to-One, Master–Slave Operation

Figure 7-21 shows several devices communicating with Bluetooth. The device in the top center is a mobile phone. To its left is a printer. The mobile phone user wishes to print a webpage on the printer. The user selects print, chooses the target printer, and

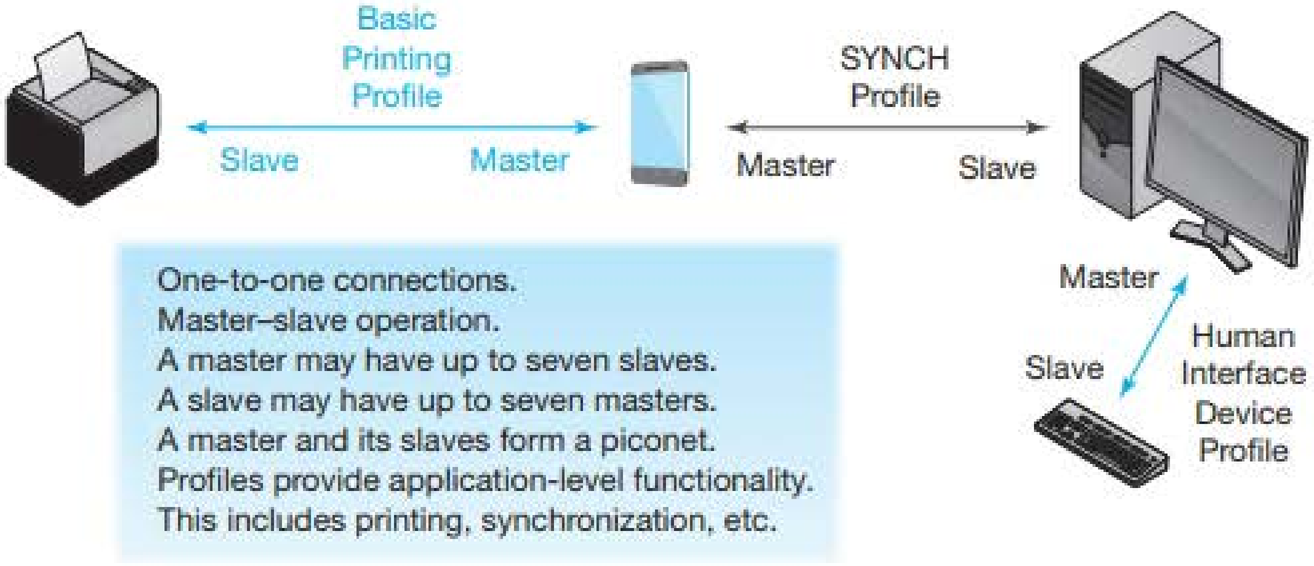
Basic
Printing
Profile
Slave
Master
SYNCH
Profile
Master
Slave
Master
Slave
Human
Interface
Device
Profile
One-to-one connections.
Master–slave operation.
A master may have up to seven slaves.
A slave may have up to seven masters.
A master and its slaves form a piconet.
Profiles provide application-level functionality.
This includes printing, synchronization, etc.

FIGURE 7-21 Bluetooth Operation

prints. The mobile is simultaneously synching files between itself and the computer on the right. At the same time, the computer is communicating with its Bluetooth wireless keyboard.

One-to-One Connections Note that Bluetooth uses **one-to-one connections** between pairs of devices. In the figure, Bluetooth implements a one-to-one connection between the mobile phone and the printer. It also implements one-to-one connections between the mobile phone and the desktop computer and between the desktop and the keyboard. Although the mobile phone connects to two devices, these are separate Bluetooth connections.

Bluetooth always uses point-to-point communication between a pair of devices.

Master–Slave Control In addition, Bluetooth always uses **master–slave control**. One device is the master, the other the slave. In the printing scenario, the mobile device is the master and the printer is the slave. The mobile phone controls the printing process.

In Bluetooth, one device is the master and the other device is the slave. The master controls the slave.

Multiple Slaves and Masters Although communication is always one-to-one, a master may have up to seven slaves simultaneously. A master and its slaves comprise a **piconet**. In Classic Bluetooth, a slave may also have up to seven masters. This means that a slave may be part of multiple piconets.

It is possible for a Bluetooth device to be a master and a slave simultaneously. Consider the relationship between the mobile phone and the desktop computer. The two are synchronizing information. The mobile phone is the master, and the desktop is the slave. However, the desktop is simultaneously a master to the keyboard.

A Bluetooth device may be a master of one device and a slave to another device simultaneously.

Test Your Understanding

15. a) What does it mean that Bluetooth uses one-to-one operation? b) Is this still true if a master communicates with four slaves simultaneously? c) Can a Bluetooth master have multiple slaves? d) Can a Bluetooth slave have two masters? e) Can a Bluetooth device be both a master and a slave simultaneously?

Bluetooth Profiles

For Wi-Fi, the 802.11 Working Group did not have to worry about applications. Desktops and laptop PCs on 802.11 WLANs already had many applications. For example, word processing programs knew how to work with printers in general, although working with a new printer usually required the computer to add the printer to its configuration and install device drivers.

However, there were no application protocols in existence for PAN applications such as wirelessly controlling keyboards, telephone headsets, printers, and other devices. Consequently, in addition to defining physical and data link layer transmission standards, the Bluetooth SIG also defined application profiles, which are called **Bluetooth profiles**. Profiles govern how devices share information and specify control messages for various uses. Figure 7-21 shows three Bluetooth profiles.

- For printing, the mobile phone uses the *basic printing profile (BPP)*. A Bluetooth device can print to any BPP compliant printer without having to install a printer driver on the Bluetooth device.
- For synchronizing information with the desktop computer, the mobile phone uses the *synchronization profile (SYNCH)*. It simply selects the computer and begins the synchronization.
- Desktop computers, in turn, use the *human interface device (HID)* profile for mice, keyboards, and other input devices. Again, there is no prior setup beyond selecting the device.

Test Your Understanding

16. a) Why would it be nice if Wi-Fi offered a basic printing profile? b) What Bluetooth profile would you use for a game joystick, based on information in the text?

Bluetooth Low Energy

In general, Bluetooth LE and Classic Bluetooth are outwardly similar. Both use one-to-one connections and master–slave control. In both, a master may also have up to seven slaves in its piconet. (The ability of a slave to serve multiple masters is still under development in Bluetooth LE.)

Similarities between the Modes

- One-to-one connections, master-slave operation
- Master may have seven slaves

Power-Efficient Design in Bluetooth Low Energy

- Usually 0.01 W to 0.5 W, compared to Classic Bluetooth's nominal 1 W (and usually toward the lower end)
- Transmits slowly over short distances
- Infrequent transmissions with deep sleep between
- Terse connection openings (100 ms for Classic Bluetooth, 3 ms for Bluetooth LE)
- Energy conservation pervades design (energy-saving spread spectrum method)

Advertising and Connections

- Small IoT devices periodically transmit advertising messages
  - Announce their existence and purpose
- To connect, the master scans, initiates an opening
- Beacons are advertising messages that include useful information
  - At airports, announcements of delays
  - In stores, a coupon as a shopper nears a department
  - Navigation directions within a building

Profiles

- Specific to small IoT devices
- Fitness trackers
- Glucose meter reading

FIGURE 7-22 Bluetooth Low Energy (Study Figure)

Advertisements and Connections There is one thing that Bluetooth LE slaves must do frequently. They must wake up and transmit a brief **advertisement message** to announce their existence and say what they can do. When the master needs a connection, it scans for such advertisements. The master then initiates a connection. The two parties then switch to master–slave communication. Fortunately, advertisement messages are brief, and there is significant time between them. This limits the power drain they create.

Beacons Bluetooth LE extends the advertisement message by adding the concept of **beacons**. These are advertising messages that include potentially useful information. Beacons can offer you a coupon when you step into a store, give your mobile phone directions for navigating through a hospital, tell you how many tickets are available for a movie near you, or inform you of flight delays in an airport. Masters can read this information from beacons without even making a connection.

Profiles Like Classic Bluetooth, Bluetooth LE has profiles. Of course, Bluetooth LE profiles reflect their use cases. In medicine, there are profiles for reading glucose meters, and there are profiles for heart rate monitors. In sports, there is a fitness tracker profile and a location and navigation profile.

Test Your Understanding

17. a) What is a typical speed, distance, and power consumption for Bluetooth LE slaves? b) What are Bluetooth LE advertising messages? c) How do Bluetooth LE beacons differ from basic advertisement messages? d) In general, how do Bluetooth LE profiles differ from Classic Bluetooth profiles? (You will have to think about this one a little.)

OTHER PROMISING IoT TRANSMISSION STANDARDS

Near Field Communication (NFC)

We have seen that when radios transmit, they produce electromagnetic waves that propagate away, taking energy with them. Very close to an antenna, there is another phenomenon, a "near field," which pulses outward a short distance, then is reabsorbed into the antenna. This near field does not propagate away from the antenna. As Figure 7-23 shows, the near field only extends a few inches from a phone with a **near field communication (NFC)** chip. However, the near field can be used for communication. One device (in the figure a mobile phone) modulates the near field to send information. The other device (in this case a point-of-sale terminal), reads changes in the phone's near field. The POS device can also manipulate the near field in a way that allows it to send information to the phone. This manipulation takes very little energy. In the most extreme case, passive **radio frequency ID (RFID)** circuits have no power at all (Figure 7-24). They use the power of the near field itself to modulate the near field to send information.

Because NFC transmission takes place at such extremely low speeds as 434 kbps, it cannot transfer very much information. Also, near field transmission distances only extend a few inches. However, this is fine for many purposes.

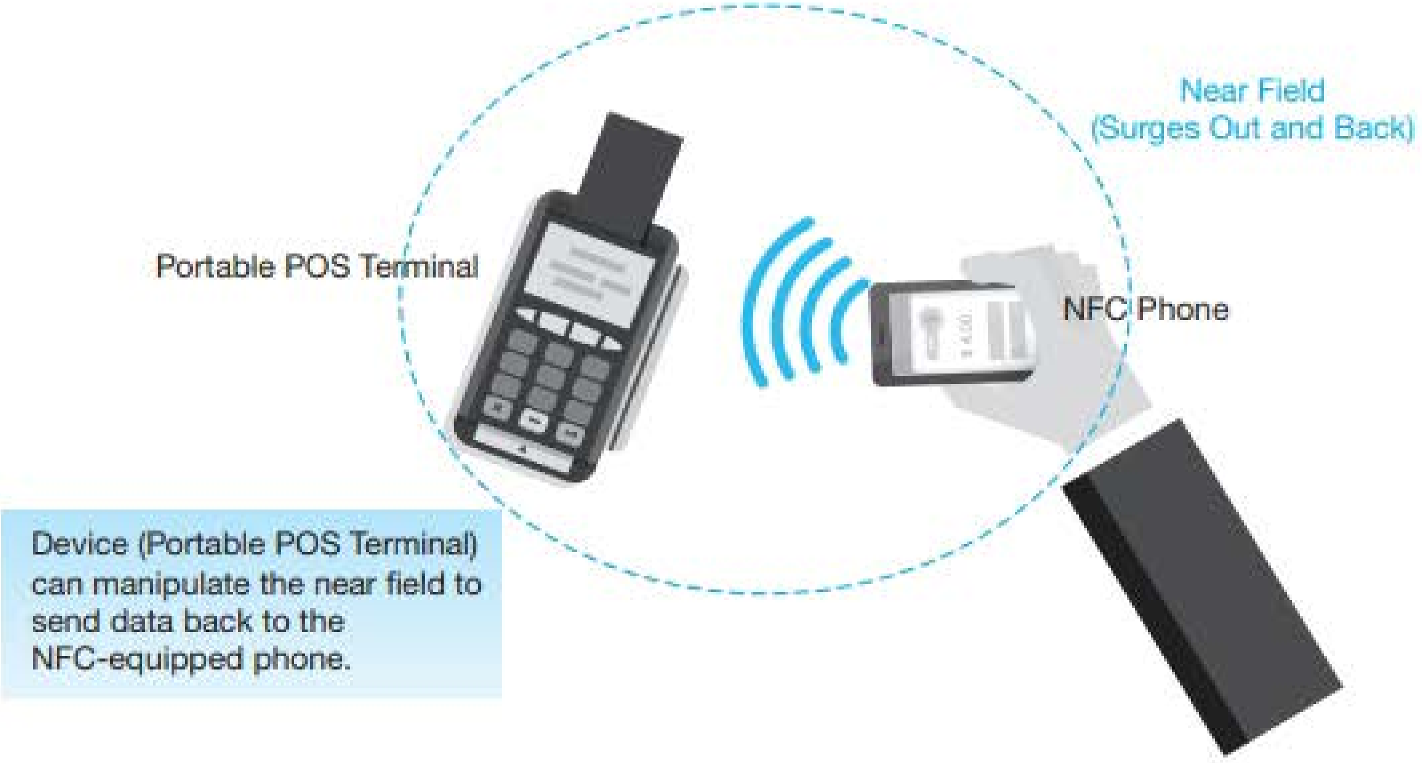

Device (Portable POS Terminal) can manipulate the near field to send data back to the NFC-equipped phone.

FIGURE 7-23 Near-Field Communication (NFC)

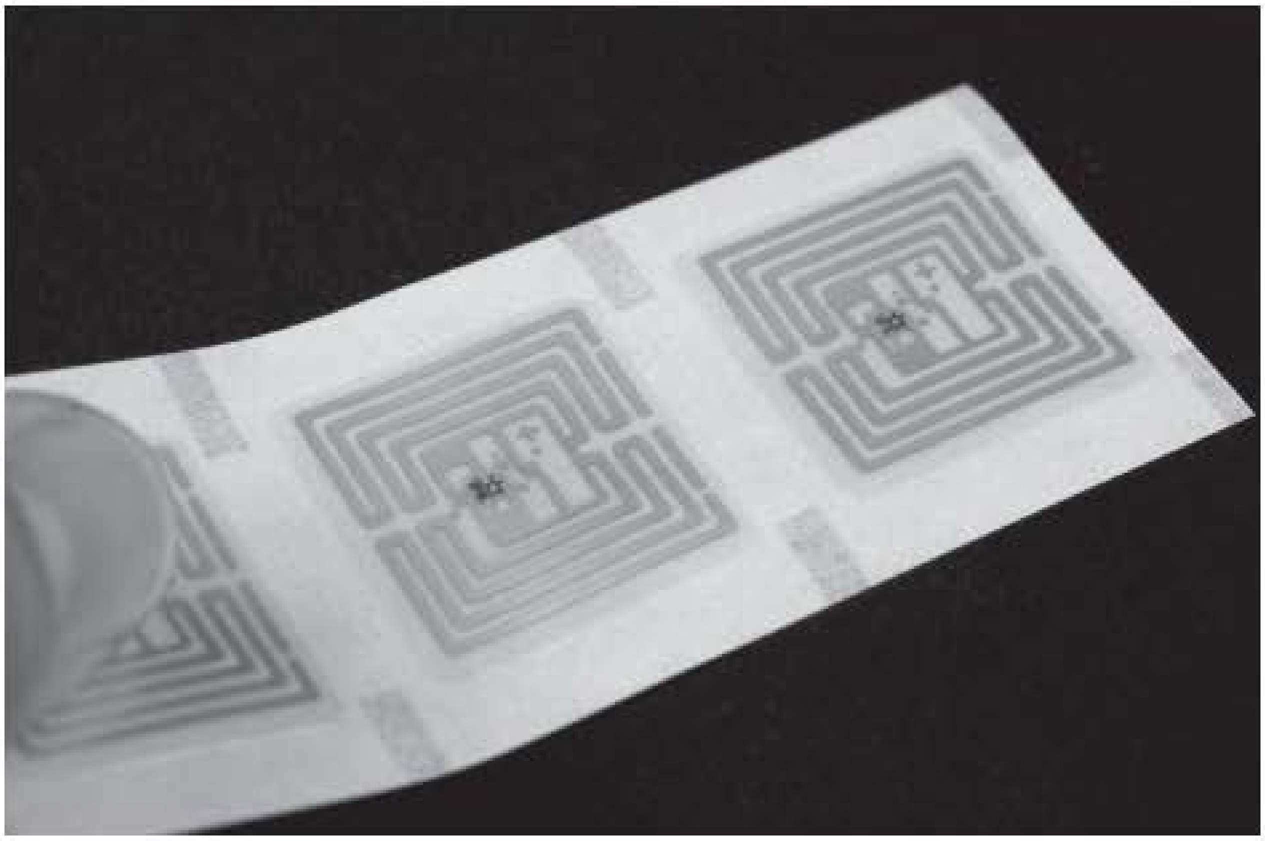

FIGURE 7-24 Passive Radio Frequency ID (RFID) Circuits for Near Field Communication

NFC standards are still in flux. All NFC protocols use transmission in the 13.56 MHz unlicensed service band created for this purpose. Its technical standards are also largely set. However, for applications such as point-of-sale payments, there are competing application standards as phone vendors and others push for dominance in a rapidly changing market.

Test Your Understanding

18. a) When two devices communicate using NFC, how close must they be? b) How does near field communication differ from normal radio communication? c) Passive RFID chips have no batteries. How can they transmit when queried? d) What is the state of NFC standards?

Wi-Fi Direct

In Figure 7-25, Wi-Fi Direct is the poster child for fast-and-far peer-to-peer communication. When you have used 802.11, it has involved an access point. However, the 802.11 standard has always included an ad hoc mode, in which two wireless Wi-Fi hosts communicate directly. This provides medium-speed communication over typical Wi-Fi distances. It has no problem connecting devices that are at different ends of a house. The Wi-Fi Alliance calls this **Wi-Fi Direct**. (Sometimes, people shorten this to Wi-Di.) Confusingly, companies that implement it on their phones, tablets, and other devices create their own name for it. This has plagued Wi-Fi Direct with marketplace confusion, and technical interoperability across vendors has been uneven.

Test Your Understanding

19. How is the access point used in Wi-Fi Direct?

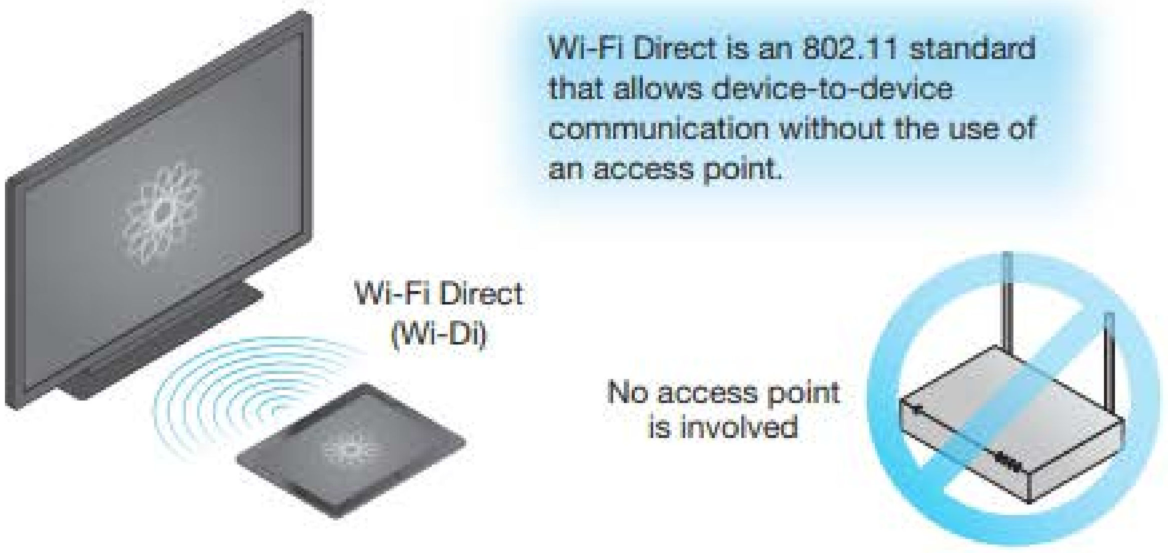
Wi-Fi Direct is an 802.11 standard that allows device-to-device communication without the use of an access point.
Wi-Fi Direct (Wi-Di)
No access point is involved

FIGURE 7-25 Wi-Fi Direct

Zigbee and Z-Wave

Near-field communication and Wi-Fi direct are designed for communication between pairs of devices. Beyond that, two standards have been created to network IoT devices in a mesh. One of these is **Zigbee** (named after the dance that bees do to communicate directions to flowers with nectar). Figure 7-26 shows a Zigbee **ad-hoc wireless network**. Ad hoc means that the network is **self-organizing**. There is no need to create a complete design in the beginning, and the network adapts automatically to changes.

Zigbee Controller (and Often Gateway)

The heart of the network is the **Zigbee controller**. The controller coordinates the network, so every Zigbee network must have one. Larger Zigbee networks may have several. In home and small business networks, the controller may also be a gateway to the Internet. ("Gateway" was an early name for "router.") In fact, the controller/gateway may actually be an Internet access router with built-in Zigbee controller functionality.

Zigbee End Devices

Zigbee end devices are IoT devices such as light switches, light bulbs, thermostats, air conditioners, door locks, and televisions. These devices must be able to communicate via the Zigbee protocol.

Zigbee Routers

End devices may connect to a controller, but they may also connect to **Zigbee routers**. Routers permit Zigbee networks to span larger distances than a single controller. For example, in Figure 7-26, Switch 1 and the Light Bulb may be too far apart to communicate directly. However, when Switch 1 transmits, its frame goes to the controller/gateway (which again may be a residential access router) to Router R1, which forwards it to the Light Bulb.[8]

[8] A number of companies now put Zigbee functionality in their access routers. In fact, some are beginning to create meshes of access routers that use Zigbee to communicate. (Others are beginning to create meshes of access routers using different protocols as well.) In the past, there have been range extenders that you could put into a distant room to extend your basic home access router's range. Mesh access routers give full access router functionality on each device in the mesh. They also are self-organizing, making them easy to install.

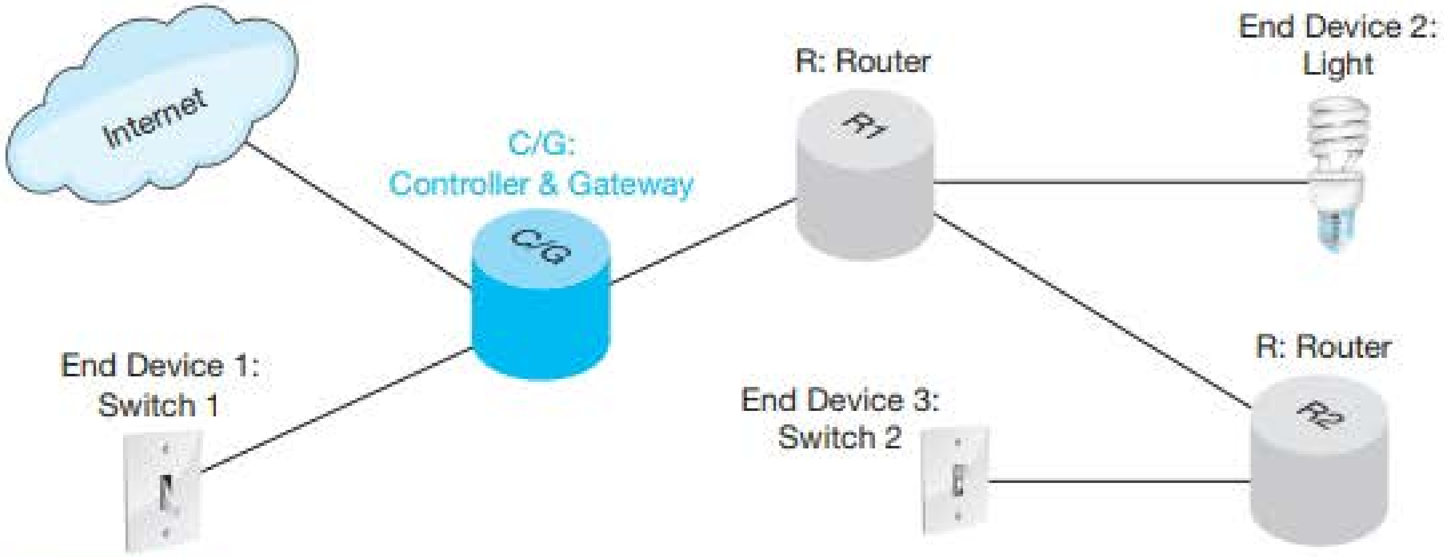

FIGURE 7-26 Zigbee Ad-Hoc Wireless Network

Dual-Band Use in Zigbee Zigbee operates in two unlicensed bands. One is the familiar 2.4 GHz unlicensed band. Another is the 800/900 MHz unlicensed band. It gets this split designation because the band is in the 800 MHz range in Europe but in the 900 MHz range in North America. The lower band can carry signals slightly farther, but the higher band can transmit signals slightly faster, albeit at the cost of slightly greater energy use.[9]

Z-Wave **Z-Wave** is a similar ad hoc wireless networking protocol. Z-Wave and Zigbee are the most popular standards for ad hoc wireless networking, but others are beginning to appear. Z-Wave is similar in speed and range to Zigbee for small-to-mid-size networks, and both have 128-bit AES encryption and other good security protections. In corporations, size limitations and other factors do become important in large ad hoc wireless networks that span large building. For example, Z-Wave only operates in the 800/900 MHz ISM bands.

Test Your Understanding

20. a) What kind of network is Zigbee used for? b) Compare the roles of Zigbee controllers, Zigbee end devices, and Zigbee routers. In what radio bands does Zigbee operate? c) What other ad hoc networking protocol is widely used? d) In what radio band or bands does it operate?

SECURITY IN THE INTERNET OF THINGS

Security is a complex situation for emerging local wireless transmission technologies. Like all wireless technologies, they are vulnerable to eavesdropping, data modification, and impersonation.

[9] An intriguing recent development has been the creation of Green Power devices in Zigbee. These are devices that do not require a battery at all. For example, light switches are powered by the act of flipping the switch. The energy of this motion is captured and used to send out a signal.

Threats

- Eavesdropping
- Data modification
- Impersonation

Cryptological Security

- Some have no cryptological security
- Example: Near field communication for reading passive RFID tags
- They rely on short transmission distances to foil eavesdroppers
- However, directional antennas and amplifiers can read signals that are far longer than distances in standards

Strength of Security

- Some have reasonably good security
- Example: Bluetooth
- However, still not as strong as 802.11i security

Device Loss or Theft

- In this age of bring your own device (BYOD) to work, this is a serious problem
- Most devices are only protected by short PINs

Maturity

- In general, new security technologies take some time to mature
- During this period, they often have vulnerabilities that must be fixed quickly
- User companies must master security for each new technology they use

FIGURE 7-27 Security in Emerging Local Wireless Technologies (Study Figure)

Some of these technologies have no cryptographic security at all. The classic example is using NFC to read passive RFID tags. These technologies assume that eavesdroppers cannot get close enough to read the information because maximum transmission distances are very small. However, distances in the standards are for normal devices. Eavesdroppers with highly directional antennas and amplifiers can intercept signals over much longer distances. Bluetooth probably has the best security among emerging wireless technologies, but its security is still weaker than 802.11i's security.

In today's world of bring your own device (BYOD) to work, emerging local wireless technologies make a worrisome corporate security situation even more problematic. For example, if devices such as mobile phones are lost or stolen, they are often protected only by brief PINs, if they are protected at all. Many of these devices contain sensitive corporate information, and even if they do not, they may allow attackers to log into sensitive servers on the corporate network.

As a rule, new security technologies tend to have vulnerabilities that take time to discover and protect against. One must hope that technology vendors will be quicker to act than attackers. In any case, companies need to fully understand security for each technology.

Test Your Understanding

21. a) Why is a short transmission range not a protection against eavesdroppers? b) Describe the state of cryptographic security for new transmission standards. c) Why is device theft or loss a serious risk?

END-OF-CHAPTER QUESTIONS

Thought Questions

7-1. In the Ms. Betsy Davis case at the beginning of the chapter, the access point on the local network did not have security. This makes a man-in-the-middle attack much easier. a) Given what you learned in this chapter, describe how it would be possible to use a man-in-the-middle attack if the legitimate access point does not implement 802.11i. b) How can you get the user to associate with your evil twin access point? (The answer is not in the text.)

7-2. a) A straight road with government-provided Internet will receive 16 access points that are 10 meters apart. About how many access points would be needed if the city decided to increase the distance to 20 meters? Just give a reasoned estimate. b) A single-story building is 100 meters by 100 meters. If access points are placed 10 meters apart on average, about how many access points will be needed? c) If the same building is also 100 meters tall, how many access points will be needed? d) Repeat Part b if the access points are 20 meters apart? Give a reasoned estimate.

7-3. (If you read the box "Expressing Power Ratios in Decibels") a) If you are told that a signal has attenuated by 20 dB, about how much has it attenuated? b) What would you say about attenuation if you were told that a signal has attenuated by 19 dB? You must approximate. c) What would you say about attenuation if you were told that a signal has attenuated by 7 dB?

7-4. Create a policy for 802.11 Wi-Fi security in a wireless network in a five-person company with one access point. This is not a trivial task. Do not just jot down a few notes. *Make it a one-page document for people in your firm to read, not something for your teacher to read.*

Perspective Questions

7-5. What was the most surprising thing you learned in this chapter?

7-6. What was the most difficult part of this chapter for you?

This page intentionally left blank

Chapter **8**

TCP/IP Internetworking I

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Define hierarchical IPv4 addresses, networks and subnets, border and internal routers, and masks.
- Given an arriving packet's destination IPv4 address, explain what the router will do with the packet based on its routing table.
- Explain the IPv4 packet header fields we did not see in earlier chapters.
- Explain the IPv6 packet's main header fields and IPv6's use of extension headers.
- Convert a 128-bit IPv6 address into canonical text notation consistent with RFC 5952.
- Explain TCP segment fields, UDP datagram fields, and TCP session closings.
- Explain why application message fragmentation is not possible with UDP.

INTRODUCTION

Switched networks and wireless networks are governed by Layer 1 and Layer 2 standards. We looked at single network standards in Chapters 5, 6, and 7. In this chapter and the next, we look at internetworking, which is governed by Layer 3 and Layer 4 standards. Figure 8-1 illustrates this situation.

We only look at TCP/IP internetworking because TCP/IP dominates the work of network professionals at the internet and transport layers. However, real-world routers cannot limit themselves to TCP/IP internetworking. Commercial routers are multiprotocol routers, which can route not only IP packets but also IPX packets, SNA

Layer	Name	Ethernet LANs	Wireless LANs	The Internet	Dominant Standards Agency(ies) / Architecture
5	Application				None
4	Transport			TCP, UDP	IETF / TCP/IP
3	Internet			IP	IETF / TCP/IP
2	Data Link	802.3	802.11		ISO and ITU-T / OSI
1	Physical	802.3	802.11		ISO and ITU-T / OSI

FIGURE 8-1 Standards Layers Recap

packets, AppleTalk packets, and other minor types of packets that we cannot cover in an introductory text book.

We examined the TCP/IP architecture in Chapter 2. We focused on IP, TCP, and UDP, although we looked at a few other TCP/IP standards. Figure 8-2 shows a few of the many standards the Internet Engineering Task Force (IETF) has created within the TCP/IP architecture. Some of the standards are shaded in this figure. We will look at them in this chapter and in Chapter 9.

Many of these are supervisory standards that govern how routers and hosts on the Internet work beyond the delivery of packets. As a worldwide network, the Internet needs many more supervisory protocols to govern it than do Ethernet and Wi-Fi networks.

Test Your Understanding

1. a) Which two layers standardize Ethernet and Wi-Fi operation? b) Which two layers standardize most of the Internet's operation? c) What do IP, TCP, and UDP govern? d) What do TCP/IP supervisory protocols govern?

5 Application	User Applications			Supervisory Applications			
	HTTP	SMTP	Many Others	DNS	BGP	Many Others	
4 Transport	TCP				UDP		
3 Internet	IPv4 and IPv6			ICMP	OSPF	EIGRP	ARP
2 Data Link	None: Use OSI Standards						
1 Physical	None: Use OSI Standards						

Note: Shaded protocols are discussed in this chapter and in Chapter 9.

FIGURE 8-2 Major TCP/IP Standards

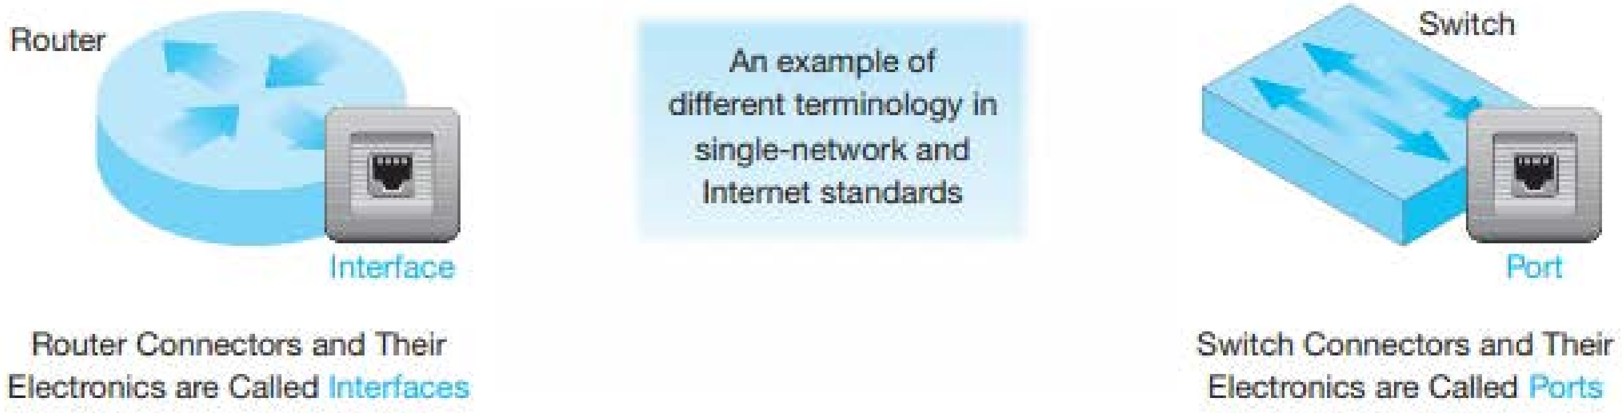
Router
Interface
Router Connectors and Their Electronics are Called Interfaces
An example of different terminology in single-network and Internet standards
Switch
Port
Switch Connectors and Their Electronics are Called Ports

FIGURE 8-3 Router Interfaces and Switch Ports

IP ROUTING

Routers make decisions about forwarding packets—which interface to send an arriving packet back out to get it closer to its destination. For routers, ports are called **interfaces**. This is another example of how terminology differs for single networks and the Internet. Single-network and Internet standards are governed by different organizations, and they use different terminology (Figure 8-3).

Router ports are called interfaces.

In this chapter, we will see that router forwarding is much more complex than the Ethernet switching. Higher complexity requires routers to do more work per arriving packet than switches do per arriving frame. Consequently, routers are more expensive than switches for a given volume of traffic. A common network adage reflects this cost difference: "Switch where you can; route where you must."

When routers forward incoming packets closer to their destination hosts, this is routing.

Test Your Understanding

2. a) What are interfaces? b) Explain the network adage "Switch where you can; route where you must."

Hierarchical IPv4 Addressing

To understand the routing of IPv4 packets, it is necessary to understand IPv4 addresses. Chapter 1 showed that IP Version 4 (IPv4) addresses are 32 bits long. However, IPv4 addresses are not simple 32-bit strings. They have internal structure, and this internal structure is important in routing.

Single Networks versus "Networks" on the Internet To understand IPv4 addressing, you need to understand what the term "network" means on the Internet (Figure 8-4). It does not mean a single network, like an Ethernet network. Rather, a network on the Internet is a collection of routers and data links owned by a

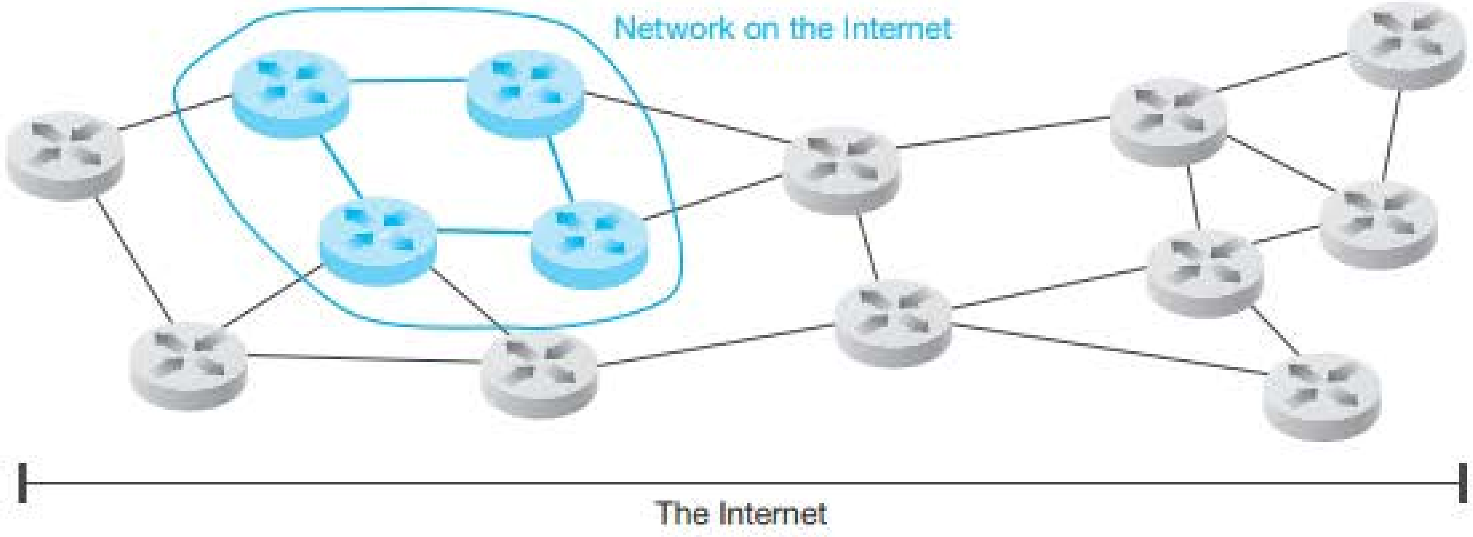

On the Internet, "network" does not mean a single network like Ethernet.
Rather, "Network" is an organizational concept.
It means the routers and data links owned and managed by a recognized organization.

FIGURE 8-4 "Network" on the Internet

recognized organization. Your home network is not a recognized network. The University of Hawai`i network is a recognized organization. So is Amazon.com. Both are end-user organizations. Internet service providers (ISPs) are also recognized organizations. ISPs are not end-user organizations.

On the Internet, "network" does not mean a single network such as an Ethernet LAN. It is an organizational concept. It means the routers and switches owned by a recognized organization, which is an end-user organization or an ISP.

Hierarchical Addressing As Figure 8-5 shows, IPv4 addresses are **hierarchical**. They consist of three parts (groups of bits) that locate a host in progressively smaller

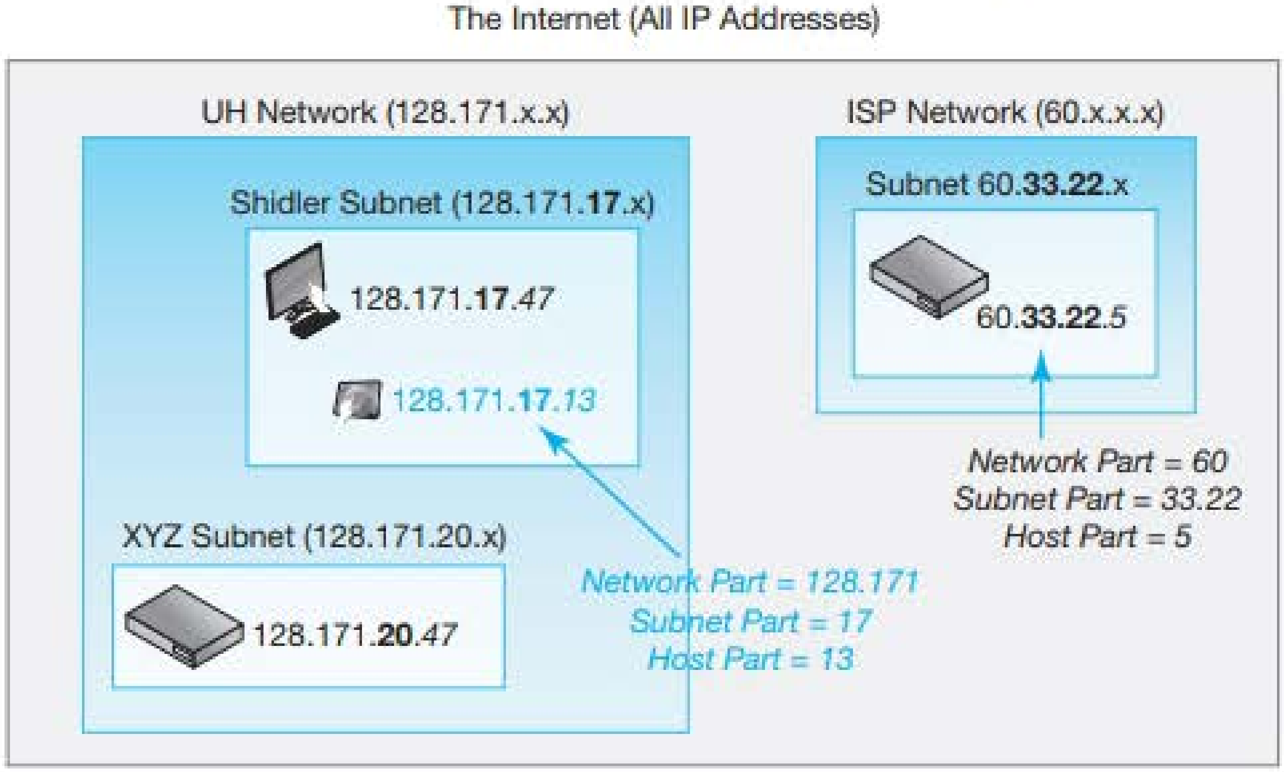

FIGURE 8-5 Hierarchical IPv4 Addresses

parts of the Internet. These are the network, subnet, and host parts. We will see later in this chapter how hierarchical IPv4 addressing simplifies routing tables. (Our examples are IPv4 addresses, but IPv6 routing works the same way but with 128-bit IPv6 addresses and masks.)

In IPv4 addressing, a part is a group of bits within the IPv4 address.

Network Part First, every IPv4 address has a **network part**, which identifies the host's recognized network on the Internet. In Figure 8-5, the network part for the University of Hawai`i Network is 128.171. All host IPv4 addresses in the University of Hawai`i Network (UH Network) begin with the network part 128.171. This is two IP 8-bit address segments. Therefore, UH Network's part is 16 bits long.

Do not get hung up on the network part being 16 bits. The UH Network is only an example. Different organizations have different network parts that range from 8 to 24 bits in length. For example, Figure 8-3 shows an ISP network with the network part 60. This network part is 8 bits long, not 16 bits.

Do not get hung up on the network part being 16 bits. This is only an example. Different organizations have different network parts that range from 8 to 24 bits in length.

Subnet Part Most large organizations further divide their networks into smaller units called **subnets**. After the network part in an IPv4 address come the bits of the **subnet part**. The subnet part bits specify a particular subnet within the network.

For instance, Figure 8-5 shows that in the IPv4 address 128.171.17.13, the first 16 bits (128.171) correspond to the network part, and the next 8 bits (17) correspond to a subnet on this network. (Subnet 17 is the Shidler College of Business subnet within the University of Hawai`i Network.) All host IPv4 addresses within this subnet begin with 128.171.17.

Again, do not get hung up on the subnet part being 8 bits long. In different organizations, subnet lengths vary widely. Keep clear in your head that the UH Network is only being used as an example. For the ISP shown in the figure, in fact, the subnet part is 16 bits long rather than 8 bits long.

Again, do not get hung up on the subnet part being 8 bits long. In different organizations, subnet lengths vary widely. Keep clear in your head that the UH Network is only being used as an example.

Host Part The remaining bits in the 32-bit IPv4 address constitute the **host part**, which specifies a particular host in a subnet. In Figure 8-5, the host part of the UH Network host is 8 bits long with a segment value of 13. This corresponds to a particular host, 128.171.17.13, on the Shidler College of Business subnet of the University of Hawai`i Network. Again, host parts in different organizations differ in length.

Variable Part Lengths Can you tell just by looking at an IPv4 address which bits correspond to the network, subnet, and host parts? The answer is no.

- For instance, if you see the IPv4 address 60.47.7.23, you may have an 8-bit network part of 60, an 8-bit subnet part of 47, and a 16-bit host part of 7.23.
- Or, you may have a network part of 16 bits, a subnet part of 8 bits, and a host part of 8 bits.
- In fact, parts may not even break conveniently at 8-bit boundaries. You may have a network part of 20 bits, a subnet part of 12 bits, and a host part of 12 bits.
- The only thing you can tell when looking at an IPv4 address is that it is 32 bits long.

Hierarchical IPv6 Address IPv6 addresses are also hierarchical and consist of three parts that are similar to those of IPv4 addresses. However, there are differences between IPv4 parts and IPv6 parts, and to discuss these, we need a better understanding of IPv6. We will look at hierarchical IPv6 addresses in the next chapter.

Test Your Understanding

3. a) What are the three parts of an IPv4 address? b) How long is each part? c) What is the total length of an IPv4 address? d) In the IPv4 address, 10.11.13.13, what is the network part? e) If you see an IPv4 address, what do you know for certain?

Routers, Networks, and Subnets

Border Routers Connect Different Networks As Figure 8-6 illustrates, networks and subnets are very important in router operation. Here we see a simple site internet. The figure shows that a **border router**'s main job is to connect different networks. This border router connects the 192.168.x.x network within the firm to the 60.x.x.x network of the firm's Internet service provider. Here, the *xs* are the remaining bits of the IPv4 address, so 192.168 and 60 are the network parts of the two networks.

A border router's main job is to connect different networks.

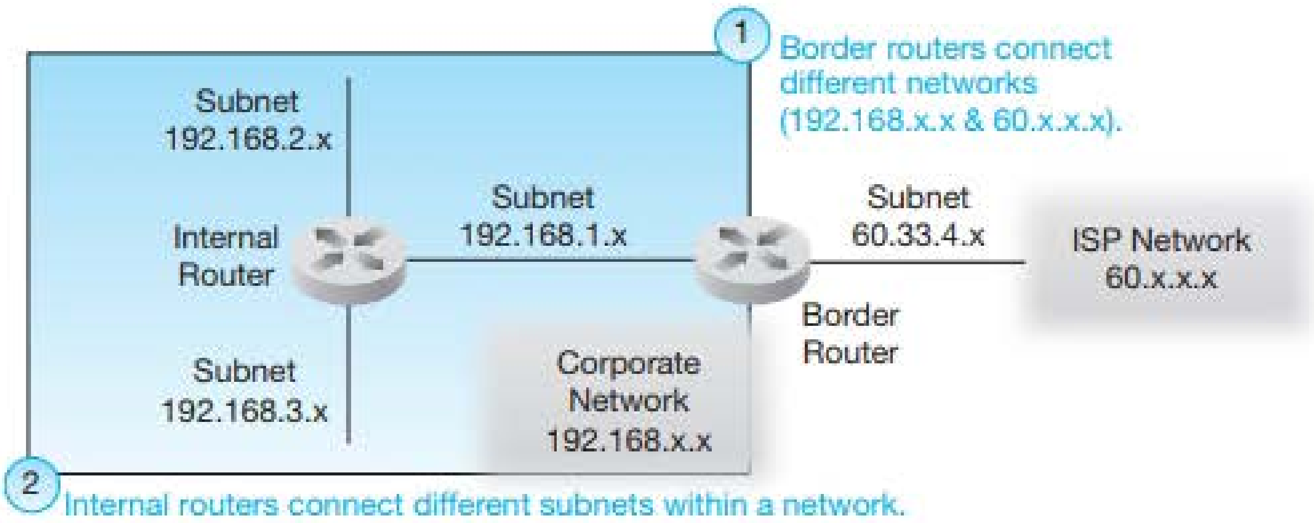
1
Border routers connect different networks (192.168.x.x & 60.x.x.x).
Subnet 192.168.2.x
Subnet 192.168.1.x
Subnet 60.33.4.x
Internal Router
ISP Network 60.x.x.x
Border Router
Subnet 192.168.3.x
Corporate Network 192.168.x.x
2
Internal routers connect different subnets within a network.

FIGURE 8-6 Border Routers, Networks, and Subnets

Internal Routers Connect Different Subnets The site network also has an internal router. An **internal router**, Figure 8-6 demonstrates, only connects different subnets within a network—in this case, the 192.168.1.x, 192.168.2.x, and 192.168.3.x subnets. Many sites have multiple internal routers to link the site's subnets.

An internal router only connects different subnets within a network.

Test Your Understanding

4. a) Connecting different networks is the main job of what type of router? b) What type of router connects different subnets?

Network and Subnet Masks

We have seen that in the University of Hawai`i network, the first 16 bits in IPv4 addresses are the network part, the next 8 are the subnet part, and the final 8 are the host part. However, because the sizes of the network, subnet, and host parts differ, routers need a way to tell the sizes of key parts. The tools that allow them to do this are masks.

32-Bit Strings Figure 8-8 illustrates how masks work. An **IPv4 mask** is a string of 32 bits, like an IPv4 address. However, a mask always begins with a series of 1s; this is always followed by a series of 0s. The total length of an IPv4 mask is always 32 bits, so if a mask begins with twelve 1s, it will end with twenty 0s.

There are two kinds of masks.

- In a network mask, the *bits in the network part of the mask are* 1s, and the remaining bits are 0s.
- In a subnet mask, the *bits of both the network and the subnet parts are 1s*, and the remaining bits are 0s.[1]

We have seen that the University of Hawai`i network part is 16 bits and the subnet part is 8 bits.

- So the network mask will have sixteen 1s followed by sixteen 0s.
- The subnet mask will have twenty-four 1s followed by eight 0s.

An IPv4 mask is 32 bits long

It begins with a series of 1s

The remaining bits are 0s

Example (Broken into octets for readability): 11111111 11110000 00000000 00000000

Prefix notation (the number of leading 1s) for this example: /12

FIGURE 8-7 IPv4 Network Mask (Study Figure)

[1] To give an analogy, to specify a state in the United States, you only need to give the name of the state. "Oklahoma" is sufficient to identify that state. For cities, you need to specify both a city and a state. There is a Portland in both Oregon and Maine, so you need to say "Portland, Oregon" to specify that city. The network part bits correspond to the state, the subnet parts to a city, so a subnet mask needs 1s in both the network and subnet parts.

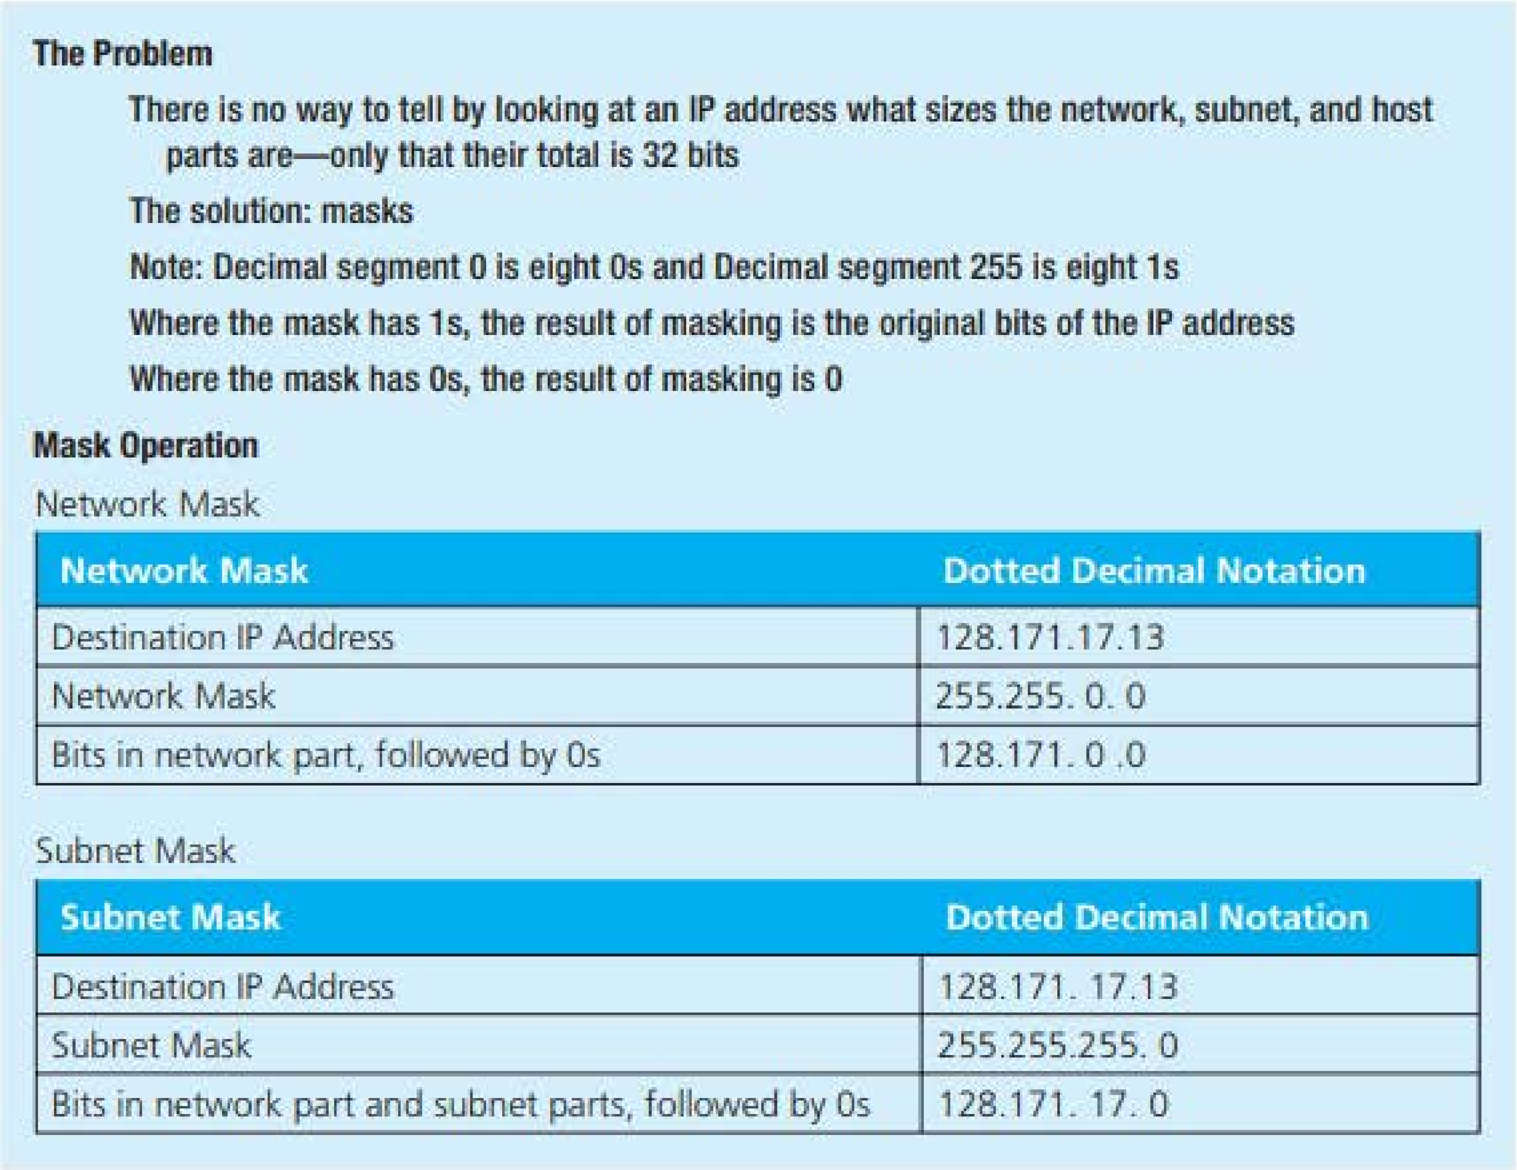

The Problem

There is no way to tell by looking at an IP address what sizes the network, subnet, and host parts are—only that their total is 32 bits

The solution: masks

Note: Decimal segment 0 is eight 0s and Decimal segment 255 is eight 1s

Where the mask has 1s, the result of masking is the original bits of the IP address

Where the mask has 0s, the result of masking is 0

Mask Operation

Network Mask

Network Mask	Dotted Decimal Notation
Destination IP Address	128.171.17.13
Network Mask	255.255. 0. 0
Bits in network part, followed by 0s	128.171. 0 .0

Subnet Mask

Subnet Mask	Dotted Decimal Notation
Destination IP Address	128.171. 17.13
Subnet Mask	255.255.255. 0
Bits in network part and subnet parts, followed by 0s	128.171. 17. 0

FIGURE 8-8 IP Networks and Subnet Masks

For example, suppose that the mask is 255.255.0.0. This means that the four 8-bit segments of the mask have the values 255, 255, 0, and 0. In dotted decimal notation eight 1s is 255 and eight 0s is 0. Therefore, the four segments have, in order, eight 1s, eight 1s, eight 0s, and eight 0s. Putting this together, the mask has sixteen 1s followed by sixteen 0s.

Prefix Notation for Masks Writing 255.255.255.0 is not very difficult, but networking professionals often use a shortcut called prefix notation. The mask 255.255.255.0 is twenty-four 1s followed by eight 0s. In prefix notation, this mask is represented as /24. Do you see the pattern? In **prefix notation**, a mask is represented by a slash followed by the number of initial 1s in the mask. What about 255.0.0.0? Yes, it is /8. Prefix notation is simpler to write than dotted decimal notation. (By the way, we call this prefix notation because it focuses on the first part of the mask—the part that is all 1s.)

In prefix notation, a mask is represented by a slash followed by the number of initial 1s in the mask.

Another advantage of prefix notation for a mask is that it is simple even if the number of leading 1s is not a multiple of eight. For example, suppose that the mask is eighteen 1s followed by fourteen 0s. The mask in prefix notation is obviously /18. What if you saw this mask in dotted decimal notation: 255.255.48.0? The first two octets are obviously all 1s. However, you would need your decimal-to-binary calculator to figure out that 48 is 110000.

Masking IPv4 Addresses Figure 8-8 shows what happens when a mask is applied to an IPv4 address, 128.171.17.13. The mask is 255.255.0.0. Where the mask has 1s, the result is the original bits of the IPv4 address. There are sixteen 1s. This is two octets. So the first two octets of the result would be 128.171. For the remaining 16 bits, which are 0s, the result of the masking is 0. So the masking result is 128.171.0.0.

Network Masks Network masks, as noted earlier, have 1s in the network part and 0s for the remaining bits. If the network mask is 255.255.0.0 and the IPv4 address is 128.171.17.13, then the result of masking is 128.171.0.0. This tells us that 128.171 is the network part.

Subnet Masks For subnet masks, in turn, the initial 1s indicate the number of bits in *both* the network and subnet parts. Therefore, if 128.171 is the network part and 17 is the subnet part, then the subnet mask will be 255.255.255.0 (/24). If you mask 128.171.17.13 with /24, you get 128.171.17.0.[2]

Test Your Understanding

5. a) How many bits are there in an IPv4 mask? b) What do the 1s in an IPv4 network mask correspond to in IPv4 addresses? c) What do the 1s in an IPv4 subnet mask correspond to in IPv4 addresses? Think carefully! d) When a network mask is applied to any IPv4 address on the network, what is the result?

6. a) A mask has eight 1s, followed by 0s. Express this mask in dotted decimal notation. b) Express this mask in prefix notation. c) In prefix notation, a mask is /16. Express this mask in dotted decimal notation. d) Express the mask /18 in dotted decimal notation. (You will need a calculator for this.)

HOW ROUTERS PROCESS PACKETS

Switching versus Routing

In Chapter 5, we saw that Ethernet switching is very simple. Ethernet switches must be organized in a hierarchy. Therefore, there is only a single possible path between any two hosts across the network. When a frame arrives, there is only one possible port to send the frame back out. In an Ethernet switching table, each Ethernet address only appears in one row. This single row can be found quickly, so an Ethernet switch does little work per frame. This makes Ethernet switching fast and inexpensive.

[2] Why not make the network part 0s and the subnet part 1s instead of making both 1s? Think of a network as a state and a subnet as a city. In the United States, there are two major cities named Portland—one in Maine and the other in Oregon. You cannot just say "Portland" to designate a city. You must give both the state and city. Analogously, there may be many subnet parts with a value of 17, so you must give both the network and subnet parts to designate a specific subnet. Another way to look at it is that if you only had 1s in the subnet part of a subnet mask, you would break the rule that masks must have a number of leading 1s followed by a number of trailing 0s. (This repeats information in the previous footnote.)

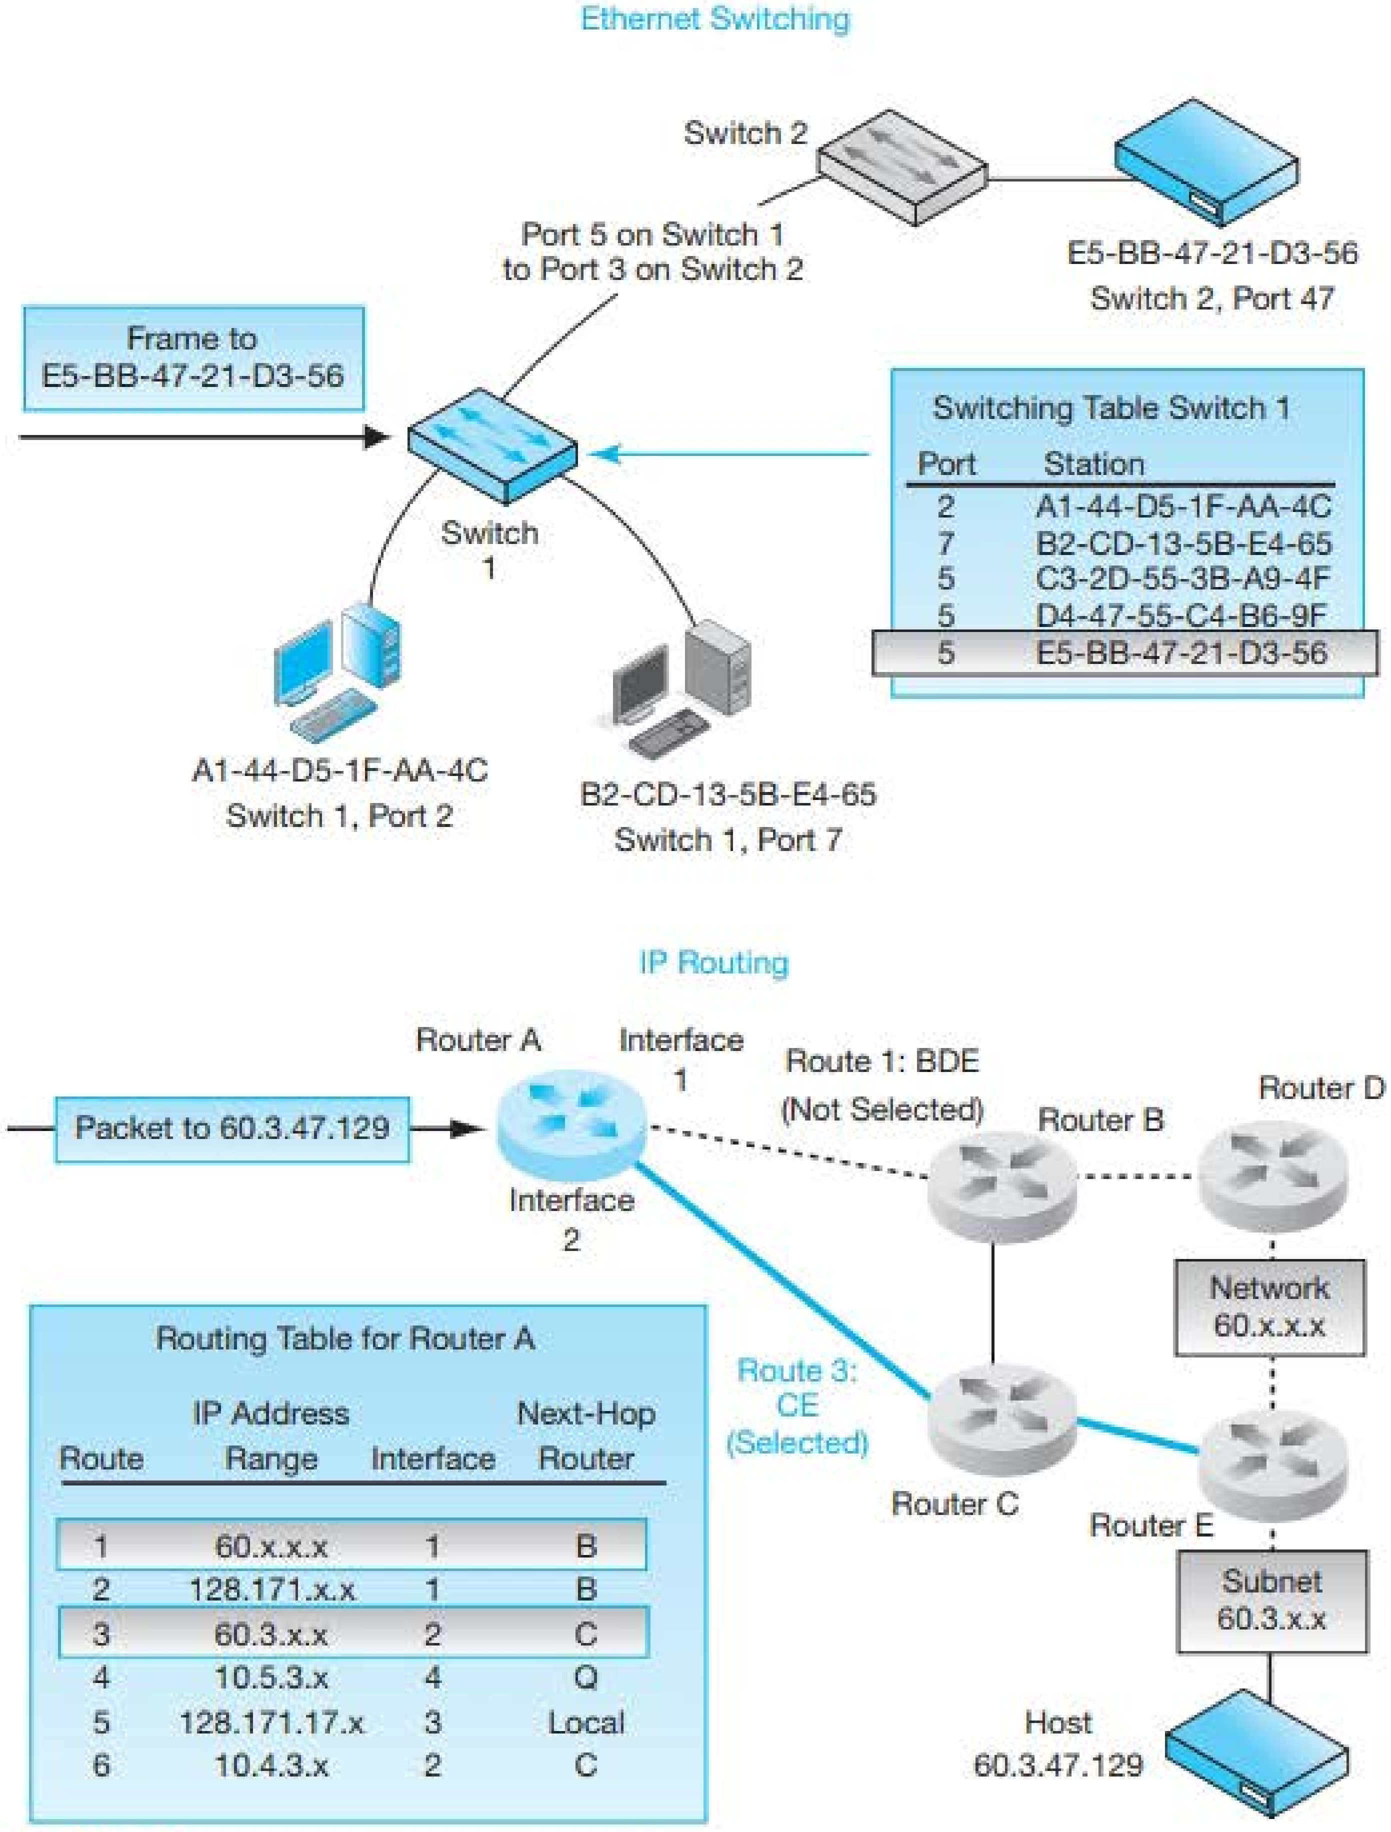

Switching Table Switch 1

Port	Station
2	A1-44-D5-1F-AA-4C
7	B2-CD-13-5B-E4-65
5	C3-2D-55-3B-A9-4F
5	D4-47-55-C4-B6-9F
5	E5-BB-47-21-D3-56

Routing Table for Router A

Route	IP Address Range	Interface	Next-Hop Router
1	60.x.x.x	1	B
2	128.171.x.x	1	B
3	60.3.x.x	2	C
4	10.5.3.x	4	Q
5	128.171.17.x	3	Local
6	10.4.3.x	2	C

FIGURE 8-9 Ethernet Switching versus IP Routing

In contrast, routers are organized in meshes. This gives more reliability because it allows many possible alternative routes between endpoints. However, in a mesh, there are multiple ways to send a packet back out to reach its destination. Figure 8-9 shows that in a routing table, several rows may match an IPv4 address. Row 1 calls for sending the packet out Interface 1 to Next-Hop Router B. Row 3, in turn, calls for sending the packet out Interface 2 to Next-Hop Router C.

Routing

Processing an individual packet and passing it closer to its destination host is called routing

The Routing Table

Each router has a routing table that it uses to make routing decisions

Routing tables have rows

Each row represents a route for a range of IP addresses—often packets going to the same network or subnet

The Routing Decision

1. Find all row matches
2. Find the best-match row
3. Send the frame out, based on information in the row

FIGURE 8-10 The Routing Process

The fact that a packet may be matched by multiple rows requires a fairly complex process to be performed on each packet. Figure 8-10 summarizes this process. To route a packet, a router must first find *all* rows that apply to an incoming packet. In fact, it will have to look at every row in the table to see if it is a match to the packet's destination IPv4 address. It must then pick the best alternative route from this list of matches. All of this requires quite a bit of work per packet, making routing much more expensive than switching per message handled.

Test Your Understanding

7. Why are routing tables more complex than Ethernet switching tables? Give a detailed answer.

Routing Table

Figure 8-11 shows a routing table. We will see how a router uses its rows and columns to make the routing decision—what to do with an arriving packet.

Rows Are Routes for All IPv4 Addresses in a Range

In the routing table, each row represents a route for all IPv4 addresses within a range of IPv4 addresses—typically addresses within a particular network or subnet. It does not specify the full route, however; it only specifies the next step in the route (either the next-hop router to handle the packet next or, on the last router, the destination host).

In the routing table, each row represents a route for all IPv4 addresses within a range of IPv4 addresses.

This is important because the routing table does not need a row for each IPv4 address as an Ethernet switching table does for EUI-48 addresses. It only needs a row for each *group* of IPv4 addresses. This means that a router needs many fewer rows than an Ethernet switch would need for the same number of addresses.

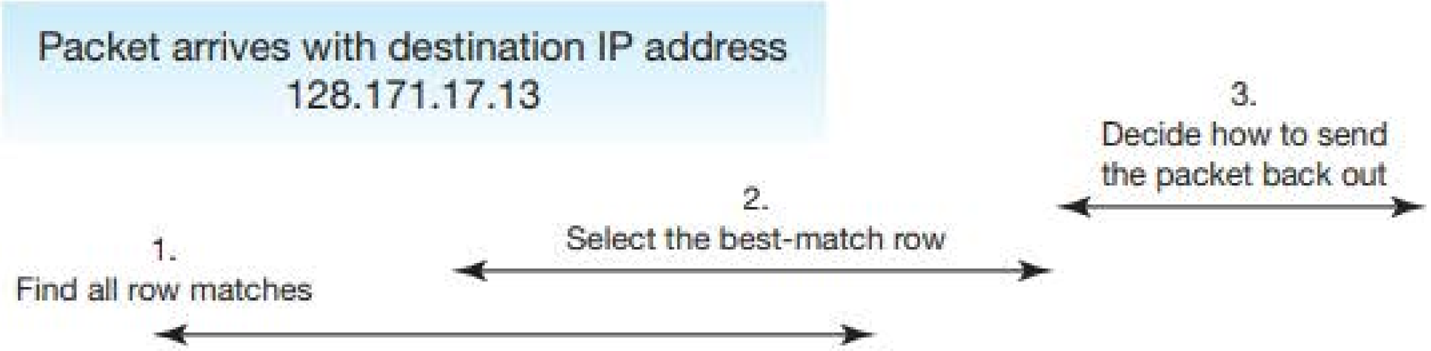

Row	Destination Network or Subnet	Mask (/Prefix)	Metric (Cost)	Interface	Next-Hop Router
1	128.171.0.0	255.255.0.0 (/16)	47	2	G
2	172.30.33.0	255.255.255.0 (/24)	0	1	Local
3	60.168.6.0	255.255.255.0 (/24)	12	2	G
4	123.0.0.0	255.0.0.0 (/8)	33	2	G
5	172.29.8.0	255.255.255.0 (/24)	34	1	F
6	172.40.6.0	255.255.255.0 (/24)	47	3	H
7	128.171.17.0	255.255.255.0 (/24)	55	3	H
8	172.29.8.0	255.255.255.0 (/24)	20	3	H
9	172.12.6.0	255.255.255.0 (/24)	23	1	F
10	172.30.12.0	255.255.255.0 (/24)	9	2	G
11	172.30.12.0	255.255.255.0 (/24)	3	3	H
12	60.168.0.0	255.255.0.0 (/16)	16	2	G
13	0.0.0.0	0.0.0.0 (/0)	5	3	H

FIGURE 8-11 Routing Table

However, there are many more IPv4 addresses on the Internet than there are Ethernet addresses on an Ethernet network. Even with rows representing groups of IPv4 addresses, core routers in the Internet backbone still have several hundred thousand rows. This is important. We will see that routers need to do calculations for *all* rows.

Test Your Understanding

8. a) In a routing table, what does a row represent? b) Do Ethernet switches have a row for each individual Ethernet address? c) Do routers have a row for each individual IPv4 address? d) What is the advantage of the answer to the previous subparts of this question?

Step 1: Finding All Row Matches

We will now see how the router uses its routing table to make routing decisions. Figure 8-12 shows that there are three very different steps. These differences can lead to confusion, so you must study this material carefully. The first step is to find

Step 1: Find All Row Matches

The router looks at the destination IP address in an arriving packet

For each row:

- Apply the row's mask to the destination IP address in the packet
- Compare the result with the row's destination value
- If the two match, the row is a match

The router must do this to ALL rows because there may be multiple matches

This step ends with a set of matching rows

Example 1: A Destination IP Address that IS in the Range

Destination IP Address of Arriving Packet	128.171.17.13
Apply the (Network) Mask	255.255.0.0
Result of Masking	128.171.0.0
Destination Column Value	128.171.0.0
Does Destination Match the Masking Result?	Yes
Conclusion	Row is a match

Example 2: A Destination IP Address that is NOT in the Destination Range

Destination IP Address of Arriving Packet	60.43.7.8
Apply the (Network) Mask	255.255.0.0
Result of Masking	60.43.0.0
Destination Column Value	128.171.0.0
Does Destination Match the Masking Result?	No
Conclusion	Not a match

Step 2: Find the Best-Match Row

The router examines the matching rows it found in Step 1 to find the best-match row

Basic rule (always used): It selects the row with the longest match (Initial 1s in the row mask) If it finds one, there is no need to go on to the tie-breaker rule

Tie Breaker (only when needed): If there is a tie on longest match, select among the tie rows based on a metric

- For cost metric, choose the row with the lowest metric value
- For speed metric, choose the row with the highest metric value

The router now knows the best-match row

Step 3: Send the Packet Back Out

Send the packet out the interface (router port) designated in the best-match row

If the address says Local, the destination host is on that interface

Sends the packet to the destination IP address in a frame

FIGURE 8-12 Steps in a Routing Decision

which of the rows in the routing table match the destination IPv4 address in an arriving packet. Due to the existence of alternative routes in a router mesh, most packets will match more than one row.

Row Number Column The first column in Figure 8-11 contains route (row) numbers. Routing tables actually do not have this column. We include it to allow us to refer to specific rows in our discussion. Again, each row specifies a route to a destination.

Row Matches How does the router know which IPv4 addresses match a row? The answer is that it uses the *Destination Network or Subnet* column and the *Mask* columns.

Suppose that all IPv4 addresses in the University of Hawai`i network should match a row. The mask would be the network mask 255.255.0.0, because the UH Network has a 16-bit network part. If this mask is applied to any UH address, the result will be 128.171.0.0. This is the value that will be in the destination column. In fact, this matches Row 1 in Figure 8-11.

Let's see how routers use these two columns in Figure 8-11. We will use two examples. This is the first:

- Suppose that a packet arrives with the IPv4 address 128.171.17.13. The router will look first at Row 1.
- In this row, the router applies the mask 255.255.0.0 to the arriving packet's destination IPv4 address, 128.171.17.13. The result is 128.171.0.0.
- Next, the router compares the masking result, 128.171.0.0, to the destination value in the row, 128.171.0.0. The two are the same, so the row is a match.

Here is the second example.

- This time, the destination IPv4 address in the arriving packet is 60.43.6.8.
- Again, the router applies the mask 255.255.0.0 in Row 1 to the destination IPv4 address, 60.43.6.8. The result is 60.43.0.0.
- Next, the router compares 60.43.0.0 to the destination value in the row, 128.171.0.0. The two are not equal. Therefore, the row is not a match.

Mask and Compare This may seem like an odd way to see if a row matches the arriving IPv4 address. A human can simply look at 60.43.7.8 and see that it does not match 128.171.0.0. However, routers do not possess human pattern-matching abilities.

On the other hand, routers (and all computers) have specialized circuitry for doing masking and comparing—the two operations that row matching requires. Thanks to this specialized circuitry, routers can blaze through hundreds of thousands of rows in a tiny fraction of a second.

In contrast, if a human saw 300,000 rows, finding matches visually would take a long time. For finding matches, stupid but fast beats smart and slow.

The Default Row The last row in Figure 8-11 has the destination 0.0.0.0 and the mask 0.0.0.0. This row will match *every* IPv4 address because masking any IPv4 address with 0.0.0.0 will give 0.0.0.0, which is the value in the Destination Field of Row 13. This row ensures that at least one row will match the destination IPv4 address of every arriving packet. It is called the **default row**. In general, a "default" is something you get if you do not have a more specific choice.

In general, a "default" is something you get if you do not have a more specific choice.

The Need to Look at *All* Rows Thanks to their mesh topology, internets have many alternative routes. Consequently, a router cannot stop the first time it finds a row match for each arriving packet because there may be a better match further on. A router must look at each and every row in the routing table to see which rows match. So far, we have seen what the router does in Row 1 of Figure 8-11. The router then goes on to Row 2 to see if it is a match by masking and comparing. After this, it goes on to Row 3, Row 4, Row 5, and so on, all the way to the final row (Row 13 in Figure 8-11).

Test Your Understanding

9. a) In Figure 8-11, how will a router test whether Row 3 matches the IPv4 address 60.168.6.7? Show the calculations in the format given in Figure 8-12. b) Is the row a match? c) Why is the last row called the default row? d) Why must a router look at all rows in a routing table? e) Which rows in Figure 8-11 match 172.30.17.6? (Don't forget the default row.) Show your calculations for rows that match. f) Which rows match 60.168.7.32? Show your calculations for rows that match. g) Which rows in Figure 8-11 match 128.171.17.13? (Show your calculations for rows that match.)

Step 2: Selecting the Best-Match Row

List of Matching Rows At the end of Step 1, the router has a list of matching rows. For a packet with the destination IPv4 address 128.171.17.13, three rows in Figure 8-11 match.

- The first is Row 1, as we have already seen.
- The second is Row 7, with a destination of 128.171.17.0 and a mask of 255.255.255.0.
- Finally, the default row (Row 13 in this figure) will always be a match.

From these, the router must select the best-match row, the row that represents the best route for an IPv4 address.

Basic Rule (Always Used): Longest Match How does the router decide whether to follow Row 1, Row 7, or Row 13? The answer is that it follows the rule of selecting the **longest match** (the longest number of initial 1s in the mask). Row 1 has a mask of 255.255.0.0, which means that it has a 16-bit match. Row 7, in turn, has the prefix /24, meaning that it has a 24-bit match. Row 13 has a prefix of 0/. Row 7 has the longest match, so the router selects Row 7 as the best match.[3]

By the way, note that the default row always has a prefix of 0/. This is the shortest possible length of match. Consequently, if *any* other row matches, its length of match will be longer, and the default row will never be chosen as the best-match row.

Tie-Breaker Rule (Only When Needed): Best Metric Value In the previous example, there was a winner for longest match. There was no need to handle a tie. However,

[3] Why the longest match rule? The answer is that the closer a route gets a packet to the destination IPv4 address, the better. Row 1 only gets the packet to the UH Network, 128.171.x.x, whereas Row 7 gets the packet all the way to the Shidler College of Business subnet of the University of Hawai`i, 128.171.17.x. This is the subnet that contains host 128.171.17.13.

what if there is a tie instead of a win? For instance, the destination IPv4 address 172.29.8.112 matches both Row 5 and Row 8 in Figure 8-11. Both have a match length of 24 bits—a tie.

In case of a tie for longest match, the tie-breaker rule is to use the **metric** column, which describes the desirability of a route. For instance, in Figure 8-11, the metric is cost. Row 5 has a cost of 34, whereas Row 8 has a cost of 20. *Lower cost is better than higher cost*, so the router selects Row 8.

In this case, the row with the *lowest* metric won. However, what would have happened if the metric had been *speed* instead of cost? *More speed is better*, so the router would choose Row 5, with the *higher* speed (34).

Test Your Understanding

10. a) Distinguish between Step 1 and Step 2 in the routing process. b) If any row other than the default row matches an IPv4 address, why will the router never choose the default row? c) Which rows in Figure 8-11 match 128.171.17.13? (Don't forget the default row.) Show your calculations for rows that match. d) Which of these is the best-match row? Justify your answer. e) What rows match 172.40.17.6? Show your calculations for rows that match. f) Which of these is the best-match row? Justify your answer. g) Which rows match 172.30.12.47? Show your calculations for rows that match. h) Which of these is the best-match row? Justify your answer. i) How would your previous answer change if the metric had been reliability?

Step 3: Sending the Packet Back Out

In Step 1, the router found all rows that matched the destination IPv4 address of the arriving packet. In Step 2, it found the best-match row. Finally, in Step 3, the router sends the packet back out.

Interface Recall that router ports are called interfaces. The fifth column in Figure 8-11 is the interface number. If a router selects a row as the best match, the router sends the packet out the interface designated in that row. If Row 1 is selected, the router will send the packet out Interface 2.

Next-Hop Router In a switch, a port connects directly to another switch or to a computer. However, a router interface connects to an entire subnet or network. Therefore, it is not enough to select an interface to send the packet out. It is also necessary to specify *a particular device* on the subnet.

In most cases, the router will send the packet on to another router, called the **next-hop router**. The next-hop router column specifies the router that should receive the packet. It will then be up to that next-hop router to decide what to do next. In Figure 8-11, the next-hop router value in Row 1 is G.[4] The default row's next-hop-router is H. This router is called the **default router**, and any packet not matching a specific row other than the default row will be sent to Router H.

In some cases, however, the destination host itself will be on the subnet out a particular interface. In that case, there is no reason to send the packet on to another

[4] In an actual router, this column would have the IPv4 address of Router H, rather than its name. However, we include the letter designation rather than the IPv4 address for ease of understanding.

router. Instead, the router will send the packet directly to the destination host. To indicate that the next destination is the destination host, the Next-Hop Router column will say *local*.

Test Your Understanding

11. a) Distinguish between Step 2 and Step 3 in routing. b) What are router ports called? c) If the router selects Row 13 as the best-match row, what interface will the router send the packet out? d) To what device? e) Why is this router called the default router? (The answer is not in the text.) f) If the router selects Row 2 as the best-match row for packet 172.30.33.6, what interface will the router send the packet out? g) To what device? (Don't say, "the local device.")

Cheating (Decision Caching)

We have discussed what happens when a packet arrives at a router. However, what will the router do if the next packet has the same destination IPv4 address? The answer is that the router *should* go through the entire process again. Even if a thousand packets arrive that are going to the same destination IPv4 address, the router should go through the entire three-step process for each of them.

As you might expect, a router might cheat or, as it is euphemistically named, cache (remember) the decision it made for a destination IPv4 address. It will then use this decision for successive IPv4 packets going to the same destination. Using a **decision cache** greatly reduces the work that a router will do for each successive packet to the same destination address (Figure 8-13).

Decision caching is not in the Internet Protocol. This is because it is not entirely safe. The Internet changes constantly as routers come and go and as links between routers change. Consequently, a cached decision that is used for too long will result in non-optimal routing or even routes that will not work and that will effectively send packets into a black hole.

Test Your Understanding

12. a) What should a router do if it receives several packets going to the same destination IPv4 address? b) How would decision caching speed the routing decision for packets after the first one? c) Why is decision caching dangerous?

Standard Routing	Decision Caching
If another packet arrives with the same destination IP address,	If another packet arrives with the same destination IP address,
Go through the entire process again.	Do what was done the last time.
Go through every row looking for matches.	
Find the best match row.	
Send the packet out the indicated interface.	

FIGURE 8-13 Standard Routing versus Decision Caching

Routing Tables for IPv6 Addresses

Routing tables for IPv6 addresses have the same columns that routing tables for IPv4 addresses have. However, the destination address in an arriving packet is a 128-bit IPv6 address, the mask is 128 bits long, and the destination network and subnet address value is 128 bits long. However, we have not looked at part lengths in hierarchical IPv6 addresses, so we cannot discuss routing tables for IPv6 addresses yet.

IN MORE DEPTH

Masking When Masks Do Not Break at 8-Bit Boundaries

All the masks we have seen up to this point have had their parts broken at 8-bit segment boundaries. For example, at the University of Hawai`i, the network part is 16 bits long, which corresponds to two segments (128.171), the subnet part is 8 bits long (17), and the host part is 8 bits long (13). All the masks in Figure 8-11 also break at 8-bit segment boundaries.

Masks that break at 8-bit boundaries are easy for humans to read. In general, you can look at a mask in the table and decide if it matches a particular IPv4 address. For instance, if the mask is 255.255.0.0 (/16), and if the destination column value is 128.171.0.0, this definitely matches the IPv4 address 128.171.45.230.

However, masks do not always break at 8-bit boundaries. For example, suppose that a row in the routing table has the destination address 3.143.12.12 and the mask is 255.248.0.0. Will the IPv4 address 3.143.12.12 match this row? At first glance, this certainly does not seem to be a match. However, it is. While it would be nice to always break IPv4 address parts into 8-bit boundaries, companies have no control over the size of their network parts, and these usually vary from 8 bits to 22 bits.

To see why the IPv4 address and destination match in this example, look at Figure 8-14. This figure shows the matching analysis when the binary representations are given for each segment. Follow the masking and you will see that the result is a match. When a mask does not break at an 8-bit boundary, you must go back to the raw 32-bit IPv4 address, mask, and destination.

Test Your Understanding

13. An arriving packet has the destination IPv4 address 128.171.180.13. Row 86 has the destination value 128.171.160.0. The row's mask is 255.255.224.0. Does this row match the destination IPv4 address? Show your work. You can use the Windows Calculator if you have a Windows PC. In Windows Vista and earlier versions of Windows, choose scientific when you open the calculator. In the Windows 7 and Windows 10 calculator, choose programmer mode.

	Dotted Decimal Notation	Segment 1	Segment 2	Segment 3	Segment 4
IPv4 address	3.143.12.12	00000011	10001111	00001100	00001100
Mask	255.248.0.0	11111111	11111000	00000000	00000000
Result	3.136.0.0	00000011	10001000	00000000	00000000
Destination	3.136.0.0	00000011	10001000	00000000	00000000
Result and Destination Match?	Yes	Yes	Yes	Yes	Yes

FIGURE 8-14 Using a Mask Whose 1s Do Not Break Down at an 8-Bit Boundary

THE INTERNET PROTOCOL VERSION 4 (IPv4) FIELDS

We have focused on IP routing. However, the Internet Protocol has other properties that networking professionals need to understand.

As noted in Chapter 1, most traffic on the Internet and private internets today is governed by the IP Version 4 standard. (There were no versions 0 through 3.) We looked at the header checksum, the source IPv4 address, and the destination IPv4 address in the first two chapters. Now we will look at the other fields in the IPv4 header.

The First Row

Figure 8-15 shows the IPv4 packet. Its first four bits constitute the **Version Number Field**. This field has the value 0100 (binary for 4). This indicates that this is an IPv4 packet. The next field gives the header length, and the last field on the first row gives the total length of the packet.[5]

Between the header and total length fields, two fields govern transmission quality. The **Differentiated Services Control Point** Field can be used for priority or other quality of service purposes. The **Explicit Congestion Notification (ECN)** Field can be used to reduce the transmission frequency between a pair of hosts to cope with congestion in the transmission system between them.

Test Your Understanding

14. a) What is the main version of the Internet Protocol in use today? b) Which field can be used to specify quality of service? c) How can the ECN Field be used?

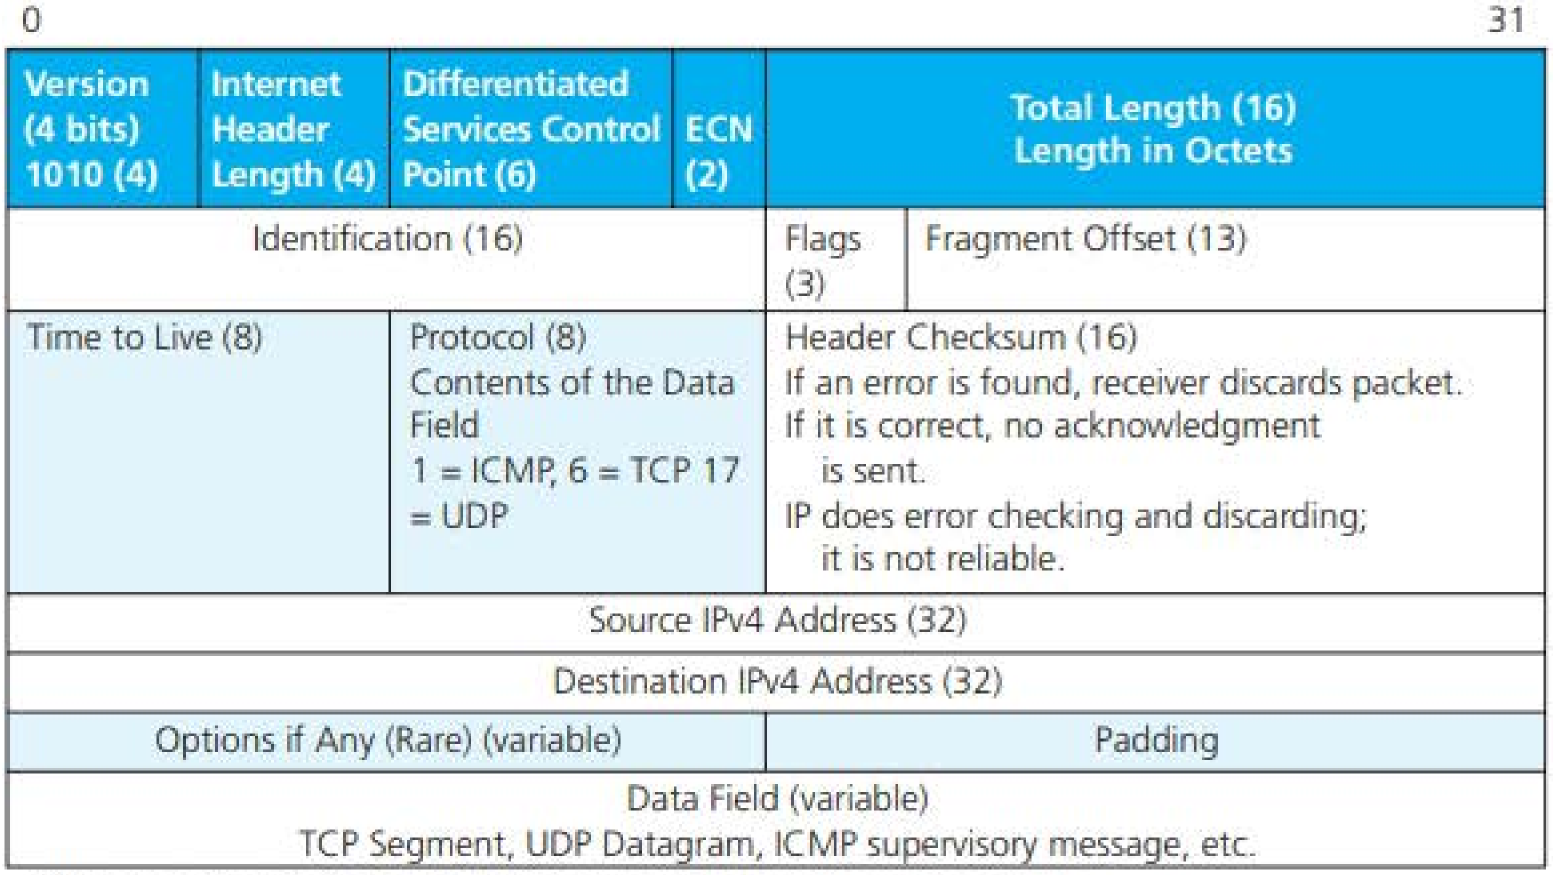

0 31

Version (4 bits) 1010 (4) | Internet Header Length (4) | Differentiated Services Control Point (6) | ECN (2) | Total Length (16) Length in Octets

Identification (16) | Flags (3) | Fragment Offset (13)

Time to Live (8) | Protocol (8) Contents of the Data Field 1 = ICMP, 6 = TCP 17 = UDP | Header Checksum (16) If an error is found, receiver discards packet. If it is correct, no acknowledgment is sent. IP does error checking and discarding; it is not reliable.

Source IPv4 Address (32)

Destination IPv4 Address (32)

Options if Any (Rare) (variable) | Padding

Data Field (variable) TCP Segment, UDP Datagram, ICMP supervisory message, etc.

Differentiated Services Control Point: To request special services, such as priority.
ECN is Explicit Congestion Notification. To notify the receiver of congestion along the route.

FIGURE 8-15 IP Version 4 (IPv4) Packet Syntax

[5] The header length field gives the length of the header in 32-bit units. The length field gives the total length of the IPv4 packet in octets.

The Second Row

TCP fragments application messages and sends them in individual packets. This has benefits that we saw in Chapters 1 and 2. When IPv4 was created, it was decided to allow routers to further fragment packets. Although this seemed like a good idea at the time, it led to many problems. Today, operating systems by default tell routers not to fragment IPv4 packets. When IPv6 was developed, packet fragmentation was not allowed at all. The second row has information that the destination host uses to reassemble fragmented packets. Given the unimportance of IPv4 packet fragmentation, we will ignore the fields in this row. It is about as useful as the human appendix, often a burst appendix at that.

Test Your Understanding

15. a) Distinguish between application message fragmentation and packet fragmentation. b) Under what circumstances would the identification, flags, and fragment offset fields be used in IP? c) Why did we not study them in detail? d) Does IPv6 allow packet fragmentation?

The Third Row

IP Time to Live (TTL) Field In the early days of the ARPANET, which was the precursor to the Internet, packets that were misaddressed would theoretically circulate endlessly among packet switches in search of their nonexistent destinations. To prevent this, IP added an ominous-sounding **Time to Live (TTL) Field** that is assigned a value by the source host. Different operating systems have different TTL defaults. Most insert the TTL value 128. Each router along the way decrements (decreases) the TTL Field by 1 when a packet arrives before going through the routing process. A router decrementing the TTL to 0 will discard the packet.

IP Protocol Field The **Protocol Field** reveals the contents of the Data Field. TCP and UDP have protocol values 6 and 17, respectively.

If the Protocol Field value is 1, the IPv4 packet carries an Internet Control Message Protocol (ICMP) message in its Data Field. As we will see later in the next chapter, IP is a lean mean routing machine with no time for supervisory messages. ICMP is TCP/IP's tool for carrying internet layer supervisory messages. After decapsulation, the internet layer process must pass the contents of the packet's Data Field up to another process.

The Protocol Field value tells the receiver which process should receive these contents. If the Protocol Field's value is 1, then the internet process will pass the contents of the Data Field to the ICMP process because these contents are an ICMP message.

Test Your Understanding

16. a) What does a router do if it receives a packet with a TTL value of 2? b) What does the next router do? c) What does the Protocol Field value tell the destination host? d) What will the destination internet process do if it sees 17 in the Data Field?

IP Options

The IPv4 header allows options. There are several possible options, and they may come in any order. Some are only read by the destination host. However, a lack of required order means that each router must look at every option to see if it applies. This is time consuming.

Test Your Understanding

17. What problem is caused by the way that IPv4 handles options?

IP VERSION 6 (IPv6)

Outgrowing IPv4

Although IPv4 continues to dominate the Internet's traffic, the Internet Assigned Numbers Authority (IANA) initially did a poor job distributing IPv4 addresses. Today, there are no more to distribute. Yet new devices like mobile phones are exploding, and each needs its own IPv6 address. This is forcing more organizations to use IPv6 addresses. Today, all firms must support IPv6.

The most fundamental change in IPv6 is the move from 32-bit addresses to 128-bit addresses. This does not produce merely four times as many addresses. Each additional bit *doubles* the number of addresses. So while there are just under 4.3 billion (4.3×10^9) IPv4 addresses, there are 3.4×10^{38} IPv6 addresses—34 undecillion. To put this in perspective, there are about seven billion people in the world today. For each person, there are 5×10^{28} IPv6 addresses. Even with the Internet of Things, IPv6 will "solve" the address availability problem for many years to come.

Test Your Understanding

18. a) What is the main problem with IPv4 that IPv6 was created to solve? b) How does IPv6 solve this problem?

IPv6

In its 1994 meeting, the IETF decided to create a new version of the Internet Protocol. The IETF called this new version **IP Version 6 (IPv6)**.[6] Over the next few years, the IPv6 standards family grew and matured. It was soon ready to be used, and many networking and computer vendors began to build IPv6 into their products.

Organizations soon found that using these new equipment capabilities, however, was a great deal more work than simply turning them on. For many years, few organizations saw the need to make the expensive upgrade to IPv6 because they had enough addresses. In addition, we will see in Chapter 9 how Network Address Translation (NAT) greatly extended the use of existing IPv4 addresses in firms, at the cost of some

[6] The IETF did define an Internet Protocol Version 5, but it was never implemented.

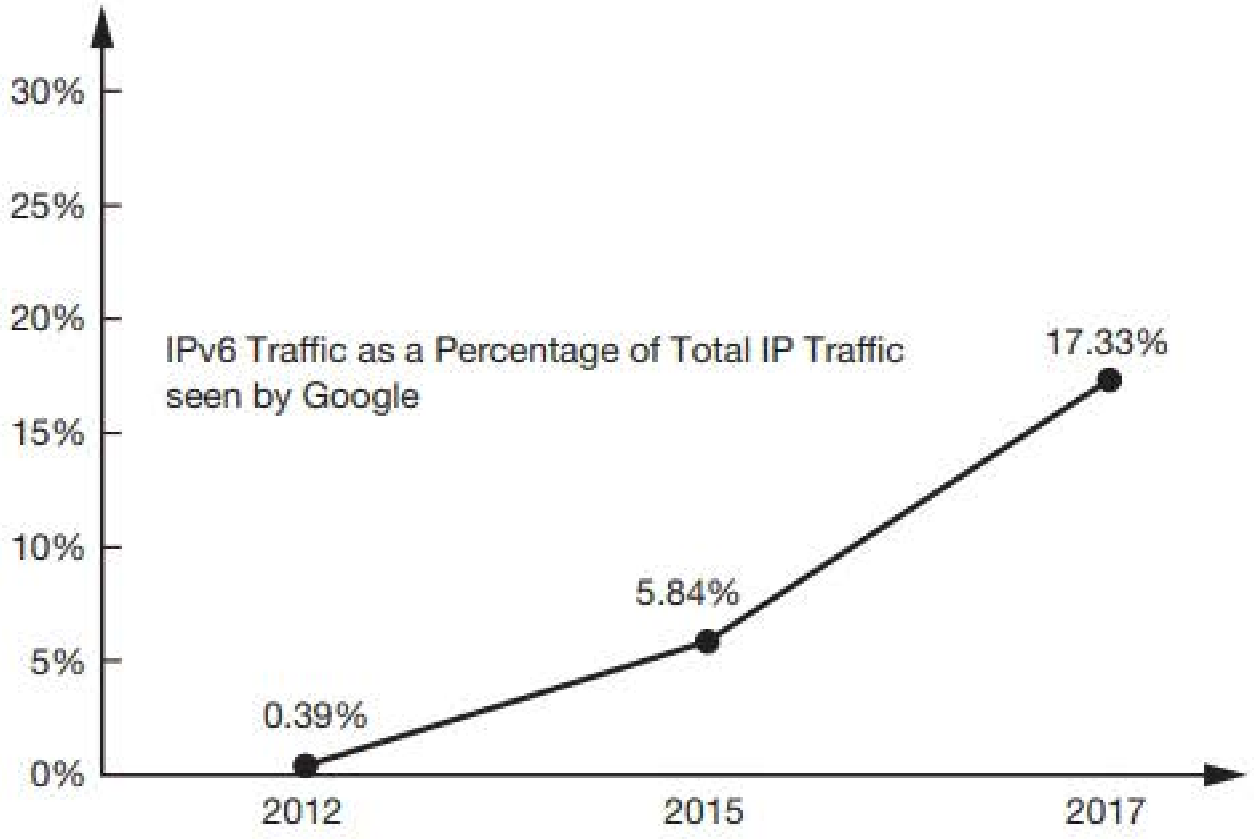

FIGURE 8-16 Explosive Growth in IPv6 Traffic

Source: Google.com.

complexity but at the gain of some security. IPv6 would have the mandatory inclusion of IPsec security functionality, but IPsec was quickly modified to work with IPv4 as well. Seeing no hard business case for upgrading, few companies did.

Now that IPv4 addresses are no longer available, however, nearly all companies are rushing to IPv6, and most have already begun to do so. Figure 8-16 shows that after years of hovering near zero penetration, IPv6 is growing explosively. In 2017, IPv6 accounted for 17 percent of all IP traffic received by Google. This is no longer a trend that can be avoided. However, companies have found that IPv6 implementation is a long and complex process. They need employees who understand this new protocol and other "v6" protocols such as ICMPv6 and DHCPv6. In addition, the tools to manage IPv6 are still less robust than those used to manage IPv4.

Test Your Understanding

19. What has been holding back the adoption of IPv6?

Writing IPv6 Addresses in Canonical Text Notation (RFC 5952)

We write IPv4 addresses for human consumption in dotted decimal notation—four segments of decimal numbers between 0 and 255. The segments are separated by dots. This gives addresses like 128.171.17.13. Humans can remember these addresses, and they are certainly easier to read and write than thirty-two 1s and 0s.

For the 128-bit addresses of IPv6, we would also like simpler ways to write them, but anything we do will still overload human memory. Consequently, when we write IPv6 addresses for human consumption, we do so to make the *reading and writing* easier. We also want to make the simplified IPv6 address searchable in text documents because they are often presented in such documents.

	128-bit IPv6 Address. Hard to write.	001000000000000100000000001001111111 1110001001100000000000000000000000000 0000000000000000000000000011001101001 1111110000111111001000
1.	Convert to hexadecimal notation; write letters in lowercase, divide four-symbol fields by colons.	001000000000000100000000001001111111 1110001001100000000000000000000000000 0000000000000000000000000011001101001 1111110000111111001000 **2001:0027:fe56:0000:0000:0000:cd3f:0fc0**
2.	Remove leading 0s from each field. However, there must be at least a single value left, so change 0000 to 0. Do not remove trailing zeroes (see last field in the right column).	2001:0027:fe56:0000:0000:0000:cd3f:0fc0 **2001:27:fe56:0:0:0:cd3f:fc0**
3.	Shorten ONE group of more than two groups of single-zero fields to two colons.	2001:27:fe56:0:0:0:cd3f:fca **2001:27:fe56::cd3f:fca**
	If there are multiple groups of more than two colons, shorten the longest.	2001:0:0:fe56:0:0:0:cd3f **2001:0:0:fe56::cd3f**
	If there is a tie for longest, choose the *first*.	2001:0:0:fe56:0:0:cd3f:fca **2001::fe56:0:0:cd3f:fca**
4.	**The final address in simplified IPv6 notation. Shortened but not short.**	**2001:27:fe56::cd3f:fca**

FIGURE 8-17 Writing IPv6 Addresses in IPv6 Canonical Text Representation Following RFC 5952

To write IPv6 addresses in the **IPv6 Canonical Text Representation**, we must follow a precise set of rules laid out in RFC 5952. Figure 8-17 shows these rules.

A 128-bit IPv6 address is shown in the following example. This is obviously difficult to write and to read.

0010000000000001000000000010011111111100101011000000000000000000000000000
000000000000000000000000011001101001111110000111111000000

Step 1: Convert to Hexadecimal Notation To simplify the IPv6 address, do not use dotted decimal notation as IPv4 does. Rather, IPv6 uses hexadecimal notation, which we saw in Chapter 5, in the context of Ethernet EUI-48 addresses.

Each "nibble" of 4 bits is converted into a hex symbol from 0 through F. A 128-bit IPv6 address, then, would be translated into 32 hex symbols (128 divided by 4).

In another annoying inconsistency in terminology, groups of IPv4 bits are collected into *segments* in dotted decimal notation, but groups of bits in IPv6 are called **fields.** This is an unfortunate use of terminology, because fields within an IPv6 are different from fields in the IPv6 header in general.

In Ethernet, we write hex symbols in pairs, separating each pair with a dash. This gives addresses like A1-B2-C3-D4-E5-F6. In IPv6, in contrast, we group hex symbols in tetrad (group of four) fields. An example of a field is *fe56*.

Note that we write the hex symbols in *lowercase* when writing hex symbols in IPv6 addresses. Each symbol is still 4 bits, so fe56 represents 16 bits. A full IPv6 address will have eight of these fields separated by *colons* (128 bits divided by 16). The following is an IPv6 address written in hexadecimal notation.

```
2001:0027:fe56:0000:0000:0000:cd3f:0fc0
```

Step 2: Remove *Leading* Zeroes from Segments This is still long. Fortunately, there are rules to help us shorten the writing of IPv6 addresses a little. The first is that in each field *any leading 0s are dropped*. This is easy to understand. If the reader sees: *27:*, this must be *:0027:*. Note that only *leading* 0s are dropped. If *trailing* 0s or 0s anywhere else were dropped too, the reader could not know if *:27:* was *:0027:*, *:2700:*, or *0270:*. Dropping leading 0s is also natural because we do that when writing decimal numbers. Here is what the IPv6 address looks like after leading 0s are dropped. (Note that the last segment is fc0, not fc.) It is much shorter.

```
2001:27:fe56:0:0:0:cd3f:fc0
```

Note that there is an exception to the rule about dropping leading zeroes. If a field consists of all zeroes (0000), shorten this to 0 instead of writing nothing. There are three such fields in this address.

Step 3: Reducing Multiple Single-Zero Fields If there are two or more single-zero fields in sequence, such as :0:0:0: in this example, you shorten this to a single pair of colons (::). So if an IPv6 address has the sequence *:0000:0000:0000:*, this can be replaced by ::. This further simplifies our IPv6 address to the following:

```
2001:27:fe56::cd3f:fca
```

Note that a *single* field with all zeros is not a *group* of all-zero fields. So if you have 2000:0:fa, you do not shorten this to 2000::fa.

If you have more than one group of single-zero fields, the following rules apply.

- First, if there is more than one group of single-zero fields, only *one* group of single-zero fields may be shortened to ::.
- Second, if there are multiple sequences of all-zero fields, the *longest* group of all-zero fields should be shortened. This just makes sense. One might as well shorten things as much as possible.
- Third, if two groups of single-zero fields tie for the longest number of all-zero groups, the *first* of these groups must be shortened.

These rules seem a little daunting, but these rules mean that everyone writes shortened IPv6 addresses the same way. Again, this is critical so that programs can test whether two addresses in written documentation are the same by comparing the text strings that are the simplified IPv6 notation.

Test Your Understanding

20. a) Why are IPv6 addresses simplified? b) Why must simplification rules be followed precisely? c) Are simplified IPv6 addresses written in uppercase or lowercase letters? d) Are simplified IPv6 addresses written with decimal or hexadecimal symbols? e) How many symbols are there in a field? f) How many bits are there in a field? g) How are fields separated? h) How many fields are there in an IPv6 address?

21. a) Write the following IPv6 address in canonical form using RFC 5952: 2001 :0ed2:056b:00d3:000c:abcd:0bcd:0fe0. b) Write the following IPv6 address in canonical form using RFC 5952: 2001:0002:0000:0000:0000:abcd:0bcd:0fe0. c) Simplify the following IPv6 address using RFC 5952: 2001:0000:0000:00fe: 0000:0000:0000:cdef. d) Simplify the following IPv6 address using RFC 5952: 2001:0000:0000:00fe:0000:0000:ba5a:cdef. e) What is the advantage of simplifying IPv6 addresses according to strict rules? f) Which RFC is used to write IPv6 addresses in canonical form?

The IPv6 Main Header

Figure 8-18 shows the IPv4 header. Actually, we will call this the **IPv6 main header** because, as we will see, an IPv6 packet can have multiple extension headers before the Data Field. The obvious difference between the IPv4 and IPv6 headers are that IPv4 addresses are 32 bits while IPv6 addresses are 128 bits.

The second difference is that the IPv6 main header, although longer, is simpler than the IPv4 header, with fewer fields for routers and hosts to consider. This relative simplicity means that routers process longer IPv6 headers faster than they process IPv4 headers. This makes them cheaper for the amount of traffic that they process.

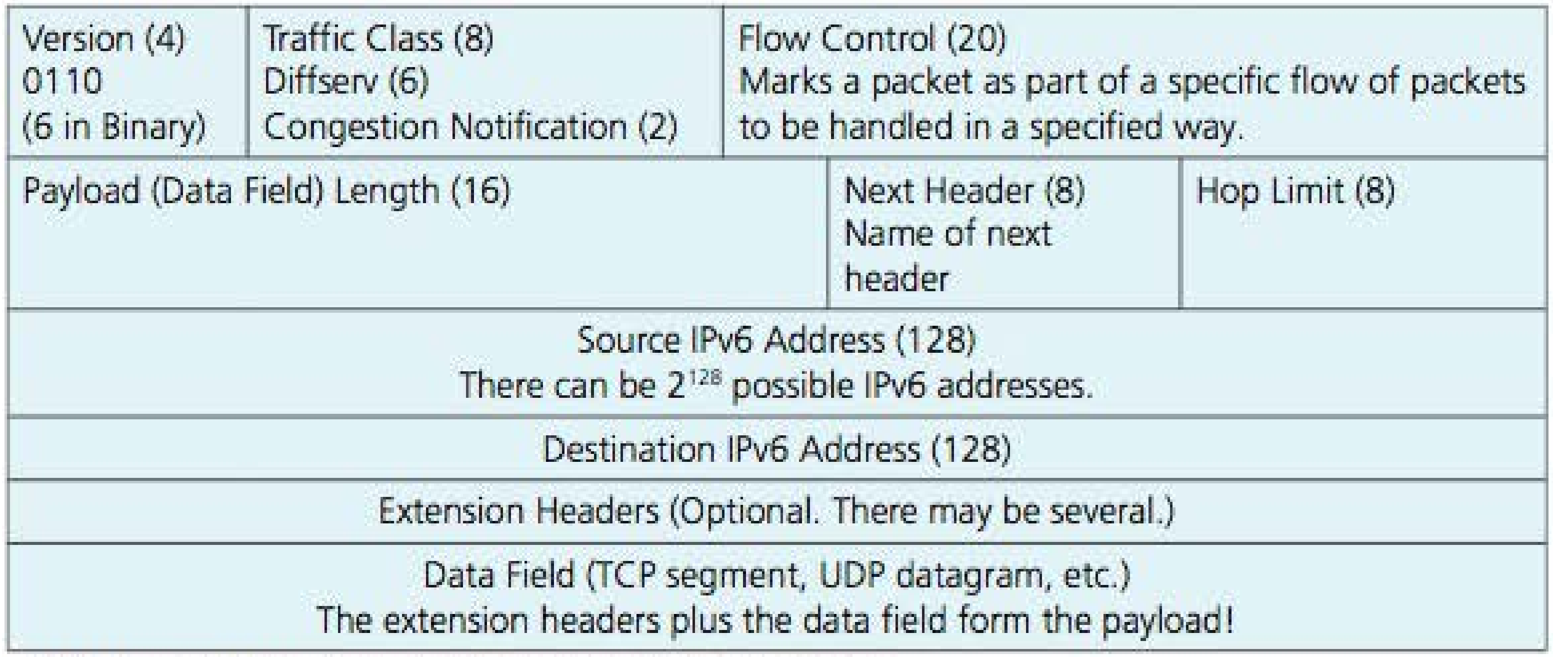

Version (4)
0110
(6 in Binary)

Traffic Class (8)
Diffserv (6)
Congestion Notification (2)

Flow Control (20)
Marks a packet as part of a specific flow of packets to be handled in a specified way.

Payload (Data Field) Length (16)

Next Header (8)
Name of next header

Hop Limit (8)

Source IPv6 Address (128)
There can be 2^{128} possible IPv6 addresses.

Destination IPv6 Address (128)

Extension Headers (Optional. There may be several.)

Data Field (TCP segment, UDP datagram, etc.)
The extension headers plus the data field form the payload!

Traffic Class Field (8) has two parts: Diffserv and Congestion Notification

Differentiated Services (Diffserv) Field describes specific special (differentiated) services requested, such as priority.

Congestion Notification notifies the receiver that congestion has been experienced along the route.

FIGURE 8-18 IP Version 6 (IPv6) Packet Syntax with Main Header, Data Field, and Possibly Extension Headers (Next Headers)

Version Number Field Both headers begin with a 4-bit **Version Number Field**. For IPv4, the field value is 0100 (four). For IPv6, it is 0110 (six).

Traffic Class and Flow Label Fields The first row of the IPv6 header also contains an 8-bit Traffic Class field and a 20-bit flow label field.[7] The two fields specify how routing will be handled in terms of priority and other quality of service matters.

- The **Traffic Class Field** has two subfields. The 6-bit **Differentiated Services (Diffserv)** subfield specifies whether *this particular packet* should be given routine best-effort service, high-priority low-latency service, or some other type of service. The last 2 bits are for congestion notification.
- The **Flow Label Field** value indicates that the packet is a member of a particular flow. The router has rules that apply to *every packet* in the flow.

Payload Length In IPv6, the **Payload Length Field** gives the length of the packet payload, which is everything beyond the 40-octet main packet header. The Payload Length Field is 16 bits long, so a payload can be up to 65,536 (2^{16}) octets long.

The Payload Length Field gives the length of the packet payload, which is everything beyond the 40-octet main packet header. It includes both extension headers and the data field.

Hop Limit Field IPv6 has a **Hop Limit Field** that is like the IPv4 time to live field. Each router along the way decrements this field's value by one, and if a router decrements it to zero, the router discards the packet.[8]

No Checksum Field? IPv4 has a Header Checksum Field to check for packet header errors. When IPv4 was created, there was a concern that if packet headers contained errors, they could cause serious problems for the Internet. Experience proved this concern to be groundless, so IPv6 has no checksum field. The computations needed to check for errors in IPv4 were taxing, even for a 20-octet header. Dropping the checksum field slashes packet handling time on routers.

Test Your Understanding

22. a) How do the Version Number Fields in IPv4 and IPv6 differ? b) What is the general purpose of the Diffserv subfield? c) Of the Flow Label Field? d) In IPv6, how can the receiver tell the length of packet? e) Does the Payload Length Field include the lengths of any extension headers in the packet? f) How is the Hop Limit Field used? g) Does IPv6 have a header checksum field? h) What is the consequence of this?

[7] In the original definition of IPv6, these fields were 4 bits and 24 bits, respectively.

[8] Internet old timers know that when IPv4 was created, the time to live value was supposed to be measured in seconds. However, this proved to be unworkable. The value was then interpreted as the maximum number of hops permitted by the packet. The hop limit field name in IPv6 recognizes this.

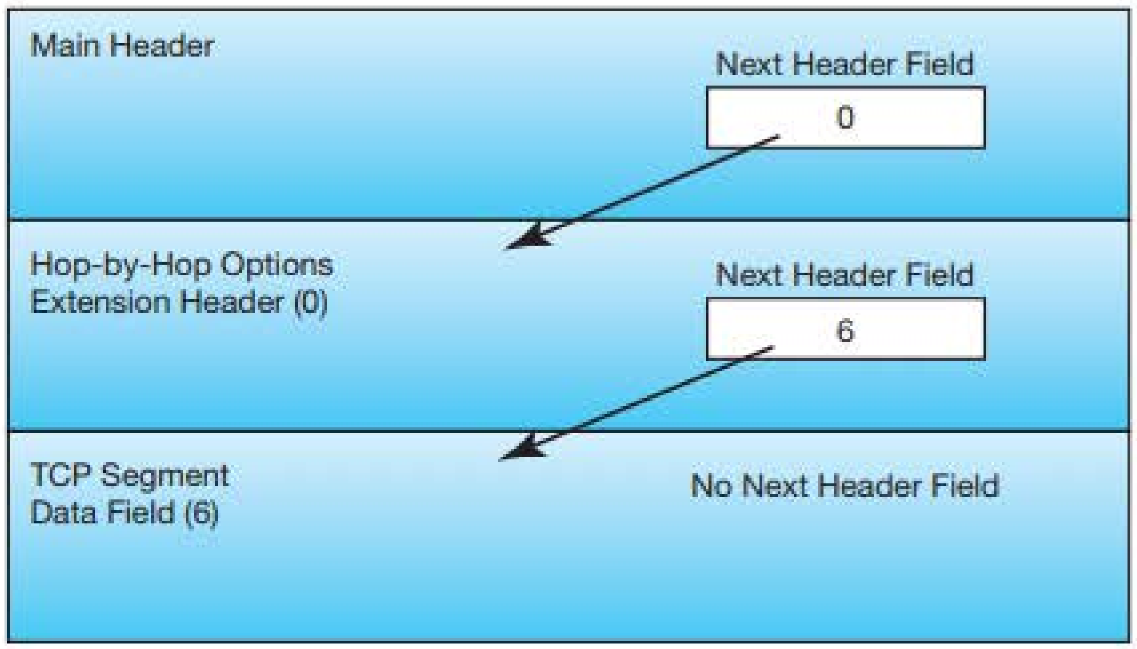

FIGURE 8-19 Main Header and Extension Headers in IPv6

Extension Headers

The IPv4 packet has option fields that allow the sender to add options. Few IPv4 packets have options, but each router must check each packet for options, and this can cost significant time, especially because many options are only relevant to the destination host.

Main Header and Extension Header IPv6 took a different approach to options. As Figure 8-19 shows, the main header can be followed by multiple **extension headers**. Each extension header has a well-defined purpose, such as providing information for security or mobile operation. Each extension header serves the role that an option does in IPv4.

Next Header Field The headers are daisy chained together based on the **Next Header Field**. The main header's Next Header Field specifies the first extension header. In Figure 8-19, the value is 0, meaning that the first extension header has hop-by-hop options that every router along the way must contend with. This is often the only extension header that routers need to deal with. That extension header's Next Header Field has the value 6, indicating that this header is followed by the TCP segment.

It is easy to confuse the terms payload and data field. The data field is the content message being delivered. The payload is everything that follows the main header. So the payload consists of both extension headers and the data field.

Figure 8-20 shows a few of the extension headers that have been defined for the Next Header Field, as well as their code values (in parentheses). The full list is much longer.

Next Header Code (Value)	
Supervisory Header	Upper Layer Messages Header
Hop-by-Hop Options (0)	TCP (6)
Destination Options (60)	UDP (17)
Mobility Header (135)	ICMP (1)
Encapsulating Security Payload Header (50)	ICMPv6 (58)

FIGURE 8-20 IPv6 Next Header Values

Test Your Understanding

23. a) Why is handling options the way that IPv4 does undesirable? b) Why is the approach of using optional extension headers desirable? c) What is often the only extension header that routers must consider? d) How does the last extension header before a UDP datagram indicate that the UDP datagram comes next? e) If you see 0 in the Next Header Field of a header, what will follow this header? f) Why are the terms *payload* and *data field* **not** synonymous?

THE TRANSMISSION CONTROL PROTOCOL (TCP)

Fields in TCP/IP Segments

In Chapter 2, we looked briefly at the syntax of TCP segments. In this section, we look at the syntax of TCP segments in more depth. When IP was designed, it was made to be a very simple "best effort" protocol (although its routing tables are complex). The IETF left more complex internetwork transmission control tasks to TCP. Consequently, network professionals need to understand TCP very well. Figure 8-21 shows the syntax of TCP segments.

Sequence Numbers TCP can handle application messages of almost any length. In Chapter 2, we saw that TCP does this by fragmenting long messages into many pieces and sending each segment in its own TCP segment. For the receiver to put the pieces of the application messages back in order, each TCP segment has a **Sequence Number Field** that gives its position in the stream of segments. The receiving TCP process puts the segments in order of increasing sequence number and reassembles the application message. The TCP process then passes the application message up to the correct application process indicated in the port number.[9]

Acknowledgment Numbers In Chapter 2, we saw that TCP uses **acknowledgments (ACKs)** to achieve reliability. If a transport process receives a TCP segment

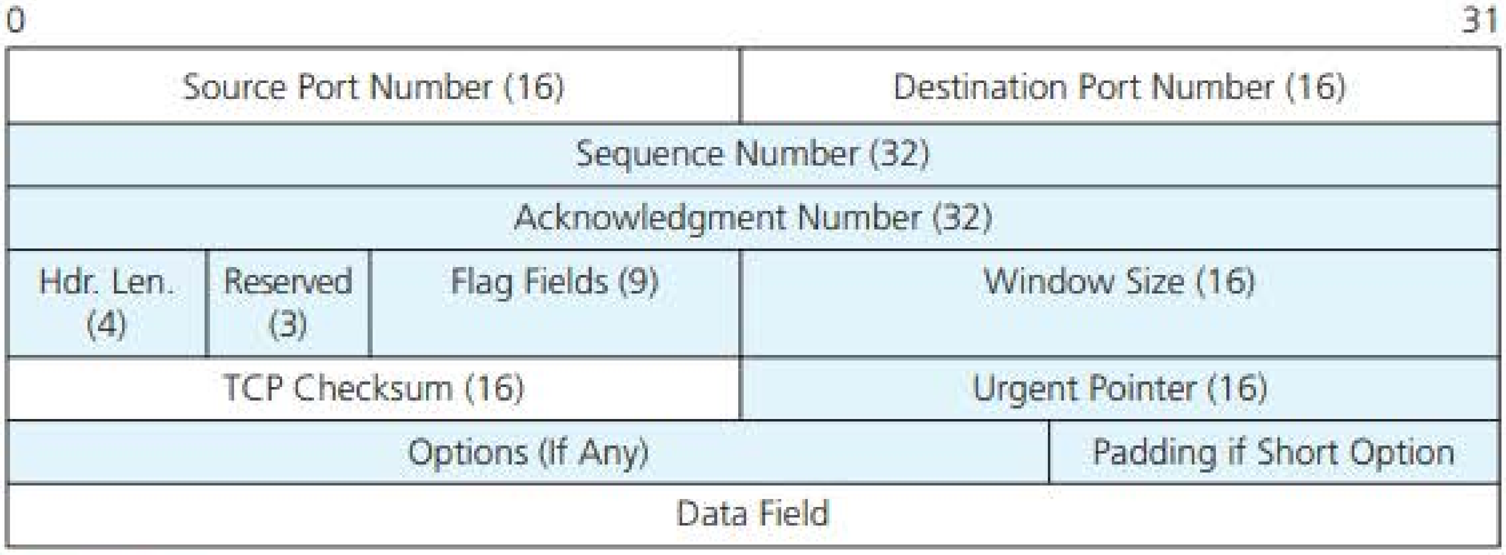

0 … 31

Source Port Number (16)		Destination Port Number (16)	
Sequence Number (32)			
Acknowledgment Number (32)			
Hdr. Len. (4)	Reserved (3)	Flag Fields (9)	Window Size (16)
TCP Checksum (16)		Urgent Pointer (16)	
Options (If Any)			Padding if Short Option
Data Field			

FIGURE 8-21 Fields in a TCP Segment

[9] Online Module A has a detailed discussion of TCP sequence and acknowledgment numbers.

correctly, it sends back a TCP segment acknowledging the reception. If the sending transport process does not receive an acknowledgment, it transmits the TCP segment again.

The **Acknowledgment Number Field** indicates which segment is being acknowledged. One might expect that if a segment has sequence number X, then the acknowledgment number in the segment that acknowledges it would also be X. Online Module A shows that the situation is more complex, but the acknowledgment number is at least related to the sequence number of the segment being acknowledged.

Flag Fields As discussed in Chapter 2, TCP has nine single-bit fields. Single-bit fields are called **flag fields**. If they have the value 1, they are said to be set. A 0 means that a flag field is not set. We saw several uses of these flag bits in Chapter 2.

- If the ACK bit is set, then the segment acknowledges another segment. If the ACK bit is set, the acknowledgment field must be filled in to indicate which message is being acknowledged.
- If the SYN (synchronization) bit is set, then the segment requests a connection opening.
- If the FIN (finish) bit is set, then the segment requests a normal connection closing.

Options Fields It is common for TCP segments to have option fields. Unfortunately, this feature was not well thought out in the original design, so there is no simple way to talk about TCP options.

Test Your Understanding

24. a) How long are sequence and acknowledgment numbers? b) How many flag fields do TCP headers have? c) If the ACK bit is set, what other field must have a value?

Openings and Abrupt TCP Closes

In Chapter 2, we saw that TCP is a connection-oriented protocol. Connection-oriented protocols have formal openings and closings. Figure 8-22 recaps normal closes and introduces a second type of close, the reset.

In Chapter 2, we looked at *normal* closings. Just as you do not simply hang up on a telephone call when you want to finish talking, if you are polite, a normal TCP close consists of two FIN segments, one in each direction, plus their acknowledgments.

However, Figure 8-22 shows that TCP also permits another type of close. This is an abrupt close. Whenever either side wishes to end a conversation, it can simply send a **TCP reset segment**. This is a segment with the **RST** (reset) flag bit set. A **reset** may occur if a problem is encountered during a connection, for security reasons, or for several other reasons.

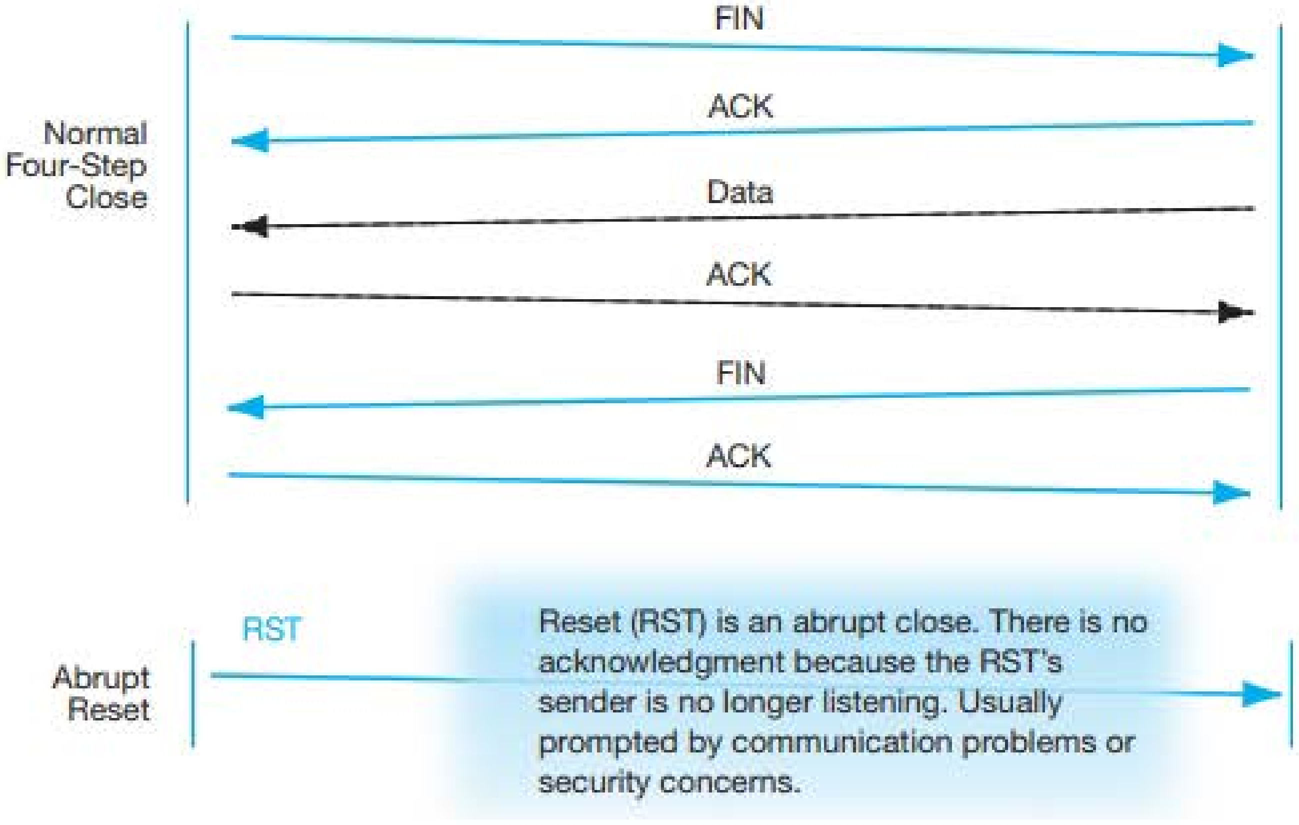

Reset (RST) is an abrupt close. There is no acknowledgment because the RST's sender is no longer listening. Usually prompted by communication problems or security concerns.

FIGURE 8-22 TCP Session Openings and Closings

Note in Figure 8-22 that an RST segment is not acknowledged. The side that sent the RST segment is not listening any longer, so acknowledging a reset would be as pointless as saying goodbye after someone has hung up on you. The RST segment is one of two segment types that are not acknowledged. As noted in Chapter 2, a segment that is nothing more than an acknowledgment (a pure acknowledgment) is not acknowledged because doing so would create an endless loop of acknowledgments.

Test Your Understanding

25. a) What is a FIN segment? b) Distinguish between four-way closes and abrupt resets. c) Why is a reset segment not acknowledged? d) What other type of segment is not acknowledged?

THE LIMITED MAXIMUM LENGTH OF USER DATAGRAM PROTOCOL (UDP) DATAGRAMS

We saw UDP in Chapter 2. This is a very simple protocol, so the discussion in that chapter is sufficient except for one point. This is the fact that UDP, unlike TCP, cannot do segmentation. The entire application message must fit into a single UDP datagram. Figure 8-23 shows that the Length Field in the UDP header is 16 bits long, so the maximum length of the UDP data field (and therefore the maximum length of an application message) is 65,536 octets. On the plus side, there is no need for sequence numbers, opening, closings, acknowledgments, or other things that require a longer header.

UDP cannot do segmentation, so an application message must fit into a single UDP datagram.

0	31
Source Port Number Field (16)	Destination Port Number Field (16)
UDP Length (in Octets) (16)	UDP Checksum (Error Discarding but No Correction) (16)
Data Field (Variable Length: 2^{16} bits give a maximum of 65,536 Octets)	

The UDP length field gives the number of octets in the data field.

UDP's length field is 16 bits.

It can represent 2^{16} possible values—65,536.

So the maximum length of the data field is 65,536 octets.

There are no sequence numbers, so longer application messages cannot be segmented and sent over several UDP datagrams.

So UDP cannot send application messages longer than 65,536 octets.

FIGURE 8-23 UDP Datagram Fields

Test Your Understanding

26. a) Why can TCP handle long application messages? b) Why can UDP not handle long application messages? c) What is the maximum application message size when UDP is used at the transport layer?

END-OF-CHAPTER QUESTIONS

Thought Questions

8-1. a) How does the postal service use hierarchical sorting? b) How does this simplify delivery decisions?

8-2. Give a non-network example of hierarchical addressing, and discuss how it reduces the amount of work needed in physical delivery. Do not use the postal service, or the telephone network.

8-3. A client PC has two simultaneous connections to the same webserver application program on a webserver. (Yes, this is possible, and in fact, it is rather common.) What will be different between the TCP segments that the client sends on the two connections? (*Hint:* Consider all the fields in a TCP segment.)

8-4. A router that has the routing table in Figure 8-11 receives an incoming IPv4 packet. The source IPv4 address in the arriving packet is 10.55.72.234. The destination IPv4 address is 10.4.6.7. The TTL value is 1. The Protocol Field value is 6. What will the router do with this packet? (*Hint:* Carefully consider all the fields in the IP and TCP headers. Think like a router.)

Perspective Questions

8-5. What was the most surprising thing you learned in this chapter?

8-6. What was the most difficult material for you in this chapter?

Chapter **8a**

Hands-On: Wireshark Packet Capture

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Use the Wireshark packet capture program at a novice level.
- Capture packets in real time.
- Analyze the packets at a novice level.

INTRODUCTION

A good way to practice what you have learned in this chapter is to look at individual packets. Packet capture programs record packets going into and out of your computer. If you capture a brief webserver interaction, you can look at header fields, TCP three-step connection starts, and other information. There are several good packet capture programs. We look at Wireshark, which is simple to use, popular, and free to download. (At least at the time of this writing.)

GETTING WIRESHARK

To get Wireshark, go to wireshark.org. Do *not* go to wireshark.com. Follow the instructions and download the program on your computer.

USING WIRESHARK

Getting Started

After installation, open the Wireshark program. You will see the opening screen. It will look like the screen in Figure 8a-1. There will be controls at the top with a blank area below them. You will soon fill this area with your packet capture.

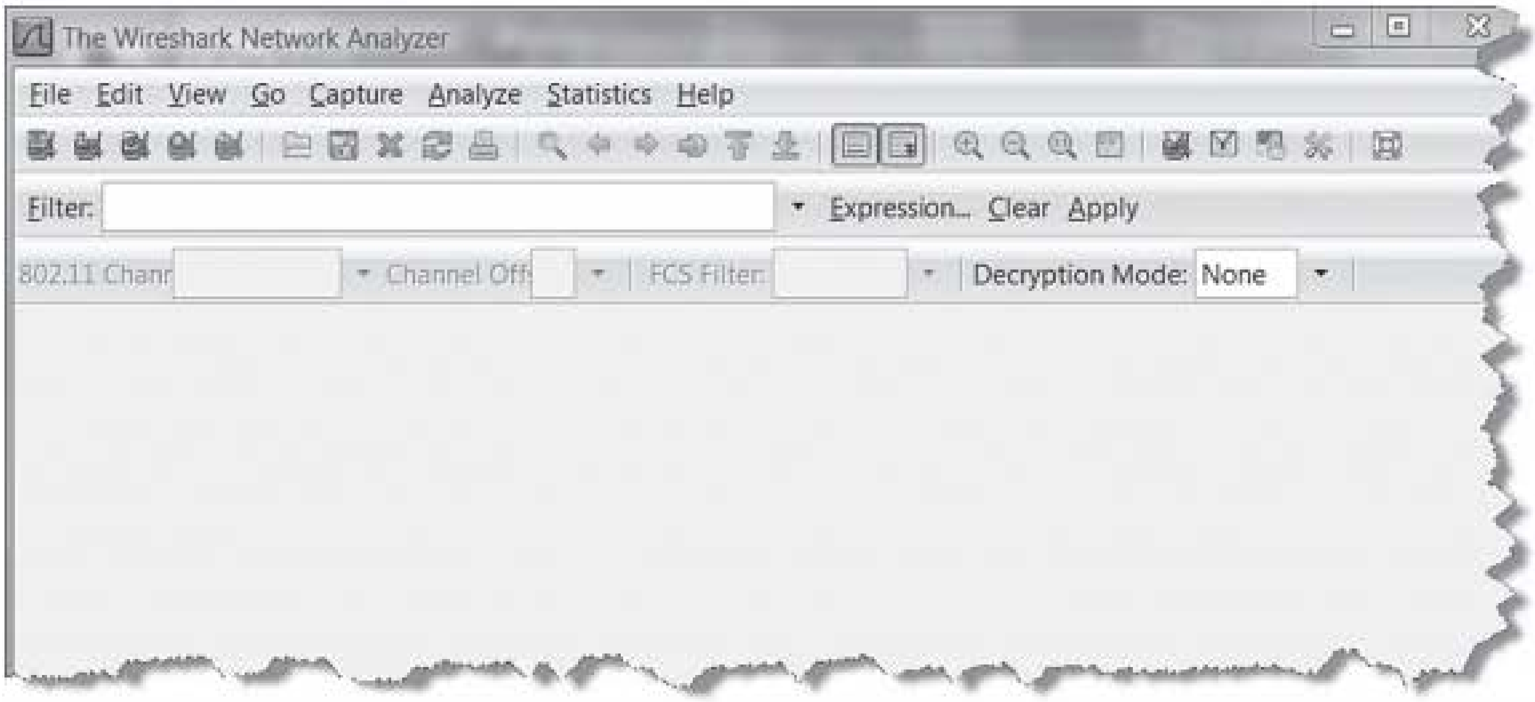
The Wireshark Network Analyzer
File Edit View Go Capture Analyze Statistics Help
Filter:
Expression... Clear Apply
802.11 Chann
Channel Off
FCS Filter:
Decryption Mode: None

FIGURE 8a-1 Initial Wireshark Screen

Starting a Packet Capture

To start a packet capture, click on the Go menu item. Then, when the Wireshark: Capture Interfaces dialog box appears, as Figure 8a-2 illustrates, select a network interface and click on Start.

Getting Data

Your browser should already be open. Switch to your browser and enter a URL. (In this example, the author went to Wikipedia.org.) This creates a flurry of packets between

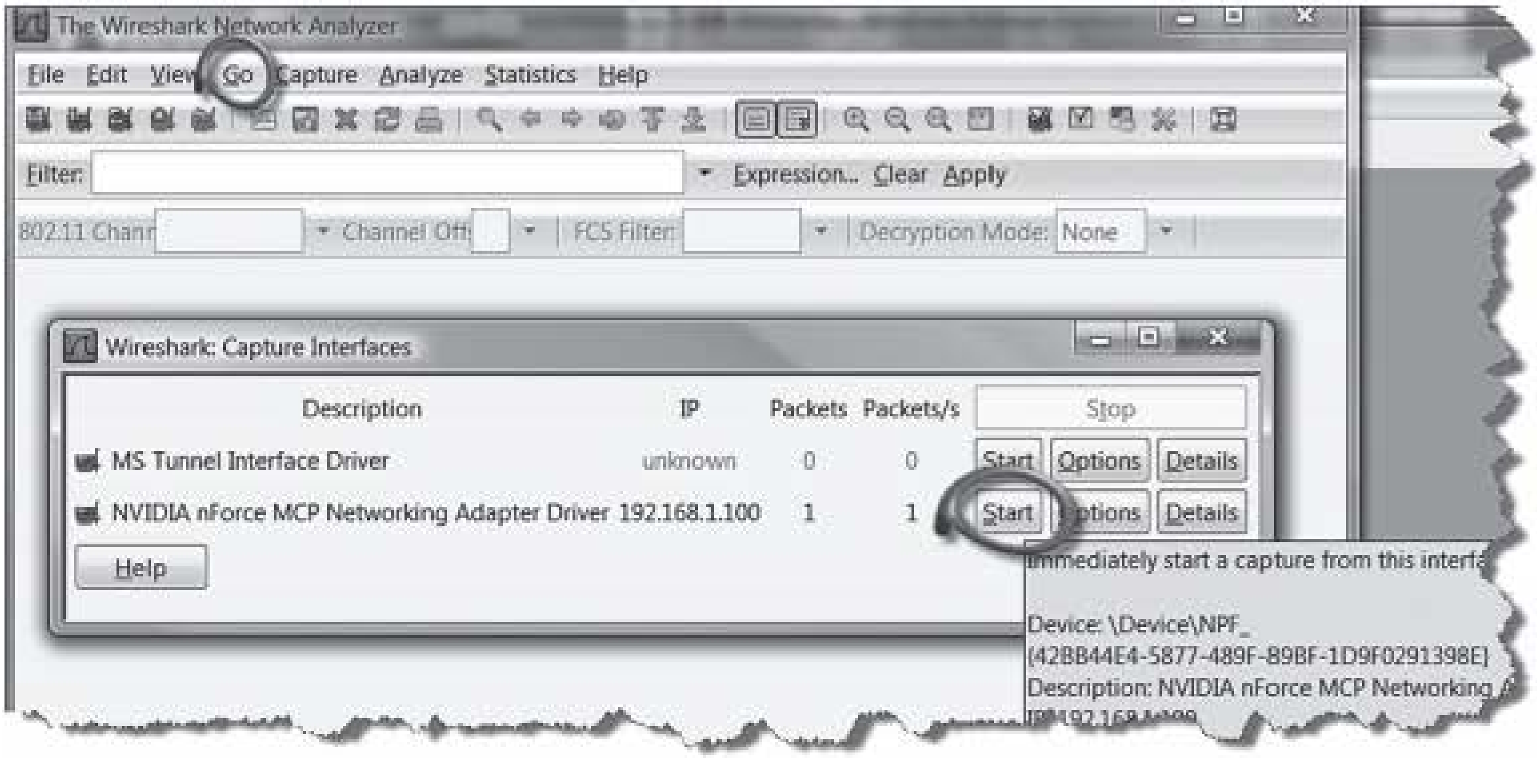
The Wireshark Network Analyzer
File Edit View Go Capture Analyze Statistics Help
Filter:
Expression... Clear Apply
802.11 Chann
Channel Off
FCS Filter:
Decryption Mode: None
Wireshark: Capture Interfaces
Stop
Start
Options
Details
Help

Description	IP	Packets	Packets/s
MS Tunnel Interface Driver	unknown	0	0
NVIDIA nForce MCP Networking Adapter Driver	192.168.1.100	1	1

Immediately start a capture from this interfa

Device: \Device\NPF_
{428B44E4-5877-489F-89BF-1D9F0291398E}
Description: NVIDIA nForce MCP Networking A
IP: 192.168.1.100

FIGURE 8a-2 Starting a Packet Capture in Wireshark

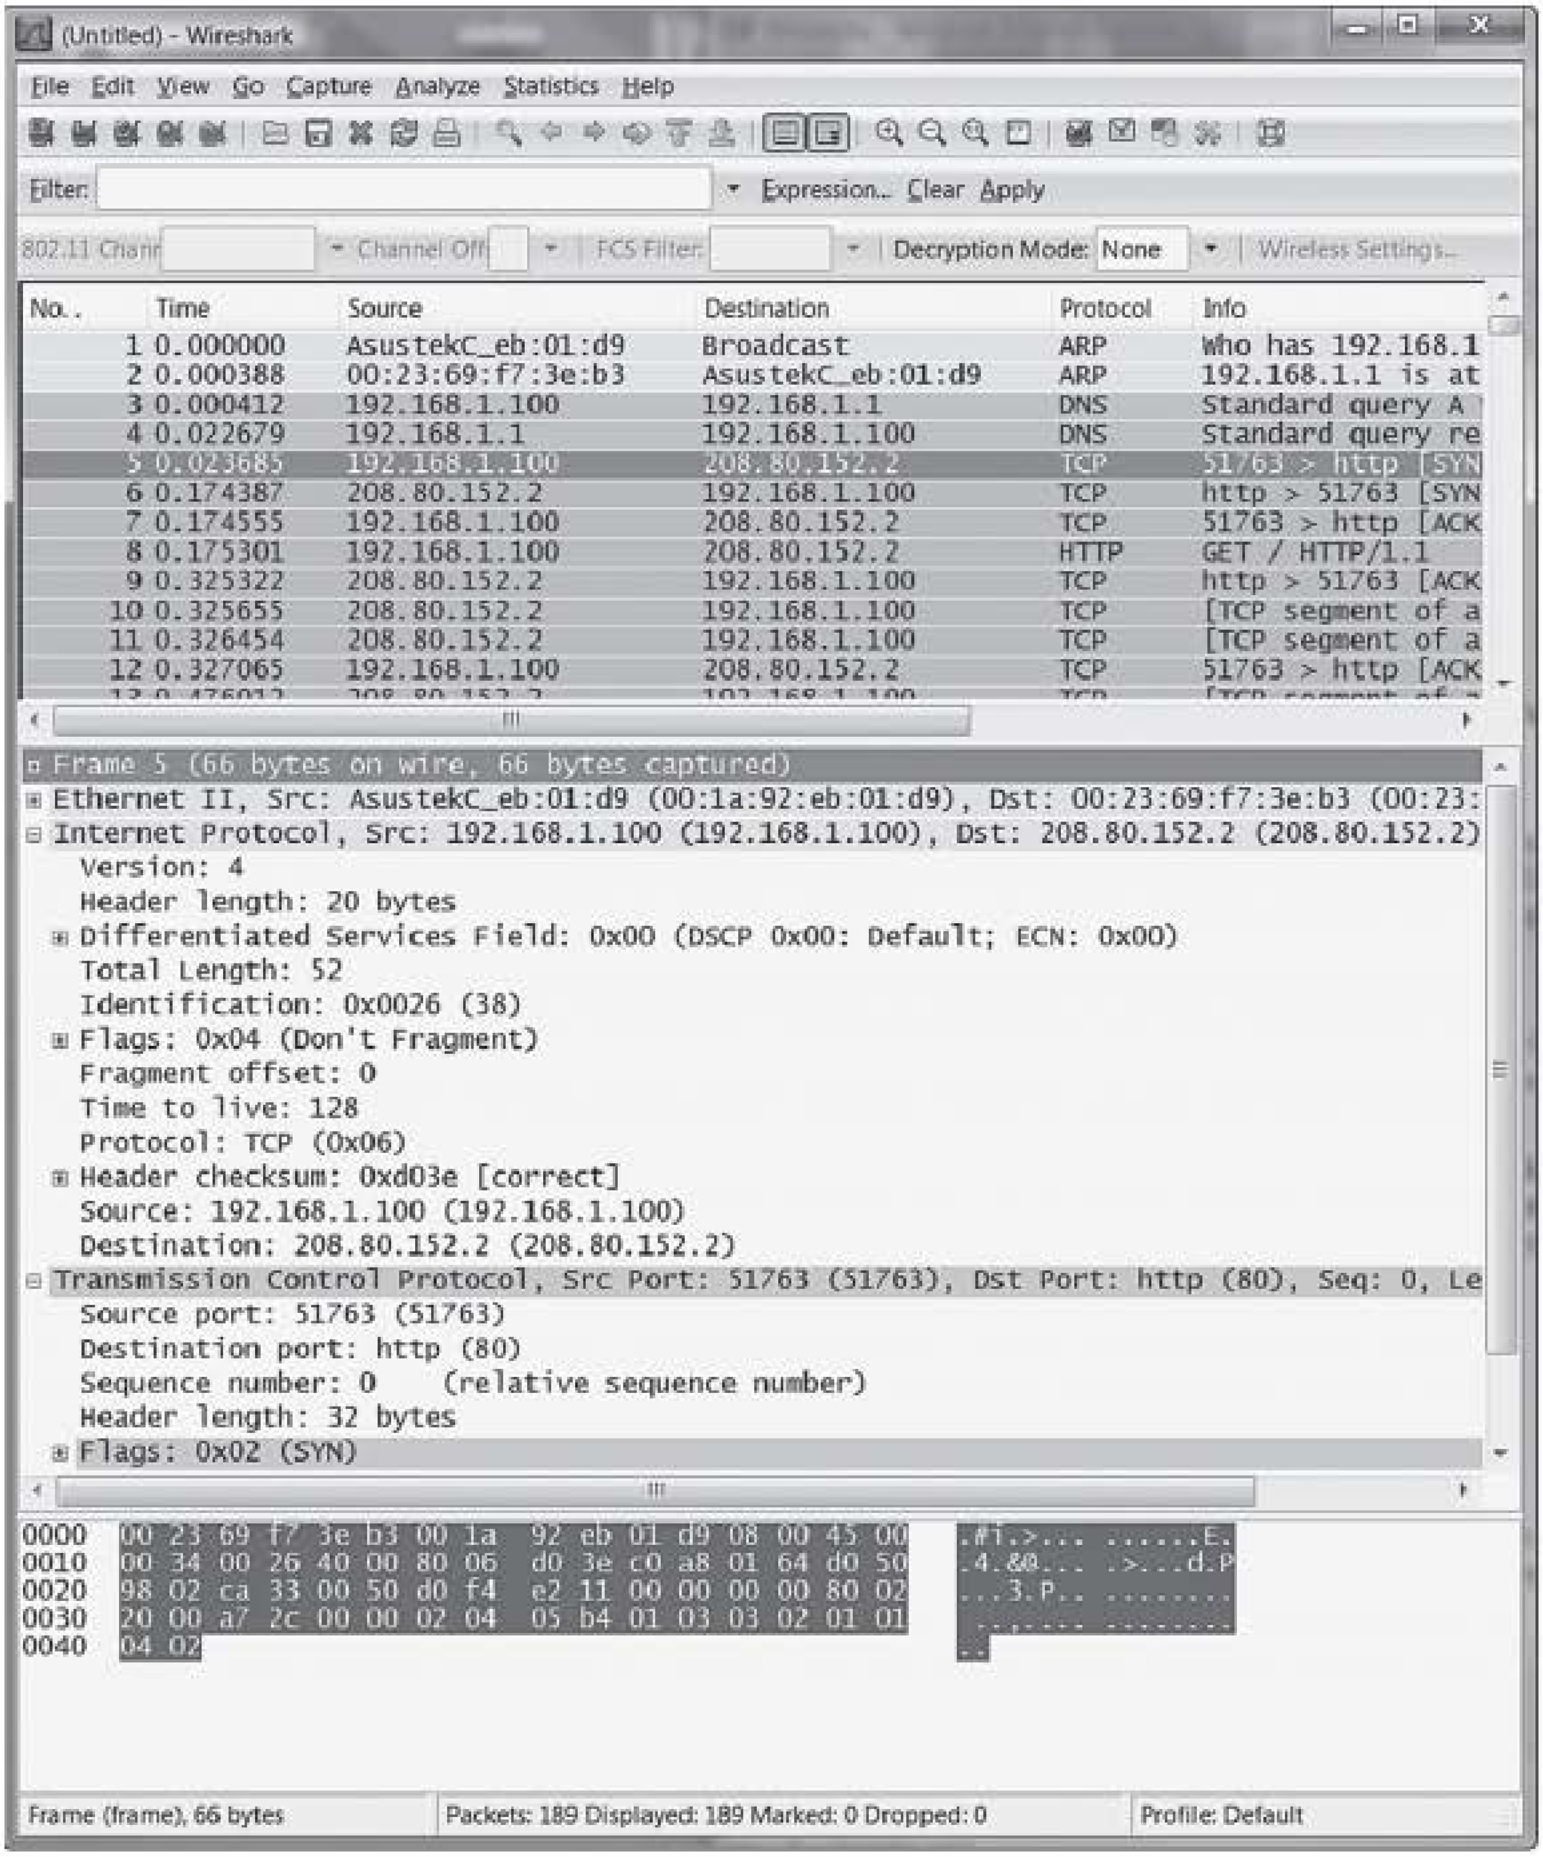

```
No. .  Time      Source              Destination         Protocol  Info
    1 0.000000  AsustekC_eb:01:d9   Broadcast           ARP       Who has 192.168.1
    2 0.000388  00:23:69:f7:3e:b3   AsustekC_eb:01:d9   ARP       192.168.1.1 is at
    3 0.000412  192.168.1.100       192.168.1.1         DNS       Standard query A
    4 0.022679  192.168.1.1         192.168.1.100       DNS       Standard query re
    5 0.023685  192.168.1.100       208.80.152.2        TCP       51763 > http [SYN
    6 0.174387  208.80.152.2        192.168.1.100       TCP       http > 51763 [SYN
    7 0.174555  192.168.1.100       208.80.152.2        TCP       51763 > http [ACK
    8 0.175301  192.168.1.100       208.80.152.2        HTTP      GET / HTTP/1.1
    9 0.325322  208.80.152.2        192.168.1.100       TCP       http > 51763 [ACK
   10 0.325655  208.80.152.2        192.168.1.100       TCP       [TCP segment of a
   11 0.326454  208.80.152.2        192.168.1.100       TCP       [TCP segment of a
   12 0.327065  192.168.1.100       208.80.152.2        TCP       51763 > http [ACK
```

```
⊞ Frame 5 (66 bytes on wire, 66 bytes captured)
⊞ Ethernet II, Src: AsustekC_eb:01:d9 (00:1a:92:eb:01:d9), Dst: 00:23:69:f7:3e:b3 (00:23:
⊟ Internet Protocol, Src: 192.168.1.100 (192.168.1.100), Dst: 208.80.152.2 (208.80.152.2)
    Version: 4
    Header length: 20 bytes
  ⊞ Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
    Total Length: 52
    Identification: 0x0026 (38)
  ⊞ Flags: 0x04 (Don't Fragment)
    Fragment offset: 0
    Time to live: 128
    Protocol: TCP (0x06)
  ⊞ Header checksum: 0xd03e [correct]
    Source: 192.168.1.100 (192.168.1.100)
    Destination: 208.80.152.2 (208.80.152.2)
⊟ Transmission Control Protocol, Src Port: 51763 (51763), Dst Port: http (80), Seq: 0, Le
    Source port: 51763 (51763)
    Destination port: http (80)
    Sequence number: 0    (relative sequence number)
    Header length: 32 bytes
  ⊞ Flags: 0x02 (SYN)
```

```
0000  00 23 69 f7 3e b3 00 1a  92 eb 01 d9 08 00 45 00   .#i.>... ......E.
0010  00 34 00 26 40 00 80 06  d0 3e c0 a8 01 64 d0 50   .4.&@... .>...d.P
0020  98 02 ca 33 00 50 d0 f4  e2 11 00 00 00 00 80 02   ...3.P.. ........
0030  20 00 a7 2c 00 00 02 04  05 b4 01 03 03 02 01 01    ..,.... ........
0040  04 02                                              ..
```

```
Frame (frame), 66 bytes    Packets: 189 Displayed: 189 Marked: 0 Dropped: 0    Profile: Default
```

FIGURE 8a-3 Collecting Data

you and the host specified in the URL. These appear on the window below the controls, as shown in Figure 8a-3.

Stopping Data Collection

To stop the data collection, click on the Capture menu item, as Figure 8a-4 shows. When the dropdown menu appears, select *Stop*. You now have a packet stream to analyze.

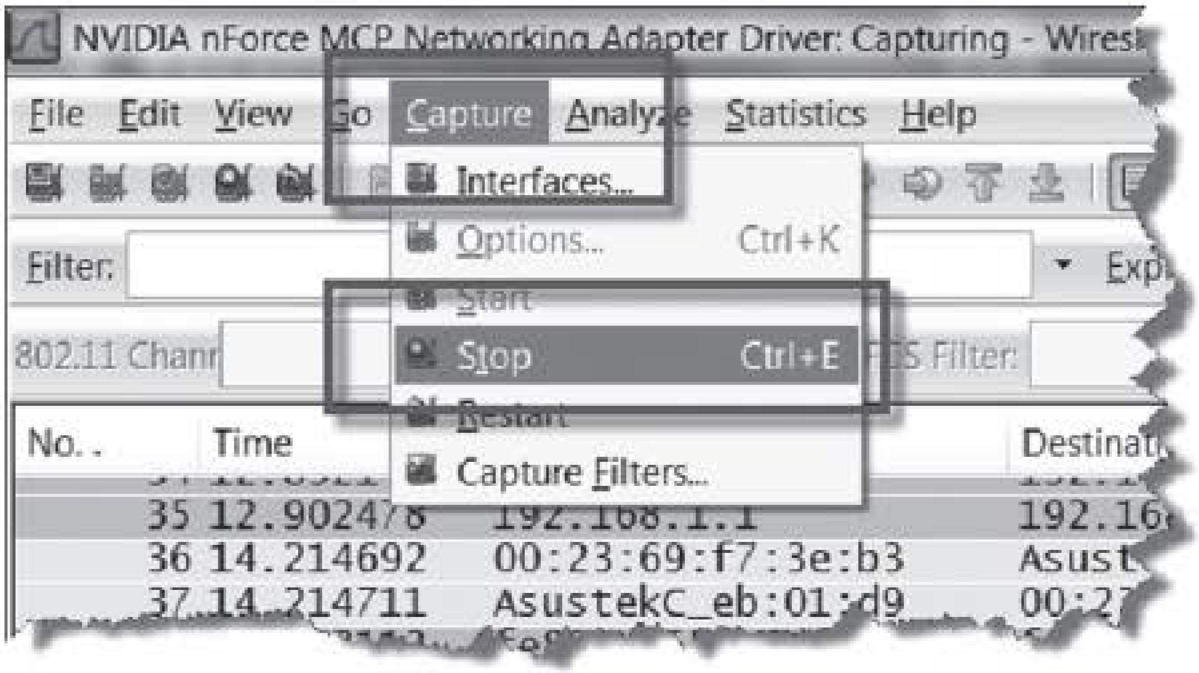
NVIDIA nForce MCP Networking Adapter Driver: Capturing - Wiresh
File Edit View Go Capture Analyze Statistics Help
Interfaces...
Options... Ctrl+K
Filter:
Start
802.11 Chann
Stop Ctrl+E
Restart
Capture Filters...
No. .
Time
Destinati

```
35 12.902478  192.168.1.1        192.16
36 14.214692  00:23:69:f7:3e:b3  Asust
37 14.214711  AsustekC_eb:01:d9  00:2
```

FIGURE 8a-4 Stopping the Data Collection

Looking at Individual Packets

Now you can begin looking at individual packets. To see how to do this, look again at Figure 8a-3.

Packet Summary Window In the upper window in the display area, you can see the packets one at a time. The capture begins with two ARP packets, which identify the data link layer address of the host with IP address 192.168.1.1.

Then comes two DNS packets. In the example, the author typed the host name Wikipedia.org in the URL. The author's computer (192.168.1.100) sent a DNS request message to its DNS server to get the IP address for Wikipedia.org. The DNS sent back the requested IP address.

Now, the author's computer opened a connection to 208.80.152.2, which is Wireshark.org's IP address.[1] It first sent a TCP SYN segment to 208.80.152.2. This is Frame 5. In Figure 8a-3, the frame has been selected.

Information about the contents of this particular frame is shown in a window below the window showing each frame on a single line. First, the window shows information on the Ethernet header and trailer. Next comes information about the IP packet, followed by information about the TCP SYN segment contained in the packet.

Window with Detailed Information on the Selected Packet The Ethernet information has been minimized. Only the source and destination MAC addresses are shown. However, information about the IP packet has been maximized. You can see the values of the individual fields in the selected packet. For example, note that the Time to Live Field in this packet has the value 128. In addition, the protocol field value indicates that the data field contains a TCP segment.

[1] If you try this, you may get a different IP address. Many firms have multiple physical webservers that they associate with a host name. A DNS response message returns the IP address of one of these physical servers.

The TCP segment information also is expanded, although only the first few fields are shown in the window. Note that the destination port is 80, indicating that the author was contacting the Wireshark.org webserver. Note also that the Flag Fields information says that the SYN bit is set, as one would expect.

To make life easier for you, Wireshark does as much translation as possible. For example, it interprets the information in the protocol field as indicating that there is a TCP segment in the packet's data field. It also indicates that Port 80 is HTTP.

The information on sequence number is highly simplified compared to the discussion in Chapter 2. This is the first TCP segment being sent. It is given the value 0 rather than its complex real value.

Hex Window The lowest window shows the contents of the packet in hexadecimal (Base 16) format. Hex is difficult for new analysts to interpret, but it is very compact compared to the information in the middle window. Experienced packet analysts quickly learn the positions of important fields and learn to read the hex symbols for that field.

Options

Figure 8a-5 shows that Wireshark capture options allow you to control what packets are captured. If you are connected to multiple external servers simultaneously, this can allow you to capture only packets for a particular connection.

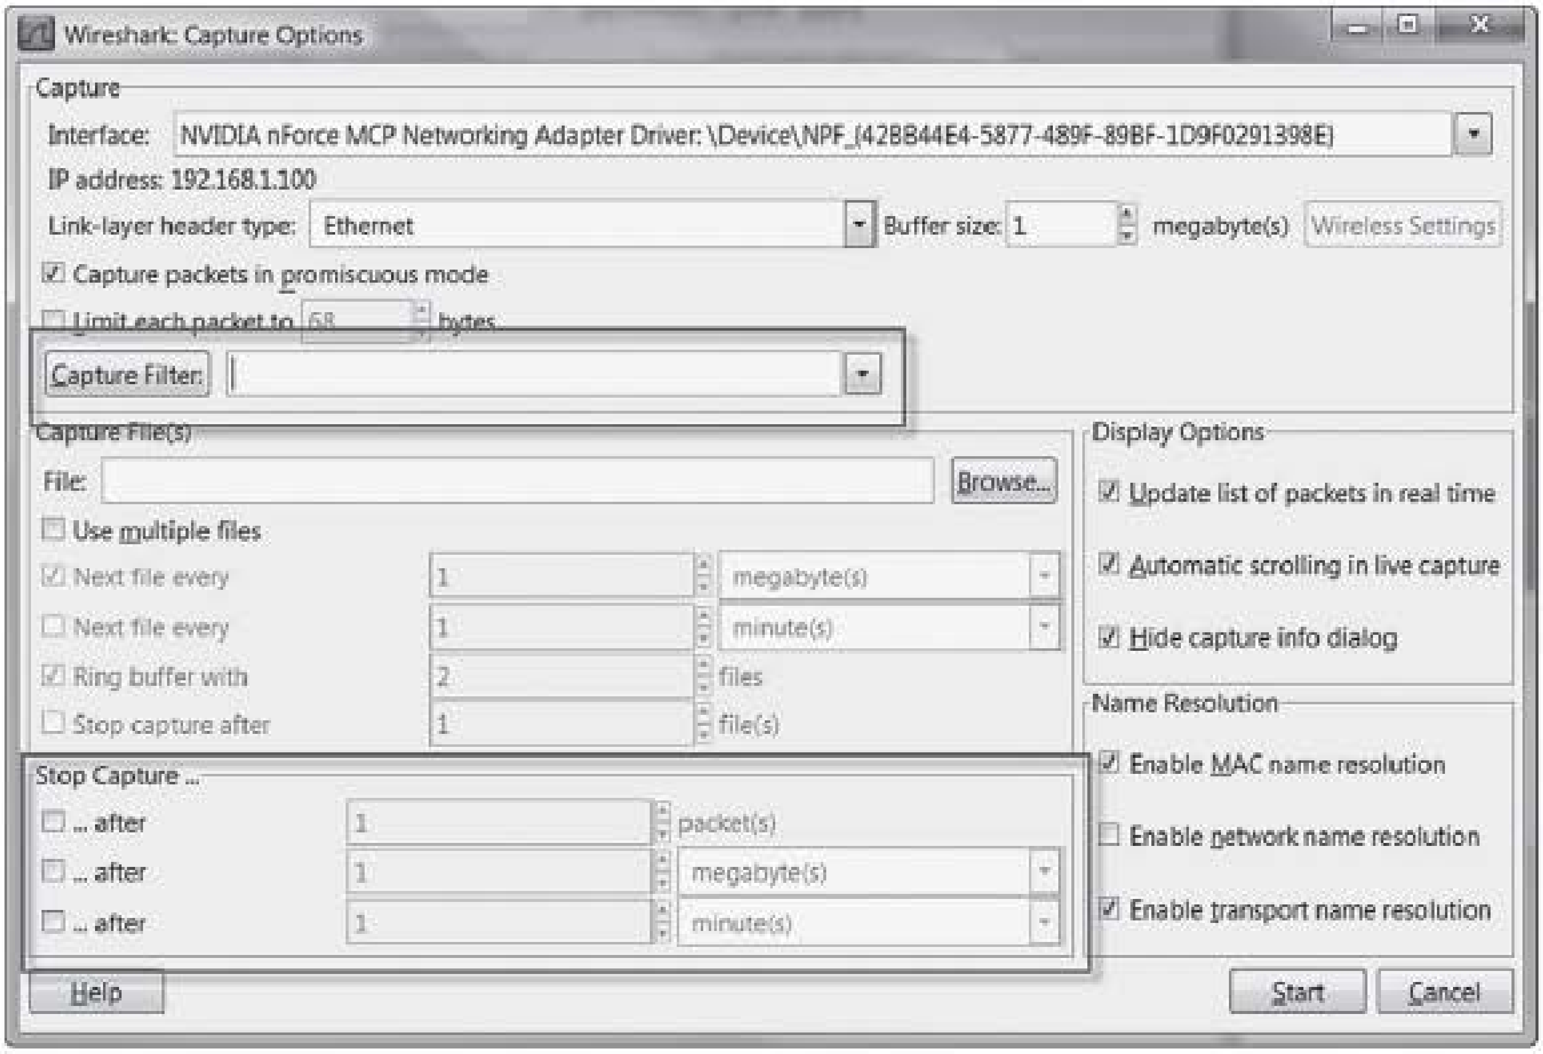

Capture

Interface: NVIDIA nForce MCP Networking Adapter Driver: \Device\NPF_{428B44E4-5877-489F-89BF-1D9F0291398E}

IP address: 192.168.1.100

Link-layer header type: Ethernet Buffer size: 1 megabyte(s)

☑ Capture packets in promiscuous mode

□ Limit each packet to 68 bytes

Capture File(s)

File:

□ Use multiple files

☑ Next file every 1 megabyte(s)

□ Next file every 1 minute(s)

☑ Ring buffer with 2 files

□ Stop capture after 1 file(s)

Stop Capture ...

□ ... after 1 packet(s)

□ ... after 1 megabyte(s)

□ ... after 1 minute(s)

Display Options

☑ Update list of packets in real time

☑ Automatic scrolling in live capture

☑ Hide capture info dialog

Name Resolution

☑ Enable MAC name resolution

□ Enable network name resolution

☑ Enable transport name resolution

FIGURE 8a-5 Wireshark Options

HANDS-ON EXERCISES

1. Do the following:

- Download Wireshark.
- Start Wireshark.
- Turn on Wireshark capture.
- Type a URL in your browser window (not Wikipedia.org).
- After a few seconds, stop the capture.
- Answer the following questions:

  1a. What URL did you use? What was the IP address of the webserver?

  1b. Find the frame in which your PC sent the SYN packet. List the source and destination IP address, the source and destination port numbers, and the header checksum.

  1c. Select the SYN/ACK packet. List the source and destination IP address, the source and destination port numbers, and the header checksum.

  1d. Select the packet that acknowledges the SYN/ACK segment. List the source and destination IP address, the source and destination port numbers, and the header checksum.

2. Change the options so that only packets you send are recorded. Do a capture. Click on the window containing Wireshark and hit Alt-Enter. This captures the window to your clipboard. Paste it into your homework.

Chapter **9**

TCP/IP Internetworking II

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Explain IPv4 subnet planning and do the calculations needed for working with subnet and host parts and deciding on part lengths.
- Do the same for IPv6.
- Explain the purposes of Network Address Translation (NAT) and how NAT operates.
- Explain in more detail than you learned in Chapter 1 how the Domain Name System (DNS) and the Dynamic Host Configuration Protocol (DHCP) operate.
- Describe the object model in the Simple Network Management Protocol (SNMP) and describe the enabling value of good security in the use of Set commands.
- Describe how the DNS was modified to deal with IPv6 addresses for host names.
- Describe how dynamic routing protocols work and how to select among alternative dynamic routing protocols.
- Describe the Internet Control Message Protocol (ICMP).
- Explain central concepts in IPsec (IP security), including its strategic importance, transport versus tunnel mode operation, ESP versus AH protection, security associations, important cryptographic methods and options, session initiation with IKE, and how IPsec compares to SSL.

INTRODUCTION

Chapter 8 covered core TCP/IP concepts. Now we focus on management and security. Although switched networks are (generally) capable of operating for long periods without intervention by network managers, TCP/IP internets require constant tuning and support. Designed to operate a worldwide service, TCP/IP supervisory protocols are extensive and complex. In addition, the TCP/IP protocols were born without security. Adding security retroactively has been difficult, but IPsec promises to add security strategically at the internet, transport, and application layers. IPsec will not solve all security problems, but its abilities are impressive.

IP SUBNETTING

IPv4 Subnet Planning

IPv4 addresses are 32 bits long. We saw in the last chapter that each organization is assigned a networking part. Organizations usually divide the remaining bits into a subnet part and a network part. Figure 9-1 shows the network parts assigned for the University of Hawai`i and a hypothetical ISP and how each decided to divide the remaining bits, over which they had control, into a subnet part and a host part. The University of Hawai`i divided the remaining bits into 8/8 subnet and host parts. The ISP divided its bits 16/8. However, these were choices. The university could have divided its 16 bits 6/10, and the ISP could have divided its 24 bits 12/12. In this section, we will see why the organization's choice of how many bits are assigned to its subnet and host parts is an important decision.

The N = 2^b–2 Rule With b bits, you can represent 2^b possibilities. Therefore, with 8 bits, you can represent 2^8 (256) possibilities. This would suggest that the university can have 256 subnets, each with 256 hosts. However, a network, subnet, or host part cannot be all 0s or all 1s.[1] Therefore, the university can have only 254 (256–2) subnets, each with only 254 hosts. Figure 9-2 illustrates these calculations.

Balancing Subnet and Host Part Sizes Selecting the sizes of the subnet and host parts is important. The larger the subnet part, the more subnets there will be. However, the larger the number of subnets there are, the fewer hosts each subnet can have. Finding a golden ratio of the two IP address part sizes requires careful thinking.

The University of Hawai`i's choice of 8-bit subnet and host parts was acceptable for many years because no college needed more than 254 hosts. In addition, the subnet

Organization	Network Part	Subnet Part	Host Part	Total Bits
University of Hawai`i Host	128.171(16 bits)	17 (8 bits)	13 (8 bits)	32
Hypothetical ISP Host	60 (8 bits)	33.22 (16 bits)	5 (8 bits)	32

FIGURE 9-1 Network, Subnet, and Host Part Length in an IPv4 Address

[1] If you have all 1s in an address part, this indicates that broadcasting should be used. All 0s parts are used by computers when they do not know their own addresses. All-zero address parts are only used as the source IP addresses in messages sent from a client to a DHCP server.

Step	Description	Example			
1	Total size of IP address (bits)	32			
2	Size of network part assigned to firm (bits)	16		8	
3	Remaining bits for firm to assign	16		24	
4	Selected subnet/host part sizes (bits)	8/8	6/10	12/12	8/16
5	Possible number of subnets (2^b-2)	254 (2^8-2)	62 (2^6-2)	4,094 $(2^{12}-2)$	254 (2^8-2)
6	Possible number of hosts per subnet (2^b-2)	254 (2^8-2)	1,022 $(2^{10}-2)$	4,094 $(2^{12}-2)$	65,534 $(2^{16}-2)$

FIGURE 9-2 IPv4 Subnetting

mask (255.255.255.0) was very simple, breaking at 8-bit boundaries. This made it easy to see which hosts were on which subnets. The host at 128.171.17.5, for instance, was on the 17th subnet. If the subnet mask did not break at an 8-bit boundary, this would not be possible, as we will see later.

Today, however, many colleges in the university have more than 254 computers, so the limit of 254 hosts has become a problem. Several colleges have now been given two subnet numbers. These colleges must connect their two subnets with a router so hosts on the two subnets can communicate. This is expensive and a little awkward.

The university might have been better served had it selected a smaller subnet part, say 6 bits. As Figure 9-2 shows, this would have allowed 62 college subnets, which probably would have been sufficient. A 6-bit subnet part would give a 10-bit host part, allowing 1,022 hosts per subnet. This would be ample for several years to come. However, it would no longer be possible to look at an IPv4 address in dotted decimal notation and see immediately what subnet it is on.

Test Your Understanding

1. a) If a subnet part is X bits long, how many subnets can you have? b) If you have a subnet part of 9 bits, how many subnets can you have? (Answer: 510 subnets) c) If you have a subnet part of 6 bits, how many subnets can you have? d) If your network part is 16 bits long, how many hosts can you have per subnet?
2. a) Your firm has an 8-bit network part. If you need at least 250 subnets, what must your subnet part size be? (Check figure: 8 bits) b) Continuing the last question part, how many hosts can you have per subnet? (Check figure: 256 hosts per subnet) c) If your firm has an 18-bit network part and you need at least 16 subnets, what must your subnet part size be? d) Continuing the last question part, how many hosts can you have per subnet? e) Your firm has a 22-bit network part. What subnet part would you select to give at least 10 subnets? f) Continuing the last question part, how many hosts can you have per subnet? g) If the University of Hawai`i had chosen a 6-bit subnet size, how many subnets could it have had? h) How many hosts per subnet? i) If the ISP had chosen a 10-bit subnet size, how many subnets could it have had? j) How many hosts per subnet?

Terminology

IPv4	IPv6	IPv4 Part Length	IPv6 Part Length for Global Unicast IPv6 Address
Network Part	Routing Prefix	Variable	Variable
Subnet Part	Subnet ID	Variable	Variable
Host Part	Interface ID	Variable	64 bits
		Total: 32 bits	Total: 128 bits

Routing Prefix and Subnet ID

Subnet ID has a fixed length of 64 bits

Total length of routing prefix and subnet ID is 64 bits

If the routing prefix is 20 bits, the subnet ID must be 44 bits long

FIGURE 9-3 IPv6 Subnetting

IPv6 Subnetting

Subnetting in IPv6 is similar, but terminology and some concepts are quite different. Figure 9-3 summarizes key changes in how IPv4 and IPv6 are divided into three parts. In most cases, this is just a matter of terminology. However, there are additional considerations in subnetting.[2]

The Three Parts Figure 9-3 shows the IPv6 counterpart of the IPv4 *network part, subnet part, and host part*. It uses similar concepts.

- The counterpart of the network part is the *routing prefix*. The **routing prefix** lets routers on the Internet route packets to an organization. It is a prefix because it is the first part.
- The equivalent of the IPv4 subnet part is the *subnet ID*. The **subnet ID** lets routers within a firm deliver packets to the correct subnets within the firm.
- The equivalent of the IPv4 host part is the *interface ID*. The **interface ID** identifies an individual host in the firm.[3]

The Fixed-Length 64-Bit Interface ID In IPv4, the size of the host part varies. In contrast, the size of the interface ID in global unicast IPv6 addresses is fixed at 64 bits. It may seem wasteful to "use up" half of all bits in the IPv6 addresses to designate a host. However, with 64 bits left for the routing prefix and the subnet ID, there are still 1.8×10^{19} possibilities for the routing prefix and subnet ID.

The size of the interface ID in global unicast IPv6 addresses is fixed at 64 bits.

[2] Technically speaking, this section refers to IPv6 global unicast addresses, which are addresses for a packet going from one host to another across the global Internet. There are other types of addresses, such as multicast addresses, which are sent from one host to multiple destination hosts. Nearly all IPv6 addresses on the Internet, however, are global unicast addresses.

[3] A host can have multiple interfaces to the Internet. This is not common for clients and servers. It is almost always the case for routers. Each router interface connects to a different network or subnet.

Routing Prefix and Subnet ID Figure 9-3 indicates that the routing prefix and subnet ID are variable in length, although their total must be 64 bits because the interface ID has already consumed 64 of the 128 bits. For example, if the routing prefix is 20 bits, the subnet ID must be 44 bits (64 bits–20 bits). If an address registrar gives a firm a short routing prefix, then the company can have a large subnet ID and can therefore have many subnets. Smaller firms, needing fewer subnets, are given longer routing prefixes.

Creating the 64-Bit Interface ID Returning to the 64-bit interface ID, it would be nice to be able to use a host's data link layer address as the interface ID. However, the most common type of data link layer address, the EUI-48 address, is only 48 bits long. For IPv6, the IEEE 802 Committee has defined a way to create a **64-bit modified extended unique identifier (EUI-64)** based on the 48-bit EUI-48 address. This modified EUI-64 address fits the interface ID part.

Creating a modified EUI-64 address with 16 more bits from a EUI-48 address requires a series of steps, which Figure 9-4 illustrates. These steps are straightforward, and there are good technical reasons for each step. However, these technical reasons are complex and irrelevant to information systems professionals.[4]

- First, express the EUI-48 address in hexadecimal notation, remove the dashes, and change all letters to lowercase. So *A0-B1-C2-D3-E4-F5* squashes down to *a0b1c2d3e4f5*.
- Second, divide the 48 bits in half. Each half has 24 bits. In this case, the first half is *a0b1c2* and the second half is *d3e4f5*.
- Third, insert the hex symbol *fffe* between the two halves.[5] This raises the 48 bits to 64 bits.

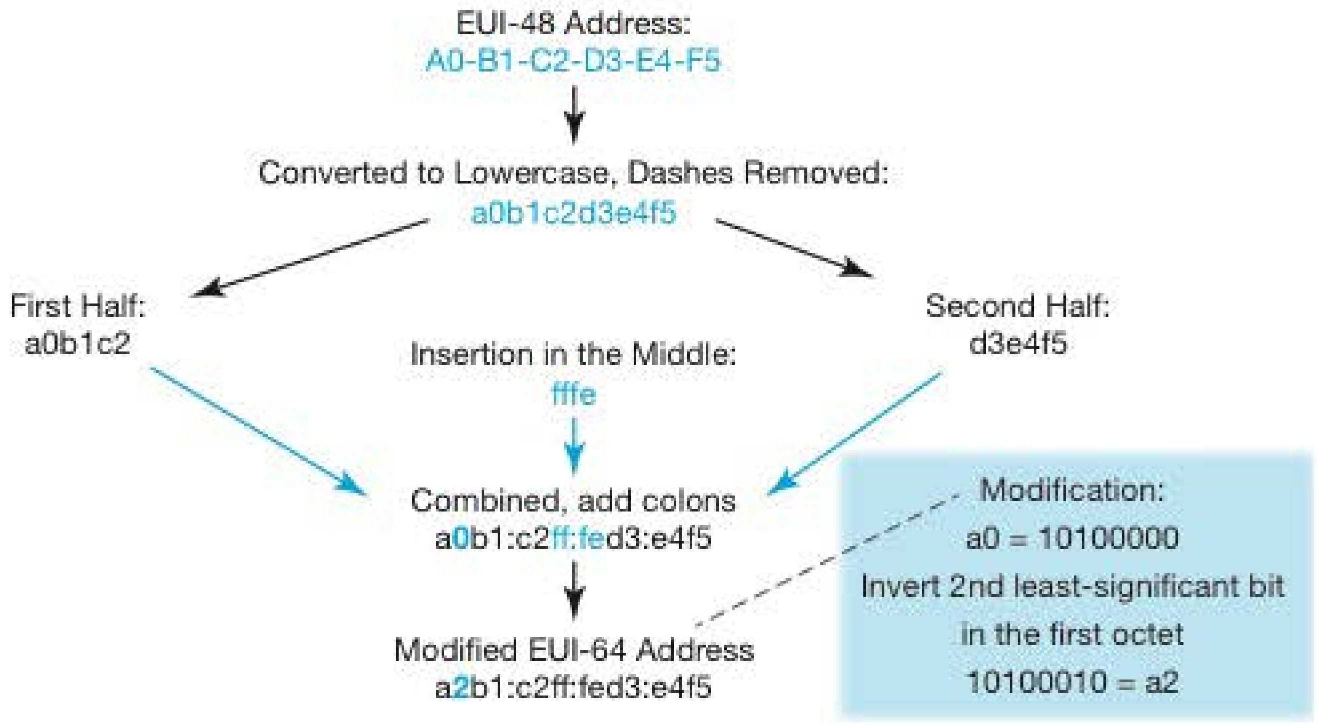

FIGURE 9-4 Converting an EUI-48 Address into a Modified EUI-64 Address

[4] Just think of them as arcane mystical protocols for joining an obscure secret society.

[5] See previous footnote.

- Fourth, write the first half, the new group, and the second half together. Now regroup them into four fields with four hex symbols apiece. Insert colons to separate these fields. The result is *a0b1:c2ff:fed3:e4f5*. (Note that a colon separates the *ff* and the *fe*. Use this as a crosscheck to make sure you have done things right.)
- Fifth, now we come to the *modified* part of the name. In this final step, the *second least-significant bit* (the second bit from the right end) in the first octet is inverted. For instance, the EUI-48 address in our example begins with a0. These two hex symbols constitute the first octet. In binary, they are *1010 0000*.[6] This must be changed to *1010 0010* by inverting the *second* least-significant bit—the bit that is the second from the right. (Inverting a bit means changing it to 1 if it is 0 and changing it to 0 if it is 1.) The inversion gives *a2* instead of *a0*. So the final modified EUI-64 is a**2***b1:c2ff:fed3:e4f5*.

Canonical Text Representation and Modified EUI-64 Addresses In Chapter 8, we saw that IPv6 addresses should be written in canonical text representation to provide a single standard way to reduce the length of written IPv6 addresses. In this chapter, we saw how to write the 64-bits interface ID of the IPv6 address in EUI-64 format.

Some students are tempted to go to canonical text representation immediately to shorten these 64-bit interface IDs as soon as they create them. However, the entire IPv6 address should be written without reduction before canonical text representation is used to shorten the address. Shortening the interface ID first can result in wrong choices being made in reducing consecutive 0000 fields to a single "::" abbreviation.

Do not convert the interface ID to canonical representation; only do so to full IPv6 addresses.

Test Your Understanding

3. a) What field in an IPv6 global unicast address corresponds to the network part of an IPv4 address? b) What field in an IPv6 global unicast address corresponds to the subnet part of an IPv4 address? c) If the subnet ID is 16 bits long, how long is the routing prefix? d) If you are a large company, do you want a long routing prefix or a short routing prefix? Explain. e) If your routing prefix is 16 bits, how long is your subnet ID? (Answer: 48 bits) f) If your routing prefix is 32 bits, how long is your subnet ID?
4. a) What field in a global unicast IP address corresponds to the host part of an IPv4 address? b) How long is this field? c) Convert the following EUI-48 address to a modified EUI-64 address: AA-00-00-FF-FF-00. (Answer: a**e**00:00ff:feff:ff00) d) Repeat for this EUI-48 address: 9B-E5-33-21-FF-0D.
5. Should you use canonical text representation to reduce the modified EUI-64 interface ID by itself, or should you do it only for the entire IPv6 address?

[6] See previous footnote.

OTHER TCP/IP STANDARDS

In this section, we look briefly at several other important TCP/IP standards that network administrators need to master.

Network Address Translation (NAT)

For security, firms must decide whether to allow people outside the corporation to learn their internal IP addresses. Doing so is a security risk. If attackers know internal IP addresses, this allows them to send attack packets from the outside world.

To prevent this, companies can use **Network Address Translation (NAT)**, which uses external IP addresses that are different from internal IP addresses used within the firm. If a sniffer learns these addresses, it cannot use this information to send attack packets to the internal IP address of a host.

NAT

- Puts false source IP addresses and port number in packets going out of the network

Expanding the Number of Available IP Addresses

- Companies receive a limited number of IP addresses from their ISPs
- There are roughly 4,000 possible ephemeral port numbers for each IP address
- So for each IP address, there can be 4,000 external connections
- If a firm is given 254 IP addresses, there can be roughly one million external connections (254 × 4,000)
- Even if each internal device averages several simultaneous external connections, there should not be a problem providing as many external IP connections as a firm desires

Security Reason for Using NAT

- External attackers can put sniffers outside the corporation
- Sniffers can learn IP addresses
- Attackers can send attacks to these addresses
- With NAT, attackers only learn false external IP addresses

Private IP addresses

- Can only be used inside firms
- 10.x.x.x
- 192.168.x.x (most popular)
- 172.16.x.x through 172.31.x.x

Transparency and Problems

- Transparent to the two hosts operating systems; each operates normally without knowing that NAT is used
- However, some applications have troubles with NAT
- There are work-arounds to these problems, but implementing NAT requires knowledge

FIGURE 9-5 Network Address Translation (NAT) (Study Figure)

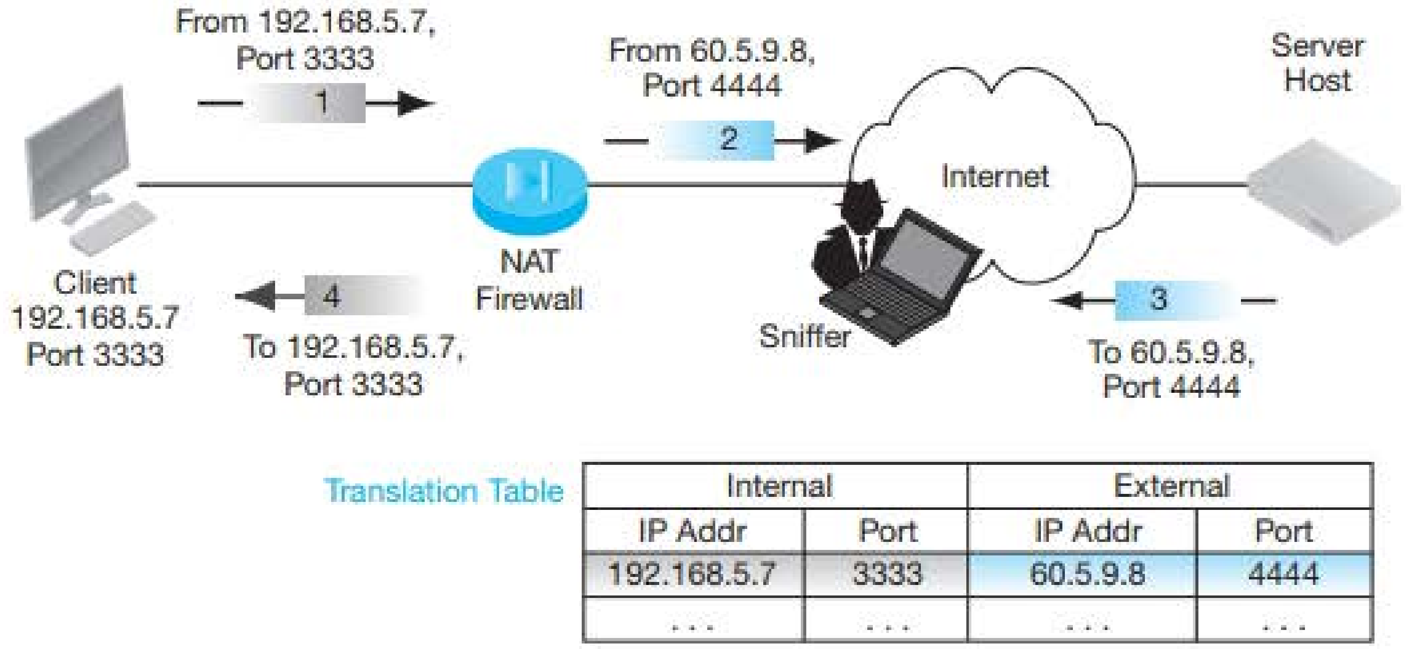

Internal		External	
IP Addr	Port	IP Addr	Port
192.168.5.7	3333	60.5.9.8	4444
. . .	. . .	. . .	. . .

FIGURE 9-6 Network Address Translation (NAT) Operation

NAT Operation Figure 9-6 shows how NAT works.

- An internal client host, 192.168.5.7, sends a packet to an external server host. The client operating system randomly generates the source port number 3333. As we saw in Chapter 2, this is an ephemeral port number that the source client host made up for this connection.
- The source socket in this packet is therefore 192.168.5.7:3333.
- When the NAT firewall at the border receives the packet, it makes up a new row in its translation table. It places the internal IP address and port number in the table. It then generates a new external source IP address and external source port number. These are 60.5.9.8 and 4444, respectively.
- Finally, the NAT firewall sends the modified packet to the external host.

Packets sent back from the external host have 60.5.9.8 in their destination IP address fields and 4444 in their destination port number fields. The NAT firewall looks these values up in its translation table, replaces the external values with the internal values, and sends them on to the client PC.

NAT and Security Figure 9-6 shows how NAT brings security. An attacker may be able to install a **sniffer program** beyond the corporation's NAT firewall. This sniffer will be able to read all packets coming out of the firm.

With NAT, an eavesdropper only learns false (external) IP addresses and false port numbers. If an attacker can attack immediately, it can send packets to the external IP addresses and port numbers, and the NAT firewall will pass them on to the internal host. However, it is rarely possible to act immediately, and NAT rows are only kept active for a few seconds or minutes. NAT provides a surprising amount of security despite its simple operation. Security professionals note that this is not very strong security, like encryption. However, it provides a substantial protection unless the attacker mounts a very sophisticated attack.

Expanding the Effective Number of IP Addresses An equally important reason for using NAT is to permit a firm to have many more internal IP addresses than its ISP gives it. Suppose that an ISP only gives a firm 254 IPv4 addresses by giving it a network part with 24 bits. In this case, the firm would not do subnetting. It would use all 8 bits for the host part. Without NAT, the firm can only have 254 PCs simultaneously using the Internet.

However, there are approximately 4,000 ephemeral client port numbers and therefore 4,000 possible external connections for *each* of the 254 public IPv4 addresses. This gives a million external connections (4,000 times 254). NAT can map these millions of connections into any combination of hosts and connections per host that it wishes. For example, it could map these connections to 100,000 internal hosts, each with 10 external connections.

Using Private IP Addresses To support NAT, the Internet Assigned Numbers Authority (IANA) has created three sets of **private IP address ranges** that can only be used *within* firms. These are the three ranges:

- 10.x.x.x
- 192.168.x.x
- 172.16.x.x through 172.31.x.x

The 192.168.x.x private IP address range is the most popular because it allows companies to use 255.255.0.0 and 255.255.255.0 network and subnet masks, respectively. These break at convenient 8-bit boundaries. However, the other two private IP address ranges also are widely used.

Transparency and Problems A nice result of the way IPsec operates is that it is transparent to the operating systems of the two hosts involved. The source host merely transmits normally, and the destination host does the same. There is no need to modify the hosts in any way.

At the same time, some applications have problems with NAT. These applications need to know the true IP addresses of the internal host. One example is IPsec, which we will see later. There are work-arounds for all these problems, but NAT requires considerable knowledge to use it effectively in corporations. Your home access router also uses NAT to allow you to have more than one internal host, but few problems occur.

Test Your Understanding

6. a) Describe NAT operation. b) What are the two benefits of NAT? c) How does NAT enhance security? d) How does NAT allow a firm to deal with a shortage of IP addresses given to it by its ISP? e) How are private IP address ranges used? f) Is NAT transparent to the operating systems of the two hosts involved? g) To all applications?

The Domain Name System (DNS)

We saw in Chapter 1 that when a user types a target host's host name, the user's PC will contact the local Domain Name System (DNS) server. This local DNS server will send the IP address for the target host back to the originating host. The user's PC can then send IP packets to the target host. In this chapter, we will add a few more elements to the picture.

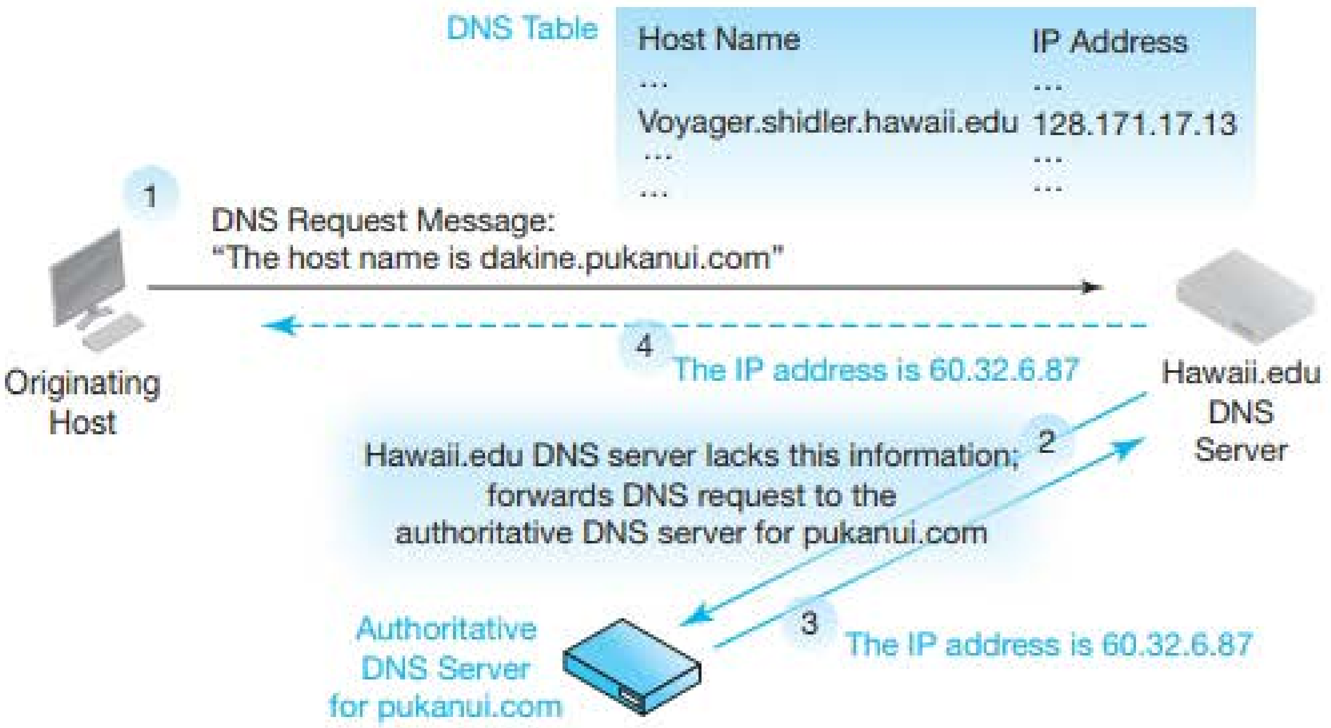
DNS Table
Host Name
IP Address
...
...
Voyager.shidler.hawaii.edu
128.171.17.13
...
...
...
...
1
DNS Request Message:
"The host name is dakine.pukanui.com"
Originating Host
4
The IP address is 60.32.6.87
Hawaii.edu DNS Server
2
Hawaii.edu DNS server lacks this information; forwards DNS request to the authoritative DNS server for pukanui.com
Authoritative DNS Server for pukanui.com
3
The IP address is 60.32.6.87

FIGURE 9-7 Domain Name System (DNS) Lookup

IP Address Lookup Figure 9-7 looks at how a DNS provides an IP address when a host sends a DNS request message specifying a host name. In the figure, the host name is dakine.pukanui.com.[7] In many cases, as we saw in Chapter 1, the local DNS server will know the IP address and send it back.

For a host name in another domain (pukanui.com instead of hawaii.edu), however, the local DNS host may not know the target host's IP address. (DNS servers are only required to know host names in their own domains, although they know many others.) In this example, the Hawaii.edu DNS server does not know the IP address of dakine.pukanui.com. To satisfy the originating host's request, the Hawaii.edu DNS server finds the **authoritative DNS server** for the domain containing the host name. In the figure, dakine.pukanui.com's DNS server is authoritative for the pukanui.com domain. The local Hawaii.edu DNS server will pass the DNS request message to this authoritative server. The pukanui.com DNS server will look up the IP address for dakine.pukanui.com and send it back to the local DNS server. The Hawaii.edu DNS server will in turn send the IP address to the originating host that sent the DNS request.

Test Your Understanding

7. a) What server will your local DNS server contact if it does not know the IP address of a host? b) Does the client know that his or her local DNS server contacted another DNS server to obtain an IP address? (This requires some thought and an answer beyond a simple yes or no.)

[7] I've been asked to explain this. *Dakine* in Hawai`i's Pidgeon is "that kind," a term to use when you can't remember what something is called ("Eh, hand me dakine."). It is pronounced "duh kin" with a long I and emphasis on the second syllable. In Hawai`ian, *puka* is a hole or empty space ("I have a puka in my shirt). *Nui* means big (opunuinui is an extra-big stomach.) So pukanui is big empty space. The first author actually owned pukanui.com. In my defense, it was cheap and I needed an example for the book.

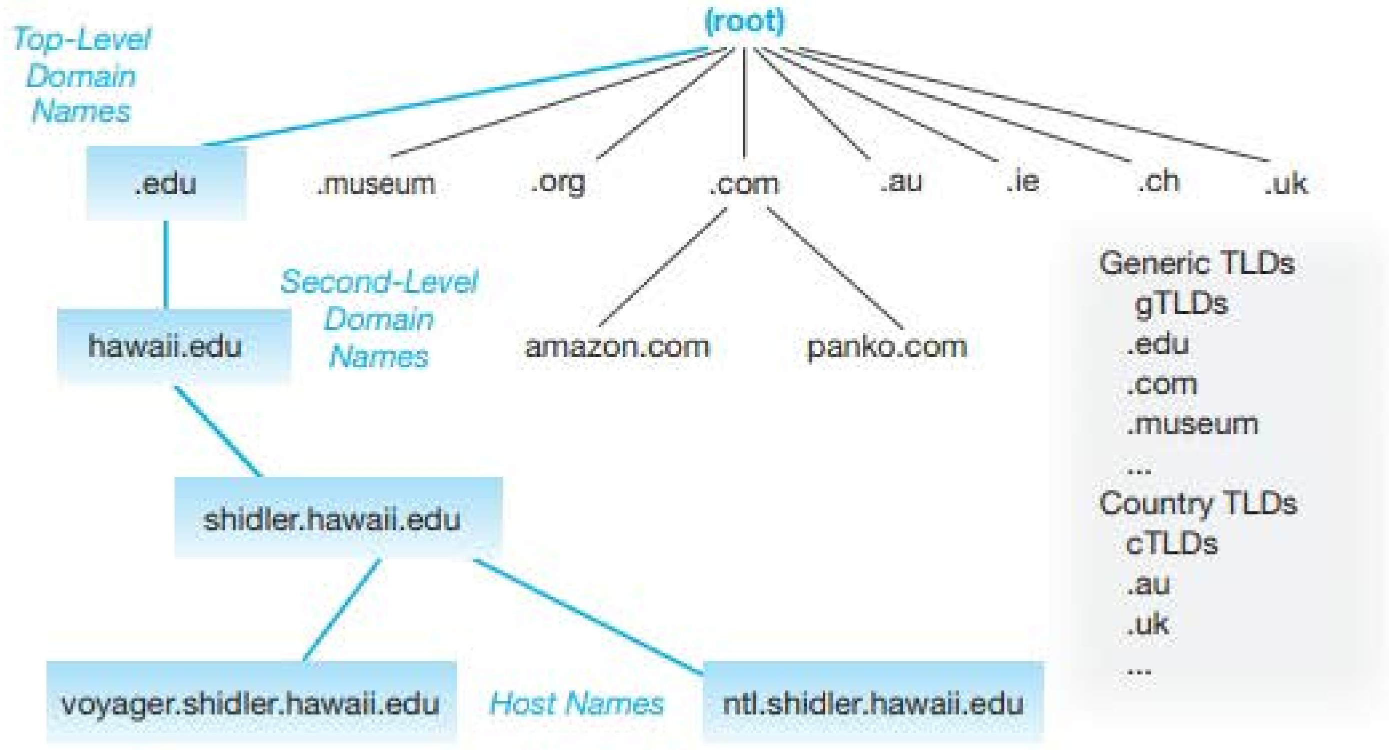

FIGURE 9-8 Domain Name System (DNS) Hierarchy

What Are Domains? Figure 9-8 shows that the **Domain Name System (DNS)** and its servers are not limited to providing IP addresses for host names. More generally, DNS is a general system for naming domains. A **domain** is a group of resources (routers, single networks, and hosts) under the control of an organization.

> *A domain is a group of resources (routers, single networks, and hosts) under the control of an organization.*

Root The figure shows that domains form a hierarchy, with host names at the bottom of the hierarchy. At the top of the DNS hierarchy is the root, which consists of all domain names. Thirteen **root DNS servers** keep overview information for the system.

Top-Level Domains Under the root are **top-level domains (TLDs)** that categorize the domain in one of two ways.

- **Country top-level domains (cTLDs)** specify the country of the domain owner. Examples are .uk, .ca, .ie, .au, .jp, .nl, .tv, .md, and .ch.
- **Generic top-level domains (gTLDs)** specify that the organization owning the name is a particular type of organization. The first gTLDs included .com, .edu, .net, .info, .gov, and .org. Later, the IANA added several more gTLDs, such as .name and .museum. In 2012, ICANN opened the naming system widely, permitting any organization to propose new generic top-level domains.

Note the distinction between the root and top-level domains. The root consists of all domains. It is not named as a level, however. If you are familiar with the UNIX operating system, the root directory concept is similar.

Also, note that it is possible for a domain to have two top-level designations, for instance, AAAA.com.ie. Most organizations, however, tend to use either a country TLD or a generic TLD.

Second-Level Domains Under top-level domains are **second-level domains**, which usually specify a particular organization (microsoft.com, hawaii.edu, tulsa.edu, cnn.com, etc.). Sometimes, however, specific products, such as movies, get their own second-level domain names. Competition for good second-level domain names is fierce. Organizations and individuals compete fiercely to get memorable second-level domains because this is how the public will reach them.[8]

Organizations and individuals compete fiercely to get second-level domains because this is how the public will reach them.

Companies get second-level domain names from **domain registrars** for nominal fees. However, getting a second-level domain name is only the beginning. Each organization that receives a second-level domain name must have a DNS server to host its domain name information. Large organizations have their own internal DNS servers that contain information on all subnet and host names. Individuals and small businesses that use webhosting services depend on the webhosting company to provide this DNS service.

In addition, of course, a second-level domain name does nothing for the firm until the firm buys or rents a webserver, builds a website, and pays an ISP to connect the website to the Internet. Then, of course, there is the matter of building the website.

Lower-Level Domains Domains can be further qualified. For instance, within hawaii.edu, which is the University of Hawai`i's second-level domain, there is a *shidler.hawaii.edu* domain. This is a third-level domain. It is the Shidler College of Business. Within shidler.hawaii.edu is *voyager.shidler.hawaii.edu*, which is a specific host within the college. It is a fourth-level domain.[9]

Test Your Understanding

8. a) Is the Domain Name System only used to send back IP addresses for given host names? Explain. b) What is a domain? c) Distinguish between the DNS root and top-level domains. d) What are the two types of top-level domains? e) Which level of domain name do corporations most wish to have? f) What are DNS root servers? g) How does a company or individual obtain a second-level domain name? h) What does a company need beyond obtaining a second-level domain name to have a website?

[8] The first author frequently gets requests to sell panko.com, usually from Japanese firms. I know. Bread crumbs. Japanese got bread from the Portuguese, and pan is the Portuguese word for bread. Ko, in Japanese, means little. Many Japanese girl's names end in ko. So panko is little bread.

[9] Host names are called fully qualified domain names (FQDNs).

Record Type	Information in the Record
A	Host Name-IPv4 Address Pair
AAAA	Host Name-IPv6 Address Pair

FIGURE 9-9 Two IP Address Lookup DNS Records for a Domain (Can Be a Single Host)

Domain Records We have seen DNS used for IP address lookups for particular host names. However, the Domain Name System holds much more information. For each domain, which can include an individual host, the DNS database contains multiple records. Each record serves a different purpose. Figure 9-9 shows the records used to look up IP addresses given host names. These are the A record for IPv4 addresses and the AAAA record for IPv6 addresses.[10]

To find an IP address for a host name, the DNS server searches through its records to find a match on the host name. It sends back both the IPv4 address and the IPv6 address associated with the host name, unless the host only has one address, in which case, the DNS server sends that address.

Test Your Understanding

9. a) Does a DNS server have one record for a particular domain (including a host), or does it have more than one? b) What is the purpose of the A record? c) What is the purpose of the AAAA record?

DHCP Servers

In Chapter 1, we saw that client PCs usually get their IP addresses from Dynamic Host Configuration Protocol (DHCP) servers. Now that we have looked at TCP/IP in more detail, we will see that DHCP servers do more than hand out IP addresses.

- Figure 9-10 shows that DHCP also provides the IP address of the default router—a router to send packets to if it does not have more specific information for sending a packet beyond the local subnet.

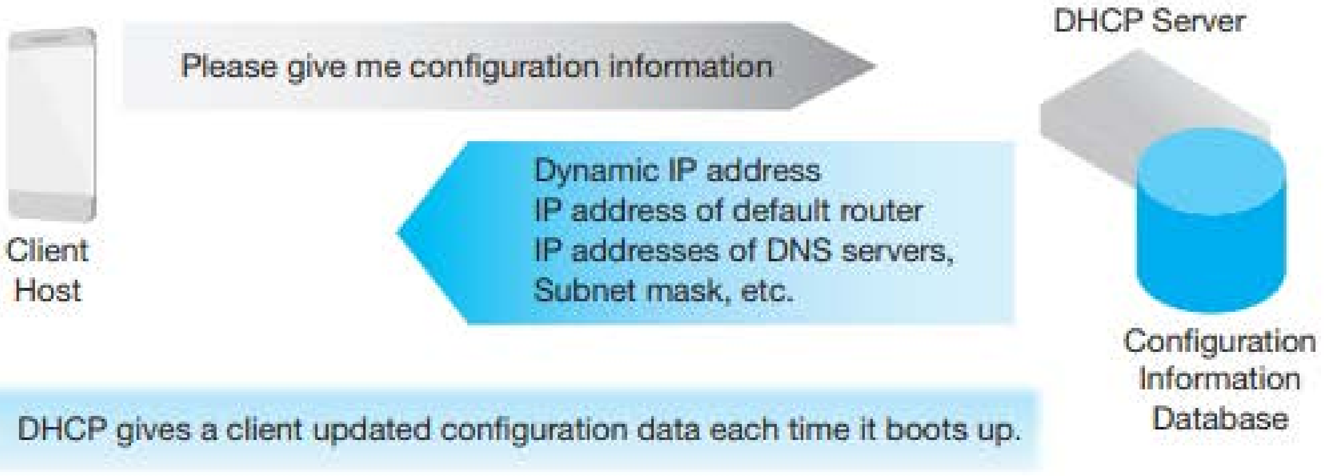
DHCP Server
Please give me configuration information
Dynamic IP address
IP address of default router
IP addresses of DNS servers,
Subnet mask, etc.
Client
Host
Configuration
Information
Database

DHCP gives a client updated configuration data each time it boots up.

FIGURE 9-10 Configuration Information in the Dynamic Host Configuration Protocol (DHCP)

[10] Four times as long, so AAAA instead of A. Standards people get bored a lot.

- It also tells the host one or more DNS server IP addresses. (Given the critical importance of DNS, most firms have multiple DNS servers.)
- Finally, DHCP tells the host its subnet mask, so that the host will know which IP addresses are in its own subnet and which requires sending the packet to a router for delivery beyond the subnet.

Up-to-Date Configuration Information DHCP guarantees that they have current configuration information each time they boot up, even if some aspects of the network have changed before booting up or if the device is moved to a different part of the network. If this configuration information had to be managed manually, all changes would cause serious extra work.

Test Your Understanding

10. a) What four pieces of configuration information does a DHCP server typically provide? b) Why is it useful to configure a client every time it boots up?

Simple Network Management Protocol (SNMP)

We saw the Simple Network Management Protocol in Chapter 4. We now look at the Simple Network Management Protocol (SNMP) in more detail, focusing on the schema of the management information base (MIB) and the security implications of the Set command.

The Management Information Base (MIB) When the manager retrieves information from agents on managed devices, it stores this information in a database called the **management information base (MIB)**. As in databases in general, "MIB"

SNMP Objects (see Figure 9-12)

- Not the managed devices themselves
- Objects are *specific pieces of information* about a managed device
- Information is stored in the management information base (MIB)

Set Commands

- Dangerous if used by attackers
- Many firms disable Set to thwart such attacks
- However, they give up the ability to manage remote resources without travel

Security as Enabler

- If a company has good security, it can enable Set
- This will save money
- In general, good security enables many network management tools that can save money and bring other benefits
- It can also enable applications that help employees do their jobs better

FIGURE 9-11 Simple Network Management Protocol (SNMP) (Study Figure)

refers both to the physical database and also to the schema (organization) of the information in the database. We will focus on the latter.

The MIB schema is not *relational*. Instead, the SNMP MIB schema is organized as a *hierarchy* of objects. The term *object* is a little confusing at first. An **object** is a piece of information about a managed device. The managed device itself is *not* an object. Figure 9-12 shows the basic schema for organizing SNMP objects.

An object is a piece of information about a managed device. The managed device itself is not an object.

- There is one set of objects for the system (switch, router, host, etc.) as a whole. For example, the manager may ask a router its system uptime—how long it has operated since its last reboot. If this is only a few minutes, the router may be suffering intermittent failures that cause it to crash and reboot frequently.
- There is also one set of IP objects, TCP or UDP objects, and ICMP objects. For example, the manager can ask the agent for a router if its routing object is On. If it is not, the router cannot act as a router. Rows discarded because of lack of memory is another useful object value to know. Also, if a router is discarding more than a tiny number of packets because its memory is full, it is time to add more memory. The number of errors will grow as traffic increases further, causing many retransmissions.
- A router may have multiple interfaces, and so will a switch (although switches usually call interfaces ports.) Each interface will have its own set of objects, including its speed and the number of errors it has experienced. If an interface has too many errors, it may have problems that need attention.

SNMP Set Security The SNMP *Set* command is very powerful. The manager can use Set to tell an agent to *change* the configuration of a managed device. If a router interface seems to be malfunctioning, for example, the manager can tell the agent to set the value of an interface to "testing." There is no need to travel to the object.

By allowing administrators to change devices remotely, the Set command can save companies a great deal of money by avoiding travel to fix problems. Unfortunately, many firms are reluctant to use Set commands because of security dangers. If attackers learn how to send Set commands to managed devices, the results could be catastrophic.

Companies that have strong security can enable Set and reap the benefits. Too often, strong security is viewed as a cost. However, strong security is also an enabler of SNMP and other systems that can save or make the organization a great deal of money.

Test Your Understanding

11. a) Explain the difference between managed devices and objects. b) List one object in each of the following areas: the system, IP, TCP, UDP, ICMP, and an interface. Explain how it might be used in network management. c) Why are firms often reluctant to use Set commands? d) How can good security be an enabler with SNMP?

Objects Are Pieces of Information About a Managed Device

- Objects are not the physical managed devices
- The SNMP MIB is organized as a hierarchy rather than as a relational database

System Objects

- System name
- System description
- System contact person
- System uptime (since last reboot)
- . . .

IP Objects

- Forwarding (for routers). Yes if forwarding (routing), No if not
- Subnet mask
- Default time to live
- Traffic statistics
- Number of discards because of resource limitations
- Number of discards because could not find route
- Number of rows in routing table
- Rows discarded because of lack of memory
- Individual row data
- . . .

TCP Objects

- Maximum/minimum retransmission time
- Maximum number of TCP connections allowed
- Opens/failed connections/resets
- Segments sent
- Segments retransmitted
- Errors in incoming segments
- No open port available errors
- Traffic data on individual connections (sockets, states)
- . . .

UDP Objects

- Errors: no application on requested port
- Traffic statistics
- . . .

ICMP Objects

- Number of errors of various types

Interface Objects (One per Interface)

- Type (e.g., 71 is 802.11)
- Status: up/down/testing
- Speed
- MTU (maximum transmission unit—the maximum packet size)
- Traffic statistics: octets, unicast/broadcast/multicast packets
- Errors: discards, unknown protocols, etc.
- . . .

FIGURE 9-12 SNMP MIB Hierarchical Object Model

Dynamic Routing Protocols

How does a router get the information in its routing table? It is possible to enter routes manually. However, that approach does not scale to the enormous size of the Internet. Instead, as Figure 9-13 shows, routers constantly exchange routing table information with one another using **dynamic routing protocols**.[11,12]

Interior Dynamic Protocols: OSPF and EIGRP Recall from Chapter 1 that the Internet consists of many networks owned by different organizations. *Within* an individual organization's network or internet, the organization decides which **interior dynamic routing protocol** to use for its internal routers, as shown in Figure 9-13. There are two popular interior dynamic routing protocols.[13] Each has relative strengths and weaknesses.

- **Open Shortest Path First (OSPF).** For interior routing, the IETF created the **Open Shortest Path First (OSPF)** dynamic routing protocol. OSPF is very efficient, having a complex metric based on a mixture of cost, throughput, and traffic delays. It also offers strong security. However, it only does TCP/IP routing. Although TCP/IP is dominant today, many corporations still have legacy protocols from other standards architectures, such as IBM's SNA architecture and Novel's SPX/IPX. Corporations cannot use OSPF for routing in these other architectures.

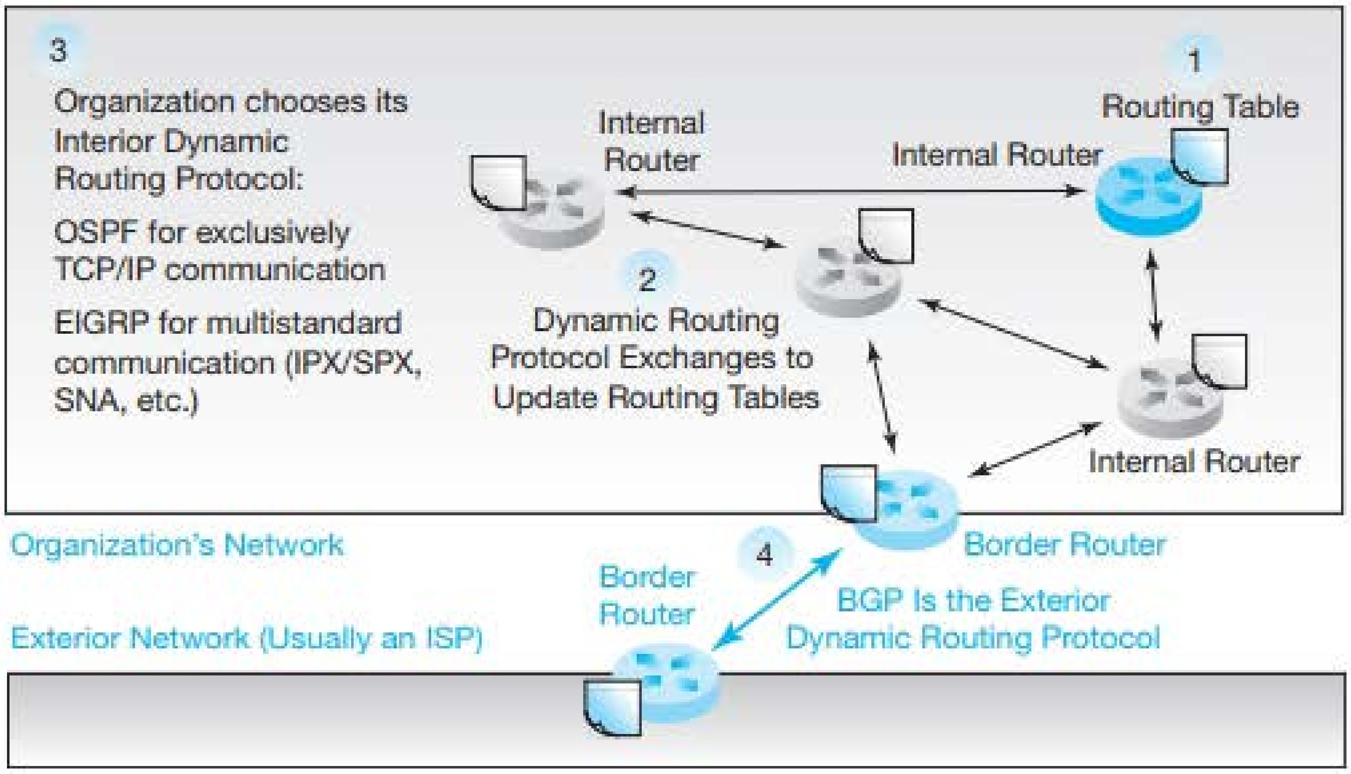
3
Organization chooses its Interior Dynamic Routing Protocol:
OSPF for exclusively TCP/IP communication
EIGRP for multistandard communication (IPX/SPX, SNA, etc.)
Internal Router
2
Dynamic Routing Protocol Exchanges to Update Routing Tables
Internal Router
1
Routing Table
Internal Router
Border Router
Organization's Network
4
Border Router
BGP Is the Exterior Dynamic Routing Protocol
Exterior Network (Usually an ISP)

FIGURE 9-13 Dynamic Routing Protocols

[11] Note that TCP/IP uses the term *routing* in two different but related ways. First, we saw earlier that the process of forwarding arriving packets is called routing. Second, the process of exchanging information for building routing tables is also called routing. The IETF sometimes is not fastidious about terminology.

[12] To give an analogy, college students talk to other students to determine which classes they should take or avoid.

[13] A third interior dynamic routing protocol is RIP, the Routing Information Protocol. RIP is simpler than OSPF or EIGRP and was once popular. However, its almost complete lack of security features makes it an unacceptable choice today. It is commonly referred to today as "rest in peace."

- **Enhanced Interior Gateway Routing Protocol (EIGRP).** Cisco Systems is the dominant manufacturer of routers. Cisco has its own proprietary interior dynamic routing protocol for large internets—the **Enhanced Interior Gateway Routing Protocol (EIGRP)**. The term **gateway** is another term for *router*. EIGRP is comparable to OSPF, but unlike OSPF, it can also route non-TCP/IP traffic.

Exterior Dynamic Protocol: BGP For communication outside the organization's network, the organization no longer has a choice. It *must* use the **exterior dynamic routing protocol** required by the external network to which it is connected. (This exterior network is usually an ISP.) The almost-universal exterior dynamic routing protocol is the **Border Gateway Protocol (BGP)**. Again, gateway is another term for router.

"Gateway" is another term for "router."

Test Your Understanding

12. a) What is the purpose of dynamic routing protocols? b) For its own network, can an organization choose its interior dynamic routing protocol? c) What is the IETF interior dynamic routing protocol? d) When might you use EIGRP as your interior dynamic routing protocol? e) May a company select the routing protocol its border router uses to communicate with the outside world? f) What is the almost-universal exterior dynamic routing protocol? g) What is a gateway?

Internet Control Message Protocol (ICMP) for Supervisory Messages at the Internet Layer

Supervisory Messages at the Internet Layer IP is only concerned with packet delivery. For supervisory messages at the internet layer, the IETF created the **Internet Control Message Protocol (ICMP)**. IP and ICMP work closely together. As Figure 9-14 shows, IP encapsulates ICMP messages in the IP data field, delivering them to their target host or router. There are no higher-layer headers or messages.

Error Advisement IP is an unreliable protocol. It offers no error correction. If the router or the destination host finds an error, it discards the packet. Although there is no retransmission, the router or host that finds the error may send an ICMP error message to the source device to inform it that an error has occurred, as in Figure 9-14. The **ICMP error advisement** message contains type and code values indicating what the problem is. For example, a host unreachable message is Type 3/Code 1.

Error advisement is not error correction. There is no mechanism within IP or ICMP for the retransmission of lost or damaged packets. ICMP error messages are only sent to help the sending process or its human user diagnose problems. They do not make IP reliable.

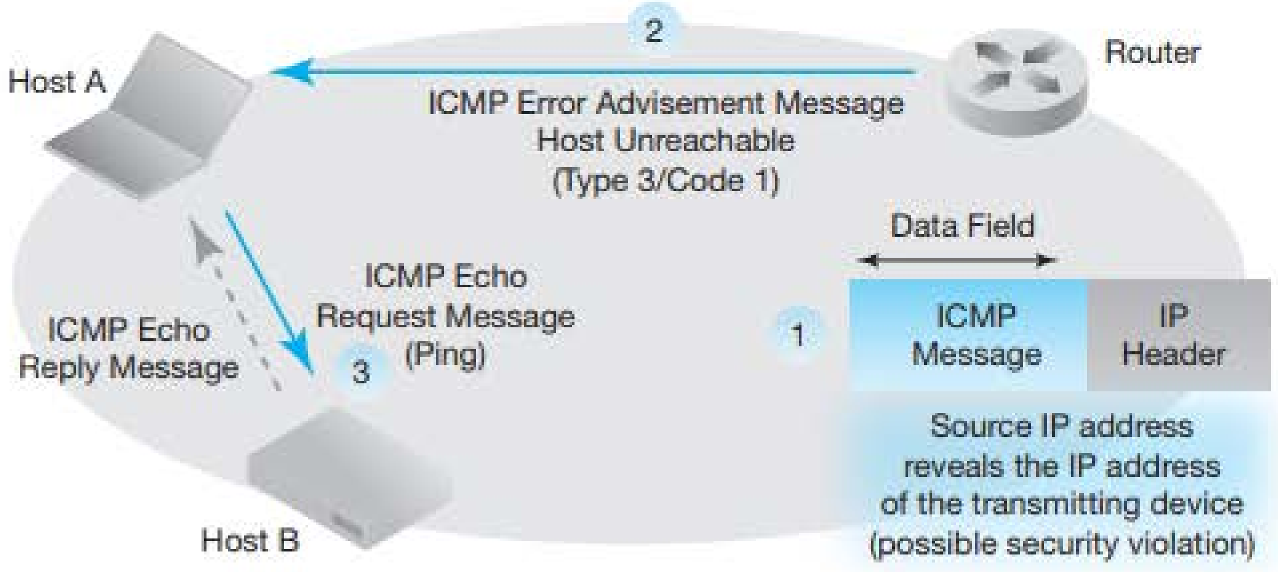

FIGURE 9-14 Internet Control Message Protocol (ICMP) For Supervisory Messages at the Internet Layer

Sending error advisement messages is not mandatory when errors occur. For security reasons, many firms filter out ICMP error advisement messages at their borders because hackers can exploit the information contained in them. Most obviously, the ICMP message will be carried in a packet that contains the IP address of the sending router or other device. If adversaries have an **exploit** (attack method) to use against routers, they have a target IP address for their attacks.

Echo (Ping) ICMP also offers **ICMP control messages**, which direct a device to change how it operates. The most widely used ICMP control messages are the ICMP **echo request** and **echo reply** messages. As we saw in Chapters 1 and 4, one host can use these messages to "ping" another host. As in the case of error response messages, the IP header for the echo reply message reveals the presence of a potential target at the source IP address. Again, many firms do not allow echo reply messages to go outside the corporation.

Test Your Understanding

13. a) For what general class of messages at what layer is ICMP used? b) Distinguish between ICMP error advisement and control messages. c) What two ICMP message types are used in ping? d) What security concern do ICMP error advisement messages and echo response messages bring?

IPsec

The Internet was born without a plan for security. Jon Postel, who edited the Internet Protocol RFC and several others, once reminisced that security threats were infrequent in the late 1970s and early 1980s. Just getting the basic protocols working took all the energy that developers had. Today, of course, security is critical, and although security in Internet protocols has improved, the standards are not everything we want, products do not implement everything the standards provide, and individual organizations have a difficult time not making mistakes in implementing the complex security facilities that are available.

Core IPsec Principles

Today, the Internet Engineering Task Force is moving at full speed to integrate strong security into its standards. Every new Request for Comments (RFC) must have a security section that lays out what security is available and what security issues remain unaddressed for the standard. More important, security has been enhanced in a broad spectrum of Internet standards. However, a piecemeal approach is confusing and leaves gaps for attackers. Many believe that the key security standards for the Internet will be those that are collectively called **Internet Protocol security (IPsec)**.

Encapsulating Security Payload (ESP) and Authentication Header (AH) Figure 9-15 shows that IPsec offers two basic protection mechanisms. One is the **Encapsulating Security Protocol (ESP)**. The other is the **Authentication Header (AH)**. The figure compares the protections they offer.

After looking at the figure, you are probably thinking, "Why on earth would anybody use AH?" ESP offers far more protections, and it includes authentication and integrity, which is the only protection offered by AH.[14] The answer, as you might suspect is, "They seldom do." Given the infrequency of AH usage, we will focus on ESP.

Given the infrequency of AH usage, we will focus on ESP.

ESP Transport and Tunnel Modes Figure 9-16 illustrates how IPsec protects communication using ESP (and AH, by the way). These are transport mode and tunnel modes. Figure 9-16 shows two packets. Both have an IP header and a data field. In both cases, the header is not protected but the data field is. Everything in it is protected without having to do anything to the different layers of content in the data field.

Protection	Encapsulating Security Payload (ESP)	Authentication Header (AH)
Authentication (and Integrity)	✓	✓
Confidentiality	✓	
Anti-Replay Protection	✓	
Other Protections	✓	

FIGURE 9-15 IPsec Encapsulating Security Payload (ESP) versus Authentication Header (AH)

[14] Historically, there were two main reasons to use AH. The first was that ESP's original design did not include authentication (which automatically gives message integrity). If you wanted both encryption and authentication, you had to use both AH and ESP. However, this has not been true for a long time. A more practical reason was that some countries outlaw encryption for confidentiality, allowing only authentication. This is rarely true today, although some countries require weak encryption. IPsec experts note that AH has some technical advantages that may be useful in specific circumstances, and they note that sometimes it is desirable to use both AH and ESP, but these situations are rare.

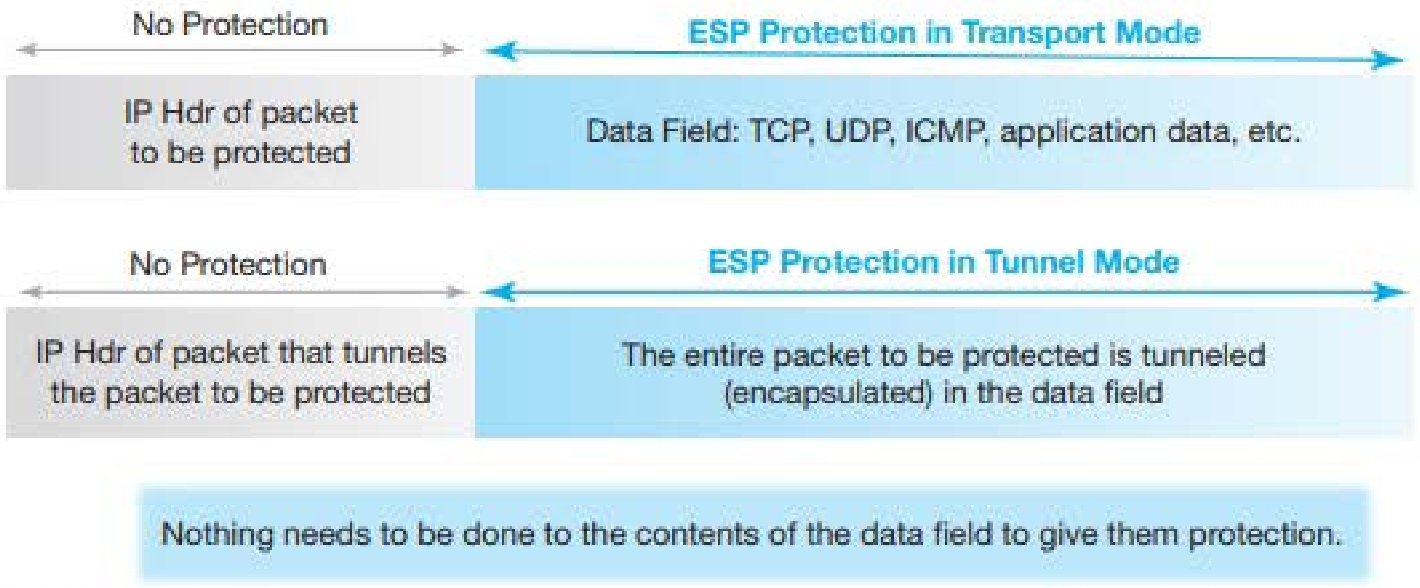
No Protection
ESP Protection in Transport Mode
IP Hdr of packet to be protected
Data Field: TCP, UDP, ICMP, application data, etc.
No Protection
ESP Protection in Tunnel Mode
IP Hdr of packet that tunnels the packet to be protected
The entire packet to be protected is tunneled (encapsulated) in the data field
Nothing needs to be done to the contents of the data field to give them protection.

FIGURE 9-16 ESP Protection in Transport and Tunnel Modes

So how are things different between the two? Look at the data fields.

- In **transport mode**, the data field holds the usual contents of an IP packet. It might contain a TCP header, a UDP header, application data, or anything else, such as an ICPM command.
- In **tunnel mode**, the data field has the entire packet protected. Tunnel mode places this packet inside another packet, which we will call the outer packet. This is called tunneling or encapsulating the packet to be protected.

Transport mode, then, leaves the header of the packet without protection. Tunnel mode fixes this limitation. Tunnel mode is the default in most IPsec implementations, but transport mode is also used somewhat.

Test Your Understanding

14. a) How does IPsec provide a great deal of protection? b) Why do we focus on ESP and not on AH? c) What is tunneling? d) Which protects more of the original IP packet, transport mode or tunnel mode? Explain how much more protection tunnel mode provides.

VPNs

IPsec creates a secure flow of packets between two endpoints. This secure flow is a **virtual private network (VPN)**. We already saw VPNs in Chapters 4 and 7. The two endpoints effectively have a secure private network connecting them. This "private network," of course, is only virtual. Figure 9-17 illustrates the three endpoint pairs that IPsec is designed to connect and protect.

- **Host-to-host VPNs** connect two hosts, often in the same site. In Figure 9-17, the VPN connects Client X with Server X. The security is handled by the two hosts, with no additional help.

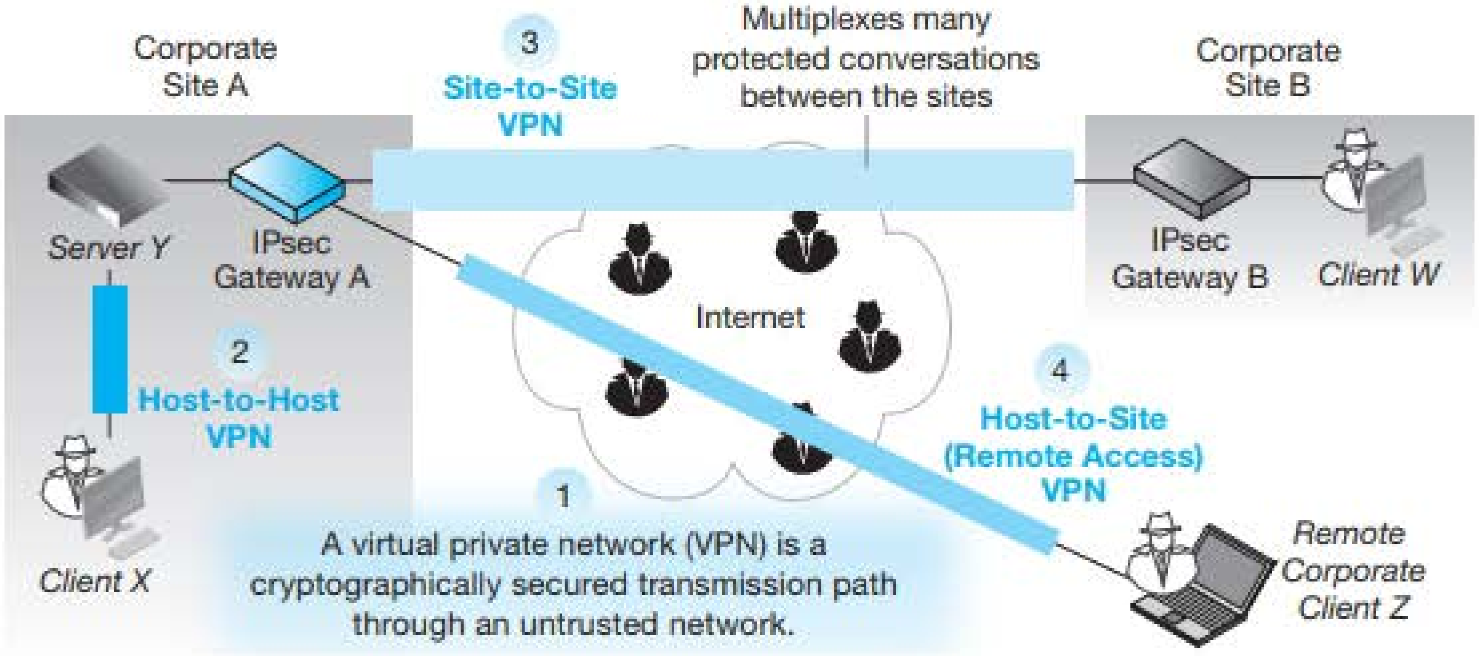
Corporate Site A
3
Site-to-Site VPN
Multiplexes many protected conversations between the sites
Corporate Site B
Server Y
IPsec Gateway A
IPsec Gateway B
Client W
Internet
2
Host-to-Host VPN
4
Host-to-Site (Remote Access) VPN
1
A virtual private network (VPN) is a cryptographically secured transmission path through an untrusted network.
Client X
Remote Corporate Client Z

FIGURE 9-17 IPsec VPNs

- **Site-to-site VPNs** connect two corporate sites.[15] The site-to-site VPN in the figure connects Corporate Site A and Corporate Site B. Host-to-Host VPNs only carry the traffic of the two hosts involved. Site-to-site VPNs, in contrast, multiplex many host-to-host transmissions between hosts at the two sites. The thicker box around the site-to-site network emphasizes this greater traffic.
- Finally, **host-to-site VPNs** protect traffic between a site and a remote corporate client who must reach the site via the Internet. Traffic is protected all the way between the remote corporate client and the site. This is also called a **remote access VPN**.

IPsec Gateways When one VPN endpoint is a site, the termination point is a device called an **IPsec gateway**. Site-to-site VPNs directly connect the IPsec gateways at the two sites. Host-to-site networks connect a remote client to the site's VPN gateway. VPN gateways can terminate many remote clients simultaneously.

Test Your Understanding

15. a) What three types of VPN does IPsec support? b) What nonhost device is a terminating point in site-to-site VPNs and host-to-site (remote access) VPNs? c) What is another term for "gateway?"

Applying ESP Protections

Error! Reference source not found. showed ESP protections broadly. In this subsection, we look at ESP in more detail.

ESP in Transport Mode Figure 9-18 shows how ESP is applied in transport mode. It shows that the sender adds an **ESP header**, an **ESP trailer**, and an **integrity check value (ICV)**. The ESP header comes after the IP packet's header, before the packet's data field. The ESP trailer and ICV come after the data field.

[15] It is also called a LAN-to-LAN or network-to-network VPN.

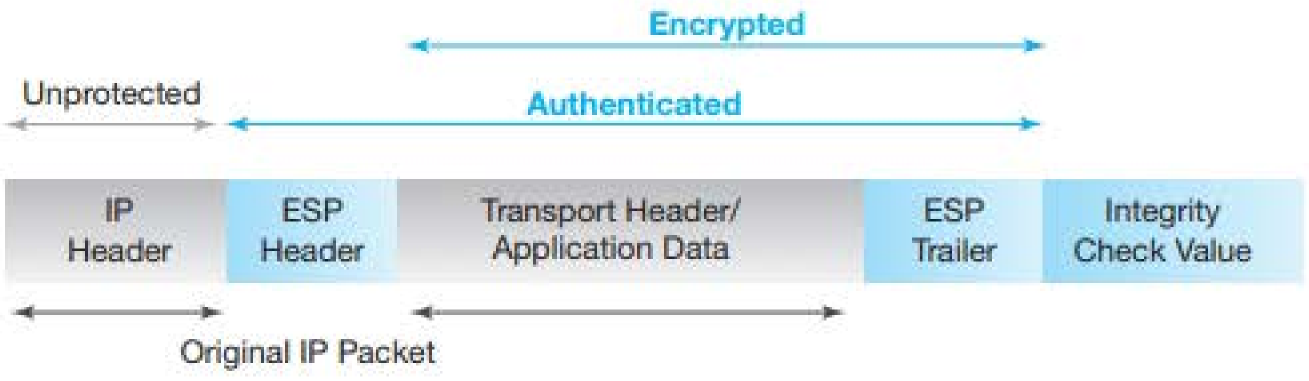

In transport mode, the packet is sent with additional fields for security.
In networking, transport means transmission.
The original packet is transported.
The IP header has no protection.
It can be read and changed en route.

FIGURE 9-18 ESP Additions to IPv4 in Transport Mode

Note that encryption only begins with the data field. ESP does not encrypt the IP header, the ESP header, or the ICV. The routers along the packet's path must be able to read the entire IP header to do their work. In turn, the destination host must be able to read the entire ESP header and integrity check value to authenticate and then decrypt an arriving packet. The ESP header is, however, authenticated.

ESP Additions in Tunnel Mode Figure 9-19 shows the additions that are made to implement ESP in tunnel mode. These again consist of an ESP header, an ESP trailer, and authentication data. The ESP header comes after the outer IP packet's header, before the outer packet's data field. The ESP trailer and authentication field come after the outer packet's data field.

This is exactly the way that EPS is implemented in transport mode. The difference is what information lies in the data field. In tunnel mode, again, the data field contains the entire original IP packet to be protected. Consequently, ESP in tunnel mode provides total protection for the protected packet.

In tunnel mode, the outer packet's header is sent in the clear. This means that each outer packet will reveal the IP address of the gateway to which the packet is going.

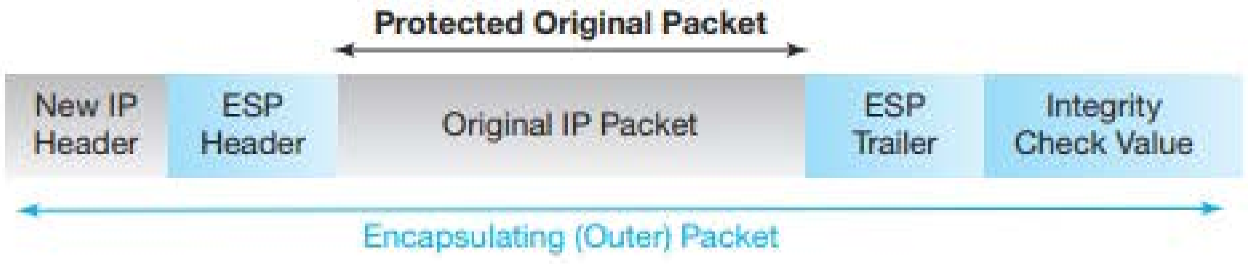

Tunneling is encapsulating a message inside another message for delivery.
In IPsec, the message being tunneled is the original IP packet.
(Tunneling is not about providing a secure "tunnel" through the Internet.)
Source and destination IP addresses in the original packet remain confidential.

FIGURE 9-19 ESP Additions to IPv4 in Tunnel Mode

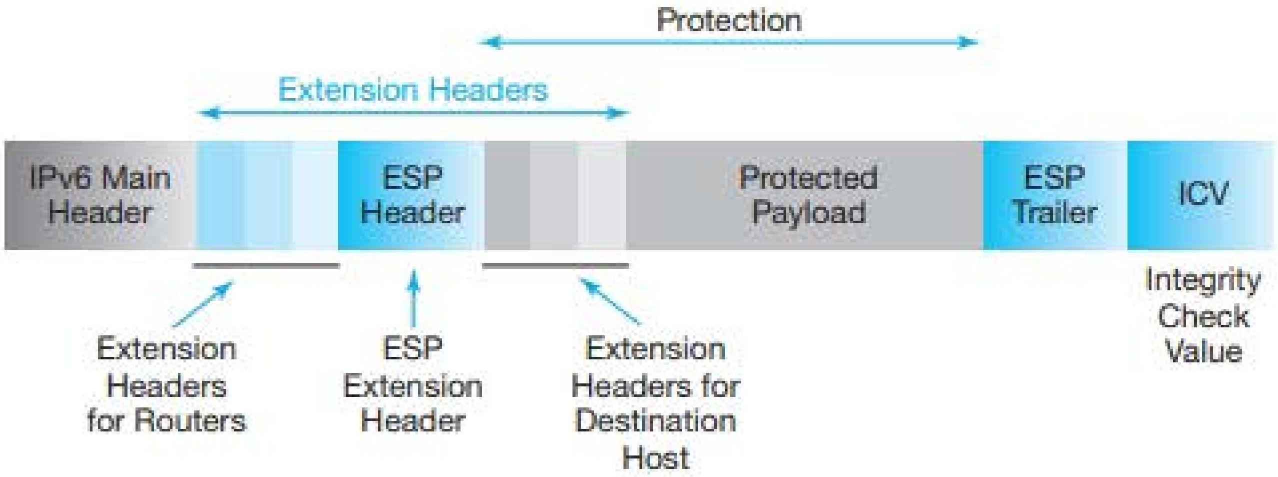

FIGURE 9-20 ESP Additions to IPv6 Packets in Tunnel Mode

However, gateway addresses are usually easy for attackers to learn anyway. Companies know that their IPsec gateways are critical single points of failure that are known to be very risky, so they are exceptionally hardened.

ESP Additions in IPv6 Figure 9-18 and Figure 9-19 show ESP additions for IPv4. Figure 9-20 shows that the main change in IPv6 comes in IPv6 extension headers. Some extension headers need to be read by routers along the packet's route. Others are only read by the destination host. The figure shows that the ESP header is normally placed after the hop-by-hop extension headers but before the destination host headers. This allows it to protect the data field (payload) plus any destination headers that routers along the way do not need to know.[16]

Test Your Understanding

16. a) What are the three added fields when IPsec ESP is used? b) What do they surround in transport mode for IPv4? c) What do they surround in tunnel mode for IPv4? d) Where is the ESP header placed in IPv6?

Security Associations (SAs)

Methods and Options When organizations implement IPsec, they want to be able to tailor it to each connection's specific situation. This requires choices, not one-size-fits-all protection. These options include whether ESP or AH will be used, whether IPsec will operate in tunnel or transport mode, what encryption method confidentiality will use, and what hashing[17] method authentication uses. Figure 9-21 shows IPsec's main encryption methods. Figure 9-22 does the same for hashing, which is a core part of authentication. Most cryptographic methods offer further options. For example, the

[16] ESP (and AH) headers were created as extension headers for IPv6. They were later added to IPv4

[17] Authentication methods require extensive processing, and this processing is directly correlated to the size of the message. Hashing addresses this problem by creating a string of bits that is much smaller than the entire packet. Hashing is applied to the long packet to produce a small bit string of fixed length (128 to 512 bits). An authentication method is *then applied to the hash, instead of the full packet*. This seems like a trick, but it gives about the same level of protection that authenticating the entire packet would do. The longer the hash, however, the greater the protection (and the longer the processing time).

Option	Name	Key Length (bits)*	Remarks
AES-192	Advanced Encryption Standard	192	Extremely strong
AES-256	Advanced Encryption Standard	256	Far stronger than AES-192
3DES	Triple Date Encryption Standard	168	Very strong but inefficient legacy standard
DES	Date Encryption Standard	56	Weak legacy standard Should not be used

*Each additional bit doubles the time needed to crack a key.

FIGURE 9-21 Common Encryption Methods and Options for Confidentiality in IPsec

Option	Name	Hash Length (bits)*	Remarks
MD-5	Message Digest	128	Weak legacy standard Should not be used
SHA-1	Secure Hash Algorithm	160	Weak legacy standard Should not be used
SHA2-224	Secure Hash Algorithm	224	Strong to extremely strong
SHA2-256	Secure Hash Algorithm	256	
SHA2-384	Secure Hash Algorithm	384	
SHA2-512	Secure Hash Algorithm	512	

*Longer hashes provide better authentication.

FIGURE 9-22 Common Hashing Methods and Options for Authentication in IPsec

AES encryption method can optionally have 128-bit keys, 192-bit keys, or 256-bit keys. These options can have major impact on the security that a method provides.

Security Associations A **security association (SA)** documents how the two parties will implement IPsec protection, including what methods they will use for different cryptographic purposes and what options will be used with these methods.

> *A security association (SA) documents how the two parties will implement IPsec protection, including what methods they will use for different cryptographic purposes and what options will be used with these methods.*

Figure 9-23 shows security associations between two hypothetical hosts.

- For transmission from Host A to Host B, the SA specifies that Host A will use ESP in tunnel mode. For confidentiality, Host A will encrypt with AES-192. Host A will use SHA2-224 for authentication. (Yes, we will look at what these terms mean a little later.)

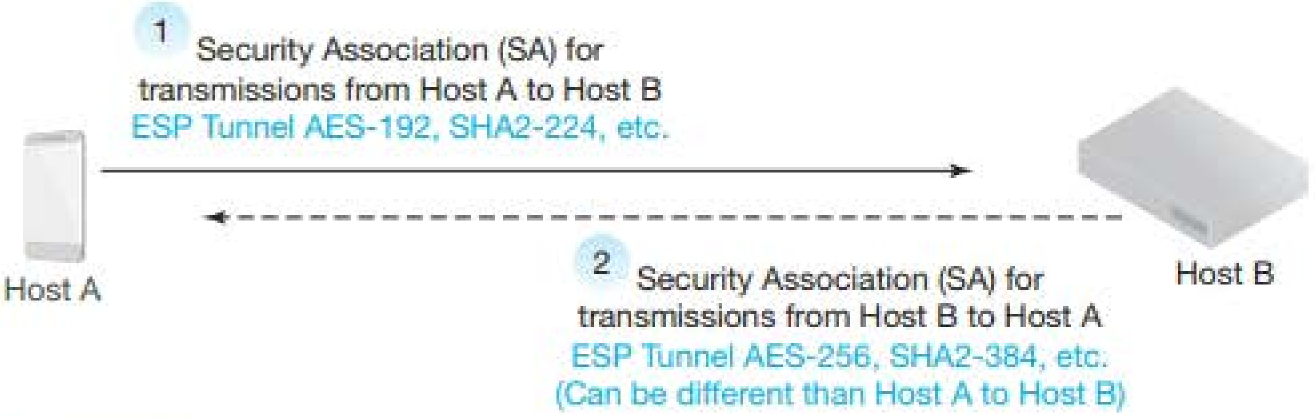
1 Security Association (SA) for
transmissions from Host A to Host B
ESP Tunnel AES-192, SHA2-224, etc.
Host A
Host B
2 Security Association (SA) for
transmissions from Host B to Host A
ESP Tunnel AES-256, SHA2-384, etc.
(Can be different than Host A to Host B)

FIGURE 9-23 Security Associations (SAs) in IPsec

- For transmission from Host B to Host A, the SA specifies that Host B will use ESP in tunnel mode. For confidentiality, Host B will encrypt with AES-256. For authentication, Host B will use SHA2-384.

Security Associations Can Be Asymmetric Note that the two security associations in Figure 9-23 are **asymmetric** (different in the two directions). The SA from Host A to Host B is weaker than the SA from Host B to Host A. (To see why this is true, 192 224, 256, and 384 are key lengths, and longer keys give stronger security even when the cryptographic method is the same.) The SAs in the two directions are often symmetrical (the same in both directions). However, they do not have to be. Sometimes, conditions call for asymmetrical security.

Security associations are often asymmetric, providing different security in the two directions.

Test Your Understanding

17. a) Distinguish between cryptographic methods and options. b) What is an SA? c) In Figure 9-23, what elements are standardized in the SAs?

Creating Security Associations

In IPsec, SSL/TLS, and other cryptographic systems, there are nearly always two stages. Figure 9-24 shows that the first is an initial handshaking (negotiating) stage. This is a very short stage in which the two parties do three things:

- Negotiate the security methods (and options) they will use in ongoing communication. This is the negotiation of the security associations the two parties will use.
- Authenticate each other.
- Securely exchange the keys they will use for ongoing communication.

Figure 9-24 shows that in IPsec this initial handshaking is governed by the **Internet Key Exchange (IKE)** protocol. IKE is an extremely complex protocol. Fortunately for you, it is well beyond what an introductory class can cover. With IKE's work done, the security associations are established, and ongoing communication begins. This stage, which accounts for nearly the entire communication session, implements the IPsec SAs negotiated by IKE.

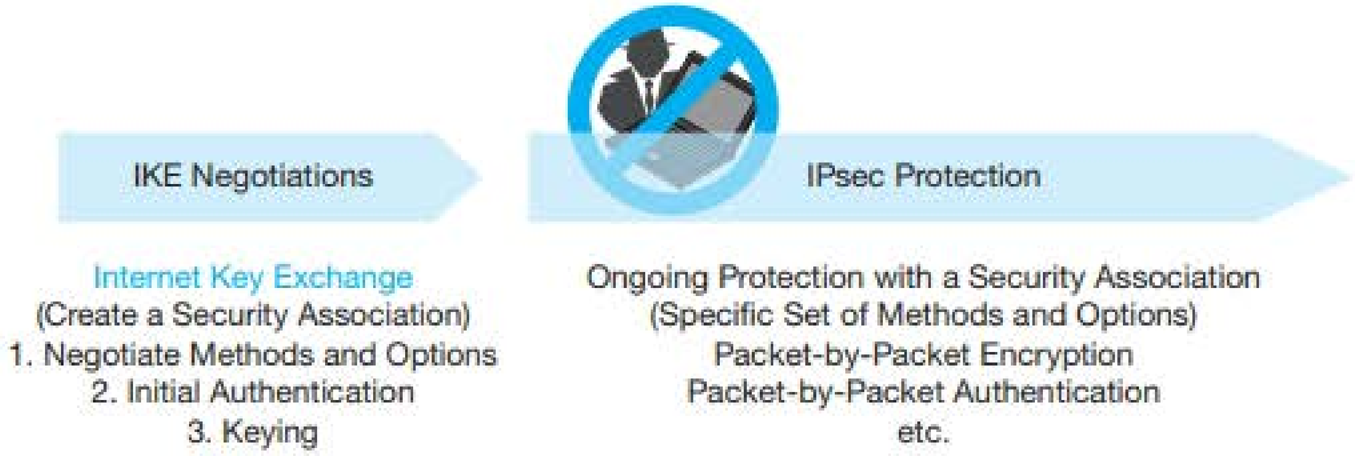
IKE Negotiations
IPsec Protection
Internet Key Exchange
(Create a Security Association)
1. Negotiate Methods and Options
2. Initial Authentication
3. Keying
Ongoing Protection with a Security Association
(Specific Set of Methods and Options)
Packet-by-Packet Encryption
Packet-by-Packet Authentication
etc.

FIGURE 9-24 Stages in IPsec Communication

Weak Methods and Options in Security Associations We have seen that IPsec offers multiple encryption methods and options. Some only provide weak security and have been cracked in practice, sometimes easily. For example, a minimum key length for encryption today is 128 bits, but DES only offers a 56-bit key. (In 1977, when DES was created, it was strong.) In authentication, both MD-5 and SHA-1 are weak and crackable today.

Companies must establish policies for not using weak algorithms in SAs. In many cases, these policies can be enforced in the technology. For example, companies that employ Microsoft servers can use Microsoft Group Policy Objects (GPOs) that enforce policies on different hosts, such as general client hosts, client hosts in highly risky operations, and so forth. These client hosts may run Windows or a Macintosh operating system. Assigning a host to a predefined group will require that host to respect the company's relevant policies for security and other matters.

Test Your Understanding

18. a) What are the two stages in IPsec protection? b) What standard is used in the first stage? c) In which stage is the SA negotiated? d) In which stage is the SA used to provide protection? e) Can SAs be different in the two directions? f) Why is it important to have and enforce policies for what cryptographic methods and options may be used in an organization?

SSL/TLS VPNs

Although IPsec is an enormously powerful tool for creating highly secure VPNs, IPsec is expensive to implement. For many purposes, companies implement VPNs using SSL/TLS, which we saw briefly in Chapter 4. Figure 9-25 compares IPsec with SSL/TLS. It shows that IPsec is a general approach to security protection, whereas SSL/TLS can only be used in some circumstances. One of these circumstances, of course, is interactions between browsers and webservers, which is very common.[18]

[18] IPsec is transparent, so a browser has no way of knowing if IPsec is being used to protect browser-server communication. This causes companies to implement SSL/TLS for many applications that can use it and require security even when IPsec is almost certainly being used.

Characteristic of VPN Technology	IPsec	SSL/TLS
Standards Organization	IETF	IETF (created by Netscape as SSL, renamed TLS by the IETF)
Layer	Layer 3	Layer 4
Built into Browsers, Webservers, and Mail Servers. So Protects These Applications at Little or No Cost.	No	Yes
Can protect any application	Yes (also protects transport-layer header and some of the IP header)	No (only SSL/TLS-aware applications such as web and e-mail)
Type of VPNs Supported in the Standard	Host-to-Host Remote Site Access Site-to-Site	Host-to-Host
Strength of Security	Excellent	Good

FIGURE 9-25 IPsec versus SSL/TLS VPNs

Test Your Understanding

19. a) List the strengths of IPsec compared to SSL/TLS. b) What is the attraction of SSL/TLS compared to IPsec?

END-OF-CHAPTER QUESTIONS

Thought Questions

9-1. Both DNS servers and DHCP servers send your client PC an IP address. What is different about these two addresses?

9-2. Assume that an average SNMP response message is 100 bytes long. Assume that a manager sends 4,000 SNMP Get commands each second. a) What percentage of a 1 Gbps LAN link's capacity would the resulting response traffic represent? b) What percentage of a 10 Mbps WAN link would the response messages represent? c) What are the management implications of your answers?

9-3. A firm is assigned the network part 128.171. It selects an 8-bit subnet part. a) Write the bits for the four octets of the IP address of the first host on the first subnet. b) Convert this answer to dotted decimal notation. (If you have forgotten how to do this, it was covered in Chapter 1.) c) Write the bits for the second host on the third subnet. (In binary, 2 is 10, and 3 is 11.) d) Convert this into dotted decimal notation. e) Write the bits for the last host on the third subnet. f) Convert this answer into dotted decimal notation. Can you tell the subnet a host is on just by looking at the dotted decimal notation representation?

9-4. A firm is assigned the network part 128.171. It selects a 10-bit subnet part. a) Draw the bits for the four octets of the IP address of the first host on the first subnet. b) Convert this answer into dotted decimal notation. (*Hint:* Use Windows Calculator.) c) Draw the bits for the second host on the third subnet. (In binary, 2 is 10, and 3 is 11.) d) Convert this into dotted decimal notation. (*Hint:* Use Windows Calculator.) e) Draw the bits for the last host on the third subnet. f) Convert this answer into dotted decimal notation. Can you tell the subnet a host is on just by looking at the dotted decimal notation representation?

9-5. a) What are the three ranges of private IP addresses? b) If a firm chooses 10.x.x.x for its internal IP addresses, how many hosts can it have internally? c) Repeat for 192.168.x.x. d) Repeat for 172.16.x.x through 172.31.x.x.

9-6. Pick one category in each category in Figure 9-12. Say how it would be used in network management.

9-7. Redo Figure 9-20 for transport mode.

9-8. After you get a second-level domain name, what more must you do to have a working website for your company?

Troubleshooting Question

9-9. Your computer sends a DNS request message to your local DNS server. After an unusually long time, your computer receives a DNS response message that the host name in your request message does not exist. This is a host you use every day. a) List problems that may have happened. (Draw the picture.) b) Which is the most likely to have cause the long delay and failure to find your host's IP address? c) How would you test it?

Hands-On Project

9-10. After Sal Aurigemma received his PhD from the University of Hawai`i, he became a professor at the University of Tulsa. There, he introduced the school to Aloha Friday, when people come to work in their colorful Aloha shirts. He got the idea of creating Aloha shirts with Tulsa's school colors and an emblem of the university on the shirt pocket. Suppose that he wants to create a company to sell school-specific Aloha shirts to other universities. He will need a company name and a second-level domain name. Go to an Internet domain name registrar. Thoughtfully come up with three appropriate and available domain names. Explain why each is good. Select one and explain why it is best.

Perspective Questions

9-11. What was the most surprising thing you learned in this chapter?

9-12. What was the most difficult thing for you in the chapter? Why was it difficult?

Chapter **9a**

Cisco's IOS Command Line Interface (CLI)

COMMAND LINE INTERFACES (CLIs)

When dumb terminals ruled the desktop (roughly when dinosaurs roamed the earth), they presented their users with command line interfaces. Figure 9a-1 shows a brief fragment from a CLI interaction. It shows that in a **command line interface (CLI)** the system gives a prompt and the user types a one-line command.

In a command line interface (CLI), the system gives a prompt and the user types a one-line command.

This example shows a small part of Cisco's CLI for **IOS**—Cisco's operating system for switches, routers, firewalls, and other devices. Configuration and management work on Cisco devices is still done primarily through this command line interface. A device administrator Telnets into the device or plugs a PC into the router. In the latter case, software then turns the administrator's expensive PC into a cheap dumb terminal.

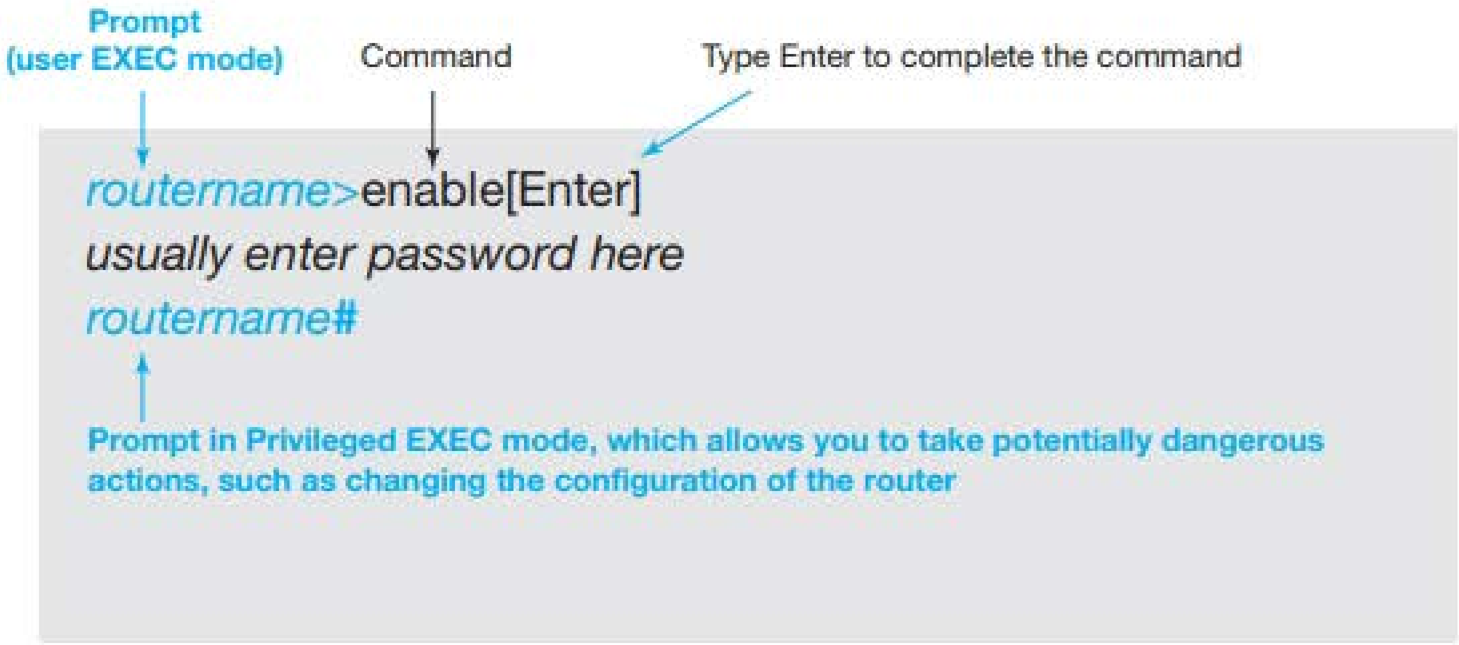

```
routername>enable[Enter]
usually enter password here
routername#
```

Prompt (user EXEC mode)
Command
Type Enter to complete the command
Prompt in Privileged EXEC mode, which allows you to take potentially dangerous actions, such as changing the configuration of the router

FIGURE 9a-1 Cisco IOS Command Line Interface (CLI)

Test Your Understanding

1. a) what is a CLI? Just spell it out. b) What is the defining characteristic of a command line interface? c) What is Cisco's operating system for routers, switches, and other devices? d) How is configuration and management work done primarily in this operating system?

CLI Essentials

The first line of the interaction in Figure 9a-1 shows the three elements in Command Line Interface (CLI) commands.

- First, there is the **prompt**. This is shown on the screen by the operating system. In the first row, the prompt is *routername*>, where router name is the name of the router being configured. The prompt ends with a right angle bracket (>). This tells the user that he or she is in *user EXEC mode*, which can execute only some IOS commands.
- Second, there is the command the administrator types. The command on the first line is a single text string, "enable." This tells IOS that the administrator wishes to enter *privileged EXEC mode*, which allows an administrator to enter all IOS CLI commands.
- Third, the user hits the Enter key to complete the command. The most common mistake that new CLI users make is forgetting to hit enter. They sit and wonder why nothing is happening. After a couple of minutes, they realize that they forgot to hit Enter. Again.

Command Modes

After the user types the first command, the system generally prompts the user for a password. Entering the *enable* command and successfully entering the password puts the administrator in *privileged EXEC mode*. To indicate this, the prompt ending changes from > to # to indicate that the administrator is in privileged EXEC mode.

In the EXEC mode, the prompt ends with >.
In privileged EXEC mode, the prompt ends with #.

Command mode is a core idea in Cisco's IOS CLI. The user is always in one of several command modes. Each command mode allows the administrator to take a specific set of actions. Different modes offer different sets of actions.[1]

Each IOS mode allows the administrator to take a specific set of actions.

[1] You can think of them as avatars with different powers in a computer game.

Test Your Understanding

2. a) What are the three things that appear on a CLI line? b) What is a command mode? c) In what command mode does the administrator start upon connecting to a router? d) In what command mode must the administrator be to give all commands? e) How does the prompt end in the user Exec command mode? f) How does the prompt end in the privileged Exec command mode? g) Why does the administrator have to always keep in mind what command mode he or she is in? This will require you to draw an inference from the material in this section.

A More Complex Cisco IOS Interaction

Figure 9a-2 shows part of a (slightly) more complex interaction. In the figure, the client logs in, goes into privileged EXEC command mode, changes the name of the router, and configures an interface. Configuring an interface requires going into a different command mode, *Configuration Interface Command Mode*, abbreviated as *config-if*. Recall that a router will normally have several interfaces (plugs) that connect to its network, so you must specify a particular interface before you configure it.

In this session, the administrator issues a series of commands.[2]

Prompt	Command	Description
routername>	Enable	The current name of the router is *routername*. ">" indicates that the user is in the restrictive user EXEC Mode. User wishes to be in Privileged EXEC Mode. Will be prompted for a password.
routername#	Config	Now in Privileged EXEC Mode (# prompt). Command to go into Global Configuration Mode.
routername(config)#	hostname Bob	In Global Configuration Mode, prompt ends in (config)#. Hostname command changes the router's name to "Bob."
Bob(config)#	interface e1	Enter Config-If (interface) Mode to configure one of the router's interfaces—the second Ethernet interface. (Interface counting begins with 0, not 1.)
Bob (config-if)#	ip address 172.30.3.100 255.255.255.0	Now in Interface Configuration Mode, the prompt ends with (config-if)#. Command assigns to the second Ethernet interface the IP address 128.171.17.13, with the subnet mask 255.255.255.0.
Bob (config-if)#	End	Ends interface configuration, goes back to Global Configuration Mode.
Bob(config)#	. . .	More IOS commands.

FIGURE 9a-2 A More Complex Cisco IOS Interaction Sequence

[2] Commands are normally shown in lowercase. However, case is not important in CLI commands.

- The administrator goes into Privileged EXEC Mode with the enable command and the correct password (not shown).
- The admin gives the Config command to enter Global Configuration Mode.
- The administrator renames the router using the hostname command. He or she changes it to "Bob."
- Bob then gives the interface command, stating that he or she wishes to configure interface e1. The e indicates that this is an Ethernet interface. Numbering begins at 0, so e1 is the router's second Ethernet interface.
- Now in Config-if Command Mode, the administrator gives the interface an IPv4 address and a subnet mask.
- End returns you to the next-higher mode. In Config-if Mode, the administrator types the End command. This puts the administrator back into Global Configuration Mode.

Note that the administrator frequently moves between modes. From the Global Configuration Mode, the administrator can switch modes to configure the router as a DHCP server, add in an access control list, specify a certificate authority to use public key authentication, and do many other advanced tasks. Mastering Cisco's CLI is a major challenge. Beyond that, knowing how to configure a router to work in complex or unusual environments takes years to master.

Test Your Understanding

3. a) If the administrator is in the Privileged EXEC Mode, what must he or she do to be able to do configuration work? b) What command mode will the administrator be in to make configuration changes? c) How does the administrator get from the Privileged EXEC Mode to the Global Configuration Mode? d) Give the syntax for the command to change the router's name. e) While in the Global Configuration Mode, give the command the administrator must use to begin configuring the first serial interface. (This will take you a little beyond the text.) f) To what will this change the prompt? g) What is the command to set an IP address and subnet mask for this interface? h) Give the command to do this for a subnet mask of sixteen 1s followed by sixteen 0s. i) How does the administrator go from this mode back to the Global Configuration Mode?

ACTIVITY

9a-1. Write the CLI prompts and commands you would use as an administrator in a session to change the third serial interfaces' IP address. Do not include unnecessary commands.

This page intentionally left blank

Chapter **10**

Carrier Wide Area Networks (WANs)

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Contrast LANs and WANs in terms of technology, diversity, economics, speed, and need for optimization.
- Describe the three carrier WAN components and the two typical business uses for carrier WANs.
- Describe how the telephone system is organized, including its hierarchy of switches. (Most carrier WAN networks use the public switched telephone network for some or all of their communication.)
- Explain and compare the ADSL and cable modem residential Internet access services and how fiber to the home is changing the residential access market.
- Discuss trends in cellular data transmission speeds.
- Distinguish between access lines and leased lines. Select a leased line for a given application speed requirement. Explain how companies use leased lines in Internet access.
- Explain how networks of leased lines, carrier Ethernet, and MPLS can be used for site-to-site communication within a firm. Discuss the relative advantages and disadvantages of each.
- Explain the capabilities of WAN optimization devices.

Albert Einstein was reportedly asked how the telegraph worked. He said it was like a very long cat with its head in one city and its tail in another. When you pull on the tail in one city, it howls in the other city. Wireless transmission is exactly the same but without the cat.

LANs AND WANs (AND MANs)

One of the most fundamental distinctions in networking is the one between local area networks (LANs) and wide area networks (WANs). Figure 10-1 shows how these two types of networks differ. We will also see how they compare to intermediate-distance networks called metropolitan area networks (MANs).

LANs versus MANs and WANs

On and Off the Customer Premises Some authors base the difference between LANs and WANs on physical distance. For instance, some say that the dividing line between LANs and WANs is one mile or one kilometer. However, the real distinction appears to be that **local area networks (LANs)** exist within a company's site, whereas **wide area networks (WANs)** connect different sites within an organization or between organizations.

Local area networks (LANs) exist within a company's site, whereas wide area networks (WANs) connect different sites within an organization or between organizations.

For LANs, then, the company owns the property and can do anything it wants. It can choose any LAN technology it wishes, and it can implement it any way it wishes.

There is no such freedom for WANs. A company cannot legally lay wires between two of its sites. (Consider how your neighbors would feel if you started laying wires

Category	Local Area Network	Metropolitan Area Network	Wide Area Network
Abbreviation	LAN	MAN	WAN
Service Area	*On customer premises* (home, apartment, office, building, campus, etc.)	*Between sites* in a metropolitan area. (city and its suburbs) A Type of WAN	*Between sites* in a region, a country, or around the world.
Implementation	Self	Carrier	Carrier
Ability to Choose Technology	High	Low	Low
Who Manages the Network?	Self	Carrier	Carrier
Price	Highly related to cost	Based on pricing strategy. Highly unpredictable	Based on pricing strategy. Highly unpredictable
Cost per Bit Transmitted	Low	Medium	High
Typical Transmission Speed	1 Gbps and more	10 Mbps to 1 Gbps	1 to 100 Mbps
Diversity of Technologies	Low: 802.3 and 802.11	Medium	High

FIGURE 10-1 LANs versus WANs (and MANs)

across their yards.) The government gives certain companies, called **carriers**,[1] permissions (**rights of way**) to lay wires in public areas and offer service to customers. In return, carriers are subject to government regulation.

When you deal with carriers, you can only get the services they offer, and you must pay their prices. Although there may be multiple carriers in an area, the total number of service choices is likely to be quite limited.

On the positive side, you do not need to hire and maintain a large staff to deal with WANs because carriers handle nearly all of the details. In contrast, if you install a LAN, you also have to maintain it. As the old saying goes, anything you own ends up owning you.

Economics Another fundamental difference between LANs and WANs stems from economics. You know that if you place a long-distance call, it will cost more than a local call. An international call will cost even more. As distance increases, the price of transmission increases. The cost per bit transmitted therefore is higher in WANs than in LANs.

You know from basic economics that as unit price increases, fewer units are demanded. Or, in normal English, when the price of an item increases, you usually buy less of it. Consequently, companies tend to purchase lower-speed WAN links than LAN links. Typically, LANs bring 1 Gbps to each desktop. WAN speeds more typically vary from 1 Mbps to about 100 Mbps. MAN speeds fall between the two.

In addition, companies spend more time optimizing their expensive WAN traffic than their relatively inexpensive LAN traffic. For example, companies may be somewhat tolerant of looking at YouTube videos on LANs, but they usually clamp down on this type of information on their WAN links. They also tend to compress data before sending across a WAN so that it can be handled with a lower-capacity WAN link.

Another aspect of economics is pricing. For LANs, you have a good idea of what installing and using a wired or wireless LAN will cost you. In carrier WANs, however, the price of services is only somewhat related to costs. Carriers change their prices strategically, for example, to encourage users to switch from one service to another. Consequently, price changes for WANs are less predictable than they are for customer-owned LAN technology.

Technologies Another difference between LANs and WANs is that LAN technology has largely settled on two related families of standards—Ethernet (802.3) for wired LANs and Wi-Fi (802.11) for wireless LANs. As we saw in Chapter 6, 802.11 WLANs are primarily used today to extend corporate Ethernet wired LANs to mobile devices.

The technological situation is more complex in wide area networking. Multiple technologies are used, including leased line data networks, public switched data networks, and wireless networks. Within these categories are multiple options. Furthermore, WAN technologies are at different stages in their life cycles, with some increasing rapidly in use and others declining.

[1] Carriers were originally called common carriers. The name reflected the fact that these carriers were required by law to provide service to anyone or any organization requesting services. Regulation was originally instituted in the railroad industry because many companies that owned railroads also owned other companies and refused to provide services to competitors of these other companies.

Test Your Understanding

1. a) Distinguish between LANs and WANs. b) What are rights of way? c) What are carriers? d) Why do you have more flexibility with LAN service than with WAN service?
2. a) Why are typical WAN speeds slower than typical LAN speeds? Give a clear and complete argument. b) Why are future WAN prices difficult to predict? c) Compare the diversity of technologies in LANs and WANs.

Other Aspects of WANs

Metropolitan Area Networks (MANs) All WANs connect sites between customer premises and cost more per bit transmitted than LANs. However, WANs differ considerably in the distances they span. Some are international and others span single nations. At the small end, some WANs are **metropolitan area networks (MANs)**, which connect sites in a city and its suburbs.

Although MANs are WANs, their relatively short distance span means that the cost per bit transmitted is lower than it is in national and international WANs. Consequently, typical transmission speeds are faster. If you have a smartphone or tablet with 3G or 4G cellular access, then you already use a MAN. Cellular networks almost always span a single MAN or even a single city. However, we will see that wired MANs are important for corporations because site-to-site traffic is large and is more efficiently transmitted over wires.

Single Networks versus Internets Some people think that LANs are single networks and that WANs are internets. However, as Figure 10-2 shows, that is not the case. Small LANs usually will be single networks, but a larger LAN, such as one on a university campus, is likely to be a local internet.

For WANs, there also can be single networks or internets. Of course, the global Internet is a WAN, and we will see that many companies use it extensively for data transmission among their premises. We also will see that companies use wide area single switched networks. These are large networks, but they are still switched single networks.

Test Your Understanding

3. a) Why do MANs have higher typical speeds than broader-scope WANs? b) Are LANs single networks or internets? c) Are WANs single networks or internets? d) Is the Internet a WAN?

Technology	LAN	WAN
Can be a single switched or wireless network?	Yes	Yes
Can be an internet?	Yes	Yes

FIGURE 10-2 Single Networks versus Internets

Carrier WAN Components and Business Uses

Figure 10-3 shows that there are three basic components to carrier wide area networks:

- First comes the customer premises with the **customer premises equipment (CPE)** needed to connect to the WAN. With mobile devices, your customer premises is wherever you are, and your mobile device is your customer premises equipment. For connecting corporate sites to wired access lines, the customer premises equipment is likely to be a border router.
- **Access links** connect the customer premises to the network core of the WAN. We will focus on wired access links because they are so prevalent. Later in the chapter, we will look at wireless access links.
- The **network core** connects access links to other access links. Again, we show it as a cloud because customers do not have to understand how it works in detail. The carrier takes care of the network core. Of course, as an IT professional, you have to understand what happens inside the cloud, and we will spend time looking at network core technologies.

The Internet connects everyone to everyone else. In contrast, **carrier WANs** primarily see two business uses. As Figure 10-3 shows, companies use carrier WANs to link their sites to the Internet and to connect their own sites together. Carrier WANs are not frequently used to connect multiple companies together because all must be customers of the same carrier WAN. When multiple companies connect with a carrier WAN, it is generally because they need more security than the Internet provides.

Test Your Understanding

4. a) List the three basic components of wide area networks. b) Are access links wired or wireless? c) What is CPE? d) What are the two common business uses for carrier WANs? e) Distinguish between the Internet and carrier WANs. f) Why are carrier WANs not often used to link multiple firms together?

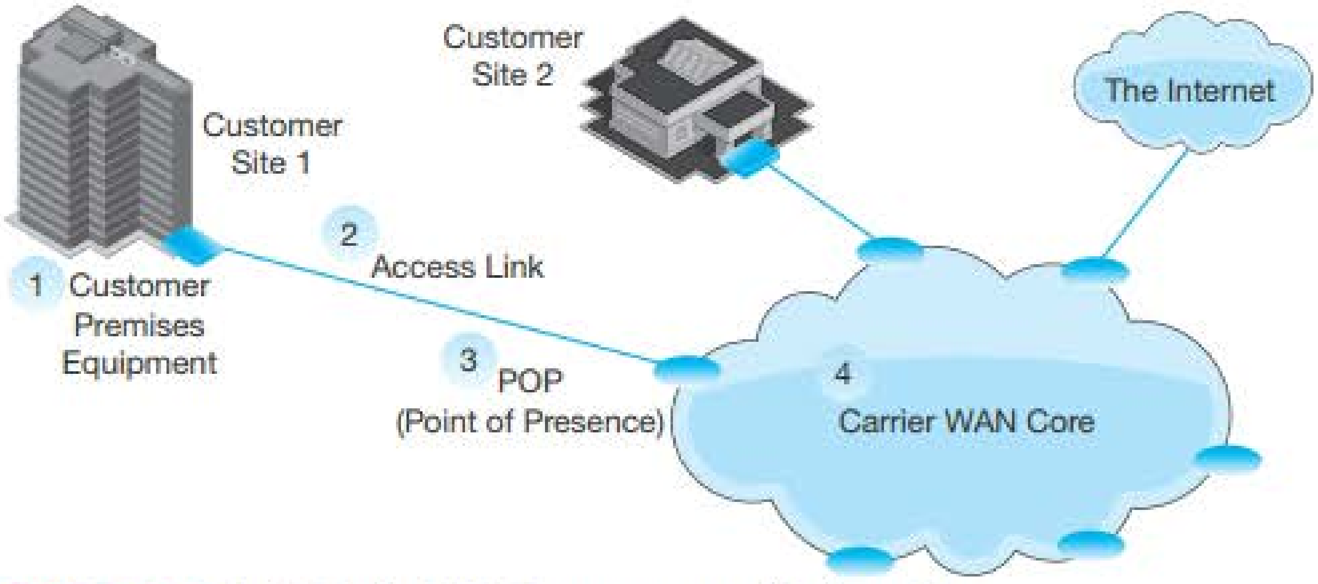

FIGURE 10-3 Basic Carrier WAN Components and Business Uses

The Telephone System

The worldwide telephone system was created by voice. However, telephone carriers now provide data service to residential and business customers. In addition, other WAN carrier providers typically find it attractive to lease their transmission lines from telephone companies. This allows WAN providers to focus on data switching.

Figure 10-4 shows the **Public Switched Telephone Network (PSTN)**, which is the official name of the telephone system. Per our discussion earlier, there is a central core, and there are access lines. The access portion of the PSTN is the **local loop**. It extends from the final telephone company switch to the customer premises.

The **PSTN Core** is a modified hierarchy of switches. **End office switches** connect the PSTN to the customer. These are usually **Class 5 switches**—the lowest in the hierarchy. For perspective, there are about 100 Class 5 end office switches in the state of Hawaii. There are fewer switches at each subsequent level. For example, Hawaii has a single Class 3 switch.

The PSTN is a modified hierarchy in the sense that, unlike Ethernet, the PSTN includes bypass trunk lines between switches that are at the same level if there is an unusually large volume of traffic between those switches. It is more efficient for such pairs of switches to communicate directly rather than involving a higher-level switch.

Test Your Understanding

5. a) Why is the PSTN important in WAN data transmission? b) What is the local loop? c) What class of switches are most end office switches? d) What is the structure of the PSTN core?

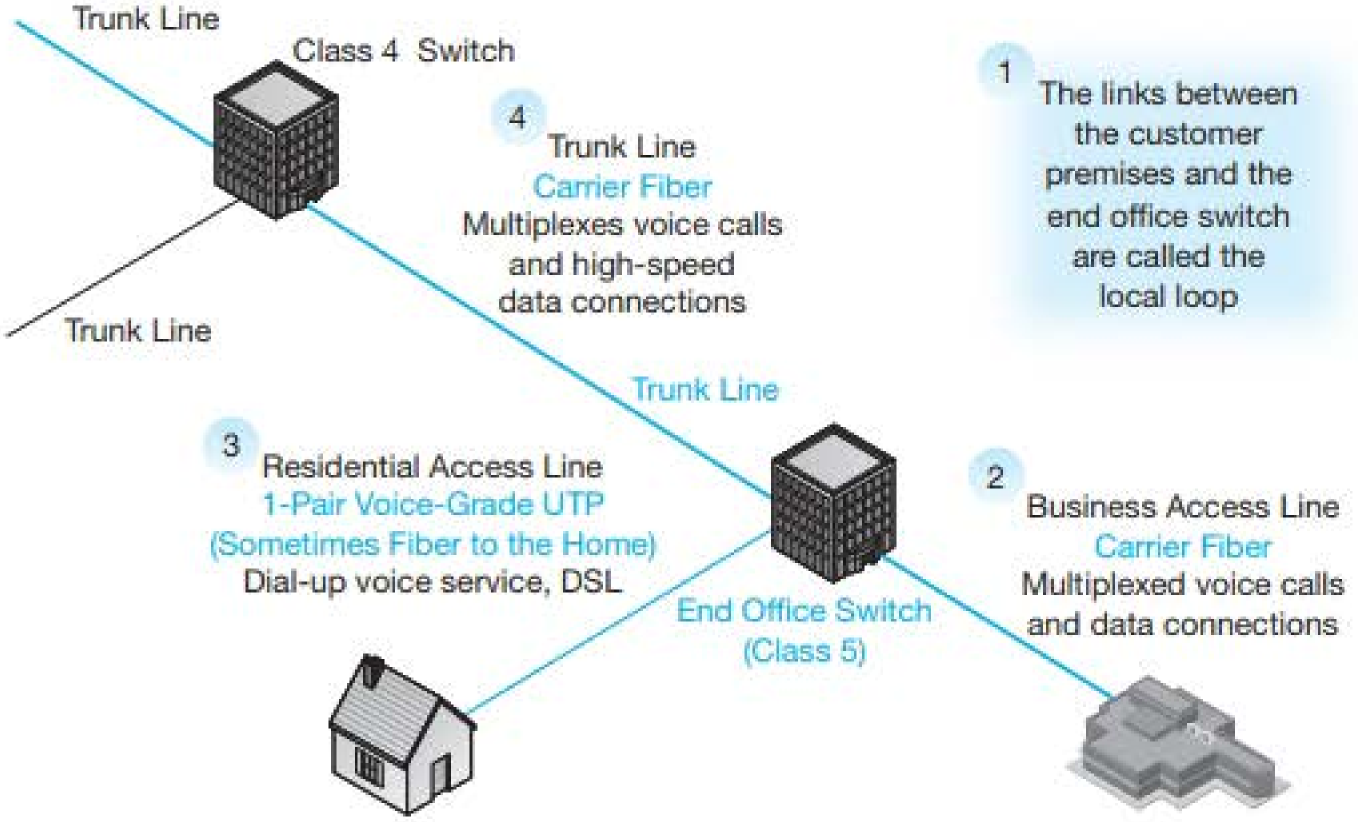
Trunk Line
Class 4 Switch
4
Trunk Line
Carrier Fiber
Multiplexes voice calls and high-speed data connections
1
The links between the customer premises and the end office switch are called the local loop
Trunk Line
Trunk Line
3
Residential Access Line
1-Pair Voice-Grade UTP
(Sometimes Fiber to the Home)
Dial-up voice service, DSL
2
Business Access Line
Carrier Fiber
Multiplexed voice calls and data connections
End Office Switch
(Class 5)

FIGURE 10-4 The Public Switched Telephone Network (PSTN)

RESIDENTIAL WIRED INTERNET ACCESS

We begin our discussion of WAN technology with residential Internet access. This will permit us to start with something familiar to most readers and give us a base of knowledge for looking at corporate WAN technologies.

Residential Asymmetric Digital Subscriber Line (ADSL) Service

Some readers are directly familiar with residential ADSL services. Figure 10-5 shows that **asymmetric digital subscriber line (ADSL)** services provide simultaneous voice and data to residential customers. Data transmission speed is asymmetric, with faster download speed than upload speed. This is reasonable. Website downloading often requires a great deal of downstream speed. So does video streaming. In contrast, few residential applications require full two-way high-speed service.

Digital Subscriber Lines Telephone companies have traditionally served residential customers with **one-pair voice-grade (1PVG) UTP** in the local loop. This single unshielded pair was created for voice, not data. It is only twisted about once a foot. However, advances in signaling algorithms have allowed telephone companies to transmit data at high speeds over these lines—while continuing to deliver voice at the same time.

The line between the end office switch and the customer is called the subscriber line. When the telephone company transmits digital signals over it, it is called a **digital subscriber line (DSL)**. These are also called DSL lines, despite the fact that expanding the acronym gives "digital subscriber line lines."

Sending data over 1-pair voice grade UTP is important because subscriber lines using this technology already run to every home and business. They have been used since the 1880s to deliver voice telephone service. There is no need to run new subscriber lines to homes in order to provide data transmission. In contrast, the business-focused leased lines that we will see later require carriers to run new transmission lines to each organization. This is extremely expensive.

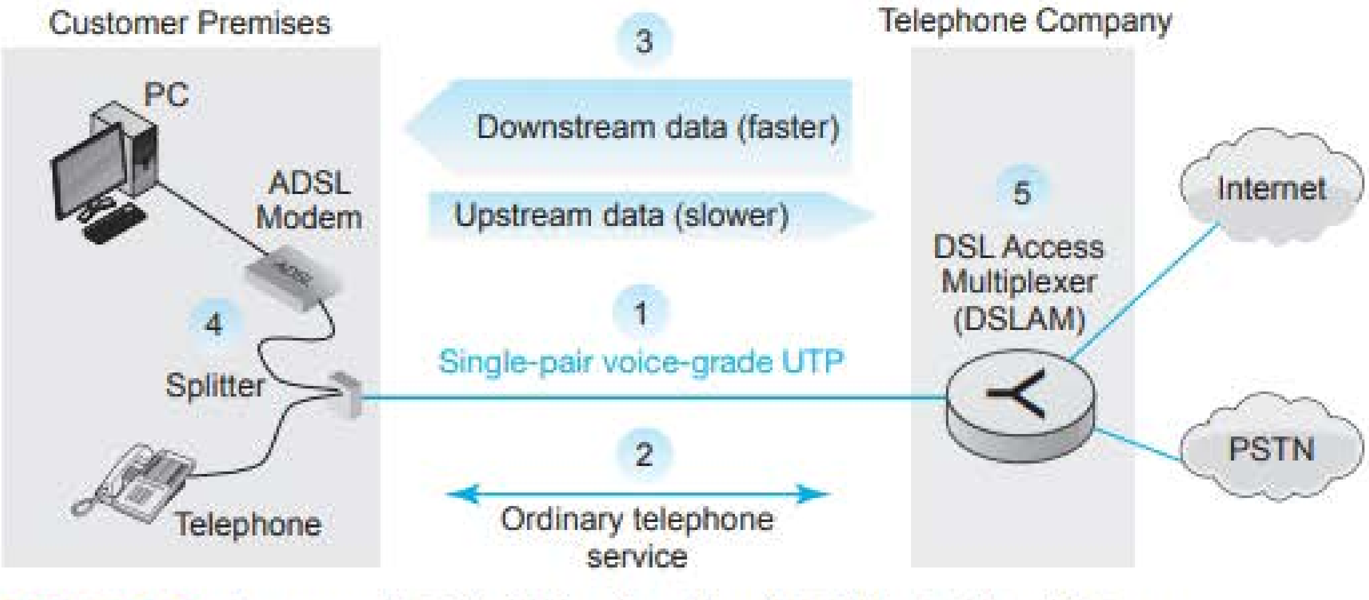

FIGURE 10-5 Asymmetric Digital Subscriber Line (ADSL) for Residential Access

Residential Customer Equipment and Service For ADSL service, a residential customer installs **ADSL modems**, although it is best to install splitters in each outlet. These **splitters** have two jacks—one for voice and one for data. Splitters separate voice and data signals, preventing possible transmission impairment.

How fast are transmission speeds in ADSL? The answer changes by the minute. In mid-2004, the first author was getting downstream speeds of just under 10 Mbps and upstream speeds a little over 2 Mbps. This is fast enough even for a high-definition movie download. ADSL vendors hope to raise downstream speeds to 100 Mbps or more in the near future. This will permit several high-definition telephone streams into the house. Faster upstream speeds will make online backup for hard disks reasonably painless.

Carrier End Office Equipment To provide ADSL, the carrier has to install a new piece of equipment at the end office switch. This is a **DSL access multiplexer (DSLAM)**. When the customer transmits, the DSLAM directs voice signals to the public switched telephone. However, when data signals arrive, the DSLAM sends it on to an ISP. The DSLAM multiplexes incoming voice and data signals onto the subscriber line.

Fiber to the Home (FTTH) Although DSL speeds today are quite fast, subscribers want to bring high-definition video into their homes, and they want multiple channels at a time. Although 1-pair voice-grade UTP is already installed, its limits are being reached. For speeds beyond about 100 Mbps, carriers are beginning to bring **fiber to the home (FTTH)**—running optical fiber from the end office switch to residential households.

Running new fiber to each household is expensive, so implementation will take time. However, by converting entire neighborhoods to FTTH at one time, carriers have been able to lower their per-house installation costs and offer more reasonable prices.

Test Your Understanding

6. a) Does residential DSL offer simultaneous voice and data service? b) Why is asymmetric speed acceptable in residential ADSL service? c) What is beneficial about transmitting data over 1-pair voice-grade UTP? d) What equipment does the customer need in his or her home? e) What is the purpose of the DSLAM? f) Why is FTTH attractive? g) How are carriers attempting to reduce the cost of installing FTTH?

Cable Modem Service

Telephone Service and Cable TV In the 1950s, **cable television** companies sprang up in the United States and several other countries, bringing television into the home. Initially, cable only brought over-the-air TV to rural areas. Later, it began to penetrate urban areas by offering far more channels than urban subscribers could receive over the air. In the 1970s, many books and articles forecast a "wired nation" in which two-way cable and the advent of 40-channel cable systems would soon turn cable into an information superhighway. (After all, it would be impossible to fill 40 channels just

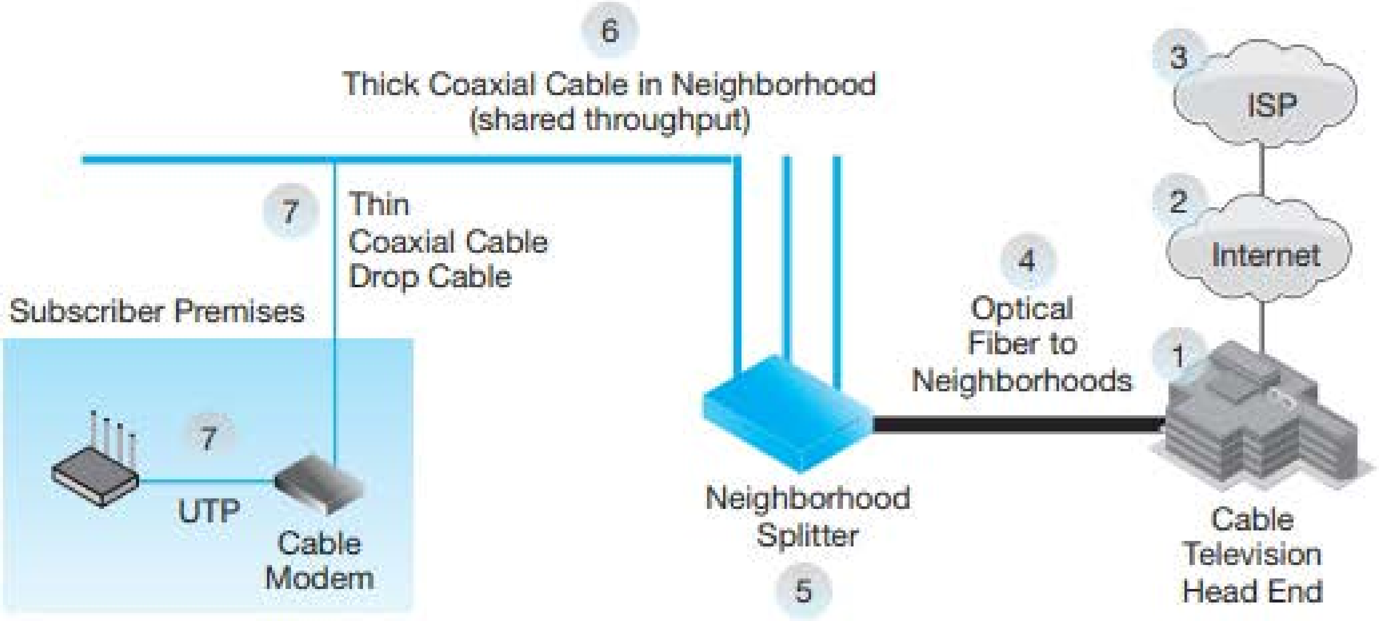
6
Thick Coaxial Cable in Neighborhood (shared throughput)
7
Thin Coaxial Cable Drop Cable
Subscriber Premises
7
UTP
Cable Modem
Neighborhood Splitter
5
4
Optical Fiber to Neighborhoods
3
ISP
2
Internet
1
Cable Television Head End

FIGURE 10-6 Cable Modem Service

with television, wouldn't it?) However, available services did not justify the heavy investment to make cable a two-way service until many years later.[2]

Figure 10-6 shows how cable television operates. The cable television operator has a central distribution point, called a **head end**. From the head end, signals travel out to neighborhoods via optical fiber.

From neighborhood splitters, signals travel through **coaxial cable**. The transmission of an electrical signal always requires *two* conductors. In UTP, the two conductors are the two wires in a pair. Figure 10-7 shows that in coaxial cable, the first conductor is a wire running through the center of a coaxial cable. The second conductor is a mesh wire tube running along the cable. The two conductors have the same axis, so the technology is called **coaxial cable**. Before the advent of high-definition HDMI cables, you typically connected your VCR to your television with coaxial cable.

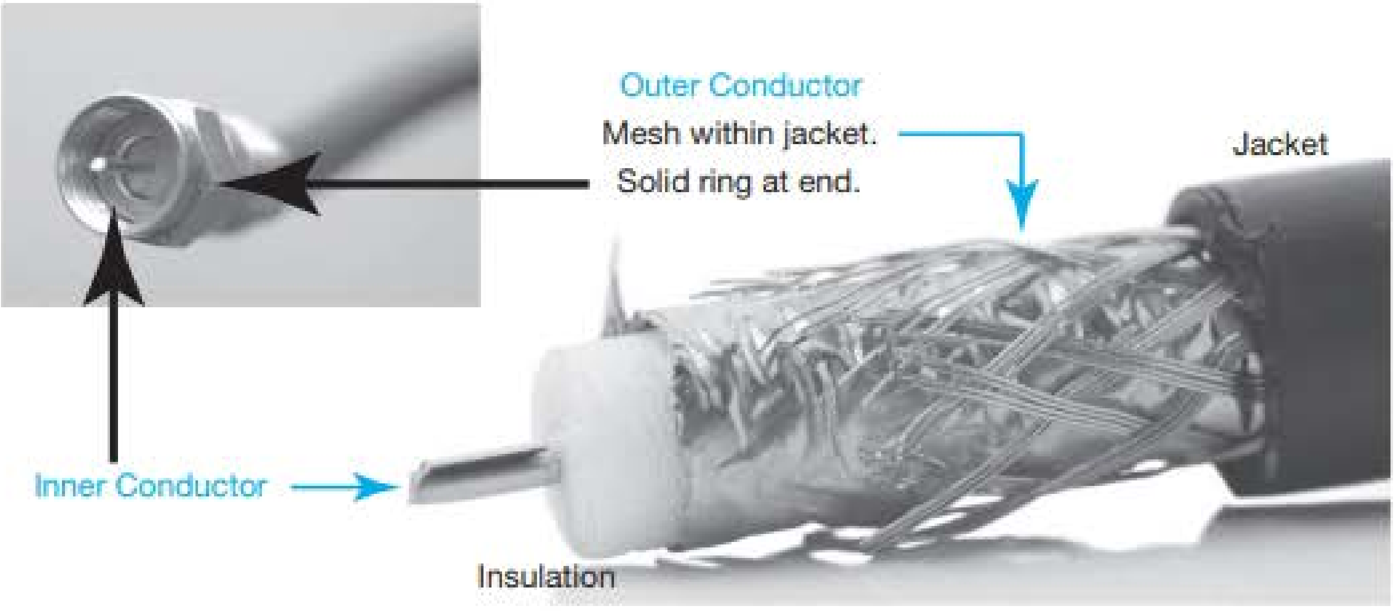

FIGURE 10-7 Coaxial Cable

[2] This was proven in the dissertation of a Stanford PhD student. The student received a contract from the White House to do the study. Unfortunately, when the study was finished, Richard Nixon was being impeached, and the Executive Office of the President of the United States refused to release the study—despite the fact that the results of the study were already widely known. The study was released a year later, and the student was able to get his doctorate.

The cable television company runs signals through the neighborhood using *thick coaxial cable* that looks like a garden hose. The access line to individual homes is a *thin coaxial cable* **drop cable**. The resident connects the drop cable to his or her television.

Cable Modem Service Cable television companies eventually moved beyond one-way television service to two-way broadband (fast) data service. For television, the repeaters that boost signals periodically along the cable run only had to boost television signals traveling downstream. Data transmission required cable companies to install **two-way amplifiers**, which could carry data in both directions. Although this was expensive, it allowed cable companies to compete in the burgeoning market for broadband service. As in the case of ADSL, cable television service was asymmetric, offering faster downstream speeds than upstream speeds.

Instead of having a DSL modem, the subscriber has a **cable modem**. In general, this cable data service is called **cable modem service**. The coaxial cable drop line goes into the cable modem. The cable modem has a USB port and an Ethernet RJ-45 connector. The subscriber plugs a computer or access router into one of the two ports.

At the cable television head end, the cable television company connects to an Internet service provider. This allows subscribers to connect to hosts on the Internet.

Test Your Understanding

7. a) What transmission media do cable television companies use? b) Why is coaxial cable called "coaxial?" c) Distinguish between the coaxial trunk cable and drop cable. d) What types of amplifiers are needed for cable data service? e) What device do customers need for cable modem service?

ADSL versus Cable Modem Service

Telephone carriers and cable television companies constantly argue about the relative advantages of their two technologies. In reality, however, things boil down to speed and cost. The situation is changing rapidly. Both are increasing speeds frequently, and both are moving to FTTH. At most points in time, ADSL has been a little cheaper and a little slower. It will be interesting to see how competition drives them to improve in the future.

Test Your Understanding

8. a) What are the important things to consider when deciding between ADSL and cable modem service for your residence? b) In the past, how has ADSL compared to cable modem service? c) Which of these two services is moving toward FTTH?

CELLULAR DATA SERVICE

ADSL and cable modem service provide wired access to the Internet by linking users to their ISPs. Cellular telephony now connects users to their ISPs while they are away from home, in the office, or in hotspots. Businesses use cellular telephone service the same way.

Cellular Service

Nearly everybody today is familiar with cellular telephony. In most industrialized countries, well over half of all households now have a cellular telephone. Many now have *only* a cellular telephone and no traditional *wireline* public switched telephone network phone.

Cells and Cellsites Figure 10-8 shows that cellular telephony divides a metropolitan service area into smaller geographical areas called **cells**. A city the size of Honolulu will have a few hundred cells.

The user has a cellular telephone (also called a mobile phone, **mobile**, or cellphone). Near the middle of each cell is a **cellsite**, which contains a transceiver (transmitter/receiver) to receive mobile phone signals and to send signals out to the mobiles. The cellsite also supervises each mobile phone's operation (setting its power level, initiating calls, terminating calls, and so forth).

Mobile Telephone Switching Office (MTSO) All of the cellsites in a cellular system connect to a **mobile telephone switching office (MTSO)**, which connects cellular customers to one another and to wired telephone users.

The MTSO also controls what happens at each of the cellsites. It determines what to do when people move from one cell to another, including deciding which cellsite should handle the transmission when the caller wishes to place a call.[3]

Cellsite Figure 10-9 shows a typical small cellsite on top of a residential building. The three large "paddles" are cellular antennas.

Handoffs If a subscriber moves from one cell to another within a city, the MTSO will implement a **handoff** from one cellsite to another. For instance, Figure 10-8 shows

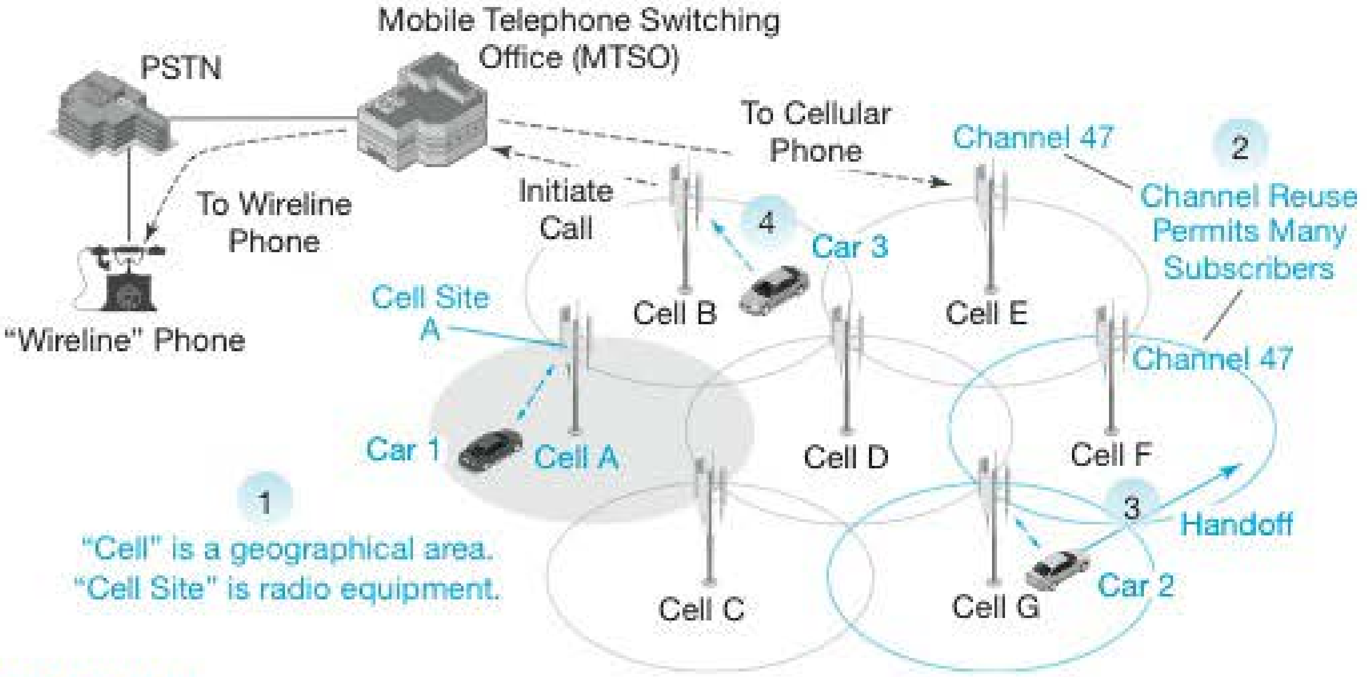

FIGURE 10-8 Cellular Telephone Service

[3] Several cellsites may hear the initial request at different loudness levels; if so, the MTSO selects a service cellsite based on signal strength, not physical distance.

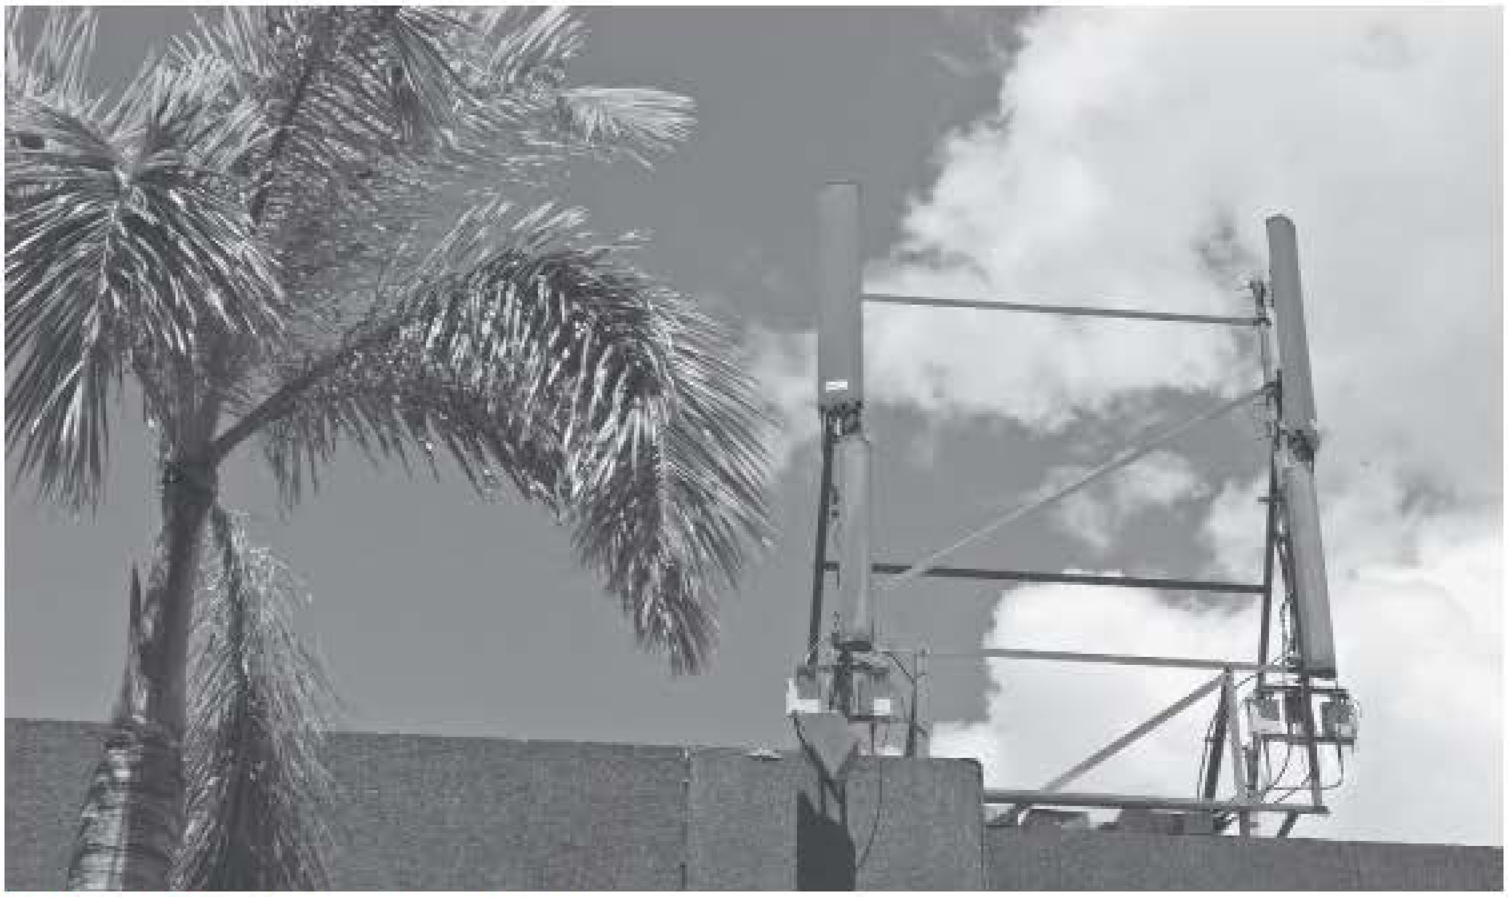

FIGURE 10-9 Cellsite with Paddle Antennas

a handoff from Cell O to Cell P. The mobile phone will change its sending and receiving channels during the handoff, but this occurs too rapidly for users to notice.[4]

Test Your Understanding

9. a) In cellular technology, what is a cell? b) What is a cellsite? c) What are the two functions of the MTSO? d) In a cellular system, distinguish between handoffs and roaming.

Why Cells?

Why not use just one central transmitter/receiver in the middle of a metropolitan area? Early pre-cellular radio telephone systems did use a single antenna, and this was much cheaper than using multiple cellsites.

The answer is **channel reuse**. The number of channels permitted by regulators is limited, and subscriber demand is heavy. Cellular telephony uses each channel multiple times, in different cells in the network. This multiplies the effective channel capacity, allowing more subscribers to be served with the limited number of channels available.[5]

Test Your Understanding

10. a) Why does cellular telephony use cells? b) What is the benefit of channel reuse?

[4] In contrast, if a subscriber leaves a metropolitan cellular system and goes to another city or country, this is roaming. To confuse matters, many carriers only call going to another city roaming if the home carrier does not offer service there.

[5] In a sense, enterprise wireless LANs with many access points are like cellular technologies. They allow users to employ the limited number of frequencies available in WLANs many times within a building.

Cellular Data Speeds

One problem in evaluating the speeds of different cellular carriers is that throughput is always considerably lower than advertised speed and varies widely within a system. There are several reasons for this.

- There is extensive overhead in cellular transmission. The data transmission rate is always less than the bit transmission rate.
- If the user is riding in a car, throughput will fall.
- If more customers use a cellsite, the cellsite may have to decrease the transmission speed to each. In particular, speed will depend on time of day.
- If the user travels into an area with an overloaded cellsite, speed will be lower.
- At greater distances from a cellsite, speed falls, just as in Wi-Fi.
- Weakened signal strength caused by transmission through buildings will also reduce speed.

Test Your Understanding

11. What factors affect what throughput an individual user will receive?

CELLULAR GENERATIONS: 3G, 4G, AND 5G

Cellular telephony has been transformed repeatedly since its birth in the 1980s. It is common to describe this evolution as a series of generations, with each generation bringing more speed and other benefits. Cellular carriers report that we are now buying fourth-generation (4G) phones and are about to see fifth-generation (5G) cellular systems.

The idea of generations began with standards agencies led by the International Telecommunications Union (ITU), which creates and authorizes standards for cellular telephony. Unfortunately, cellular carriers have largely ignored these official generations. Carrier marketing departments, for example, have used the term 4G for services that are 100 times slower than ITU's 4G standards require. Carriers are now preparing to market 5G systems; this is interesting because the ITU has not created 5G standards yet.

Corporations deal with carriers rather than the ITU-T, so we will (reluctantly) discuss generations as carriers do. Figure 10-10 shows that carrier generations have come roughly every decade since the 1980s. It also shows their typical data speeds and the new applications that each generation made possible.

- The first generation appeared in the early 1980s. These systems were limited to voice communication. Just being able to talk while walking around was revolutionary.
- The 1990s brought 2G systems. These were still primarily telephones, but they were entirely digital services. They could carry data—but only at the incredibly slow speed of 10 kbps. (This is not a typographical error. This at least permitted text messaging and text-only e-mail.
- The new century introduced mobile phones that would be familiar to today's users. 3G systems brought Internet access with speeds that at least made some

Carrier Generation	Dates	Typical Data Speed	Application Now Possible	Remarks
1G	1980s	Voice-only	Voice telephony	
2G	1990s	10 kbps	Texting and Text-Only E-Mail	First generation of digital devices and digital transmission
3G	2000s	A few hundred kilobits per second	Web Surfing	
4G	2010s (current)	10 Mbps	Video Streaming	Uses IP for transmission.
5G	2020s	100 Mbps		Low-power mode for IoT devices Low-latency for responsiveness

FIGURE 10-10 Cellular Generations (Carrier Terminology)

sense—a few hundred kilobits per second. This was far from perfect, but even slowly loading webpages on a phone were welcome. From a technical standpoint, all 4G systems run over IP. This integrates cellular data service with mainstream networking.

- We are currently in the fourth generation of cellular phones. Throughputs of 10 Mbps are common. Surfing the Internet is now comfortable. Beyond that, we can now stream television, movies, and videos of kittens.
- We are about to see a new generation of service. Most obviously, 5G should boost speed by another order of magnitude. It will also have many other technical advances. For one thing, it will offer low-speed, low-power modes to bring the Internet of Things into the cellular world. In addition, 5G cellular should slash latency. This will make interactions with remote software far more natural, slashing the last major impediment to working with remote systems.

Test Your Understanding

12. a) What does "G" stand for in cellular telephony? b) Which generation first brought decent Web access? c) Which generation is now bringing speeds of about 10 Mbps? d) What speeds can we expect from 5G? e) Why will 5G also bring energy-efficient low speeds? f) When carriers use the terms 4G and 5G, do they use it consistently with the formal standards for 4G and 5G?

WIRED BUSINESS WANs

To communicate with customers and for access to remote employees, companies use the Internet. However, they still need to use carrier WANs to reach the Internet and to connect their sites to one another. Figure 10-11 illustrates this situation.

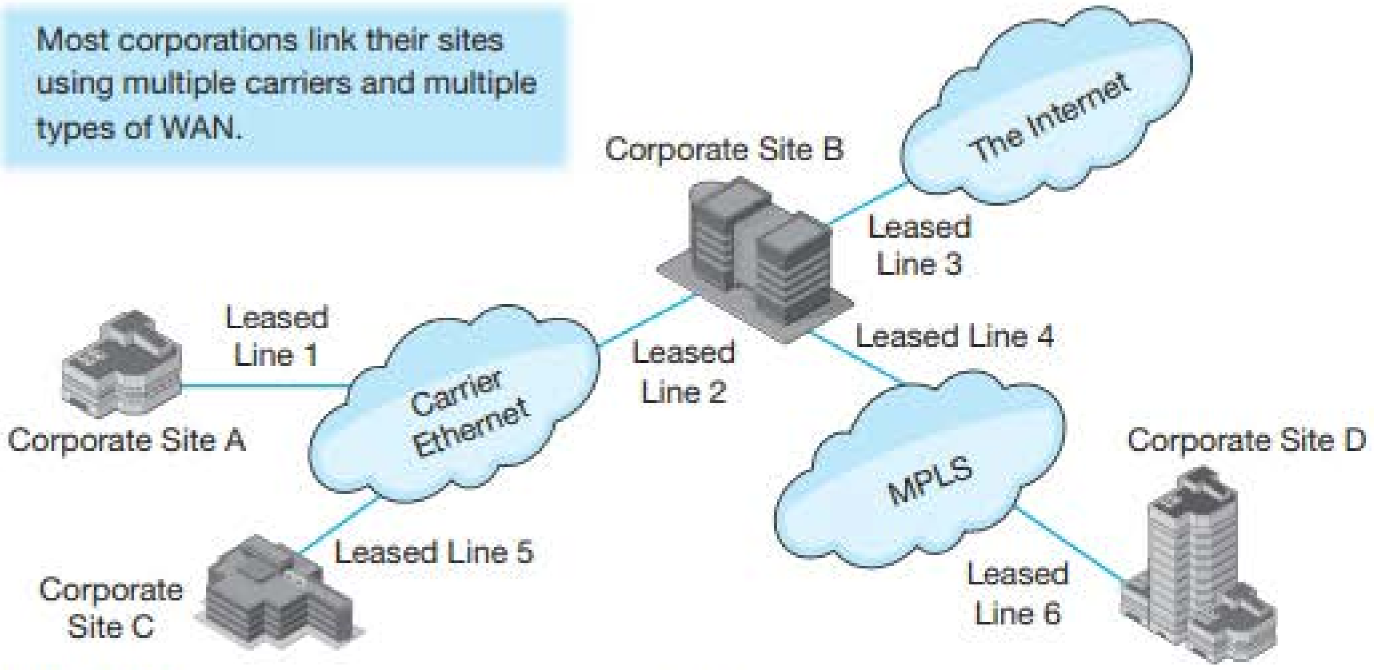

FIGURE 10-11 The Internet and Wired Carrier WANs for Business

Leased Lines

To connect to the Internet, Figure 10-11 shows that companies typically use leased lines from a carrier, most commonly the local telephone company. **Leased lines** are fast, point-to-point, always-on connections. As the name suggests, if a company wishes to use a leased line, it must sign a lease for a specified duration. Specifying the wrong speed when a leased line is ordered creates a persistent problem.

Figure 10-12 shows that a leased line is really a complex transmission path between the two points it connects. This path passes through customer access lines at the two ends and trunk lines between carrier switches along the path. To the user, however, the access line seems to be a simple data pipe all its own.

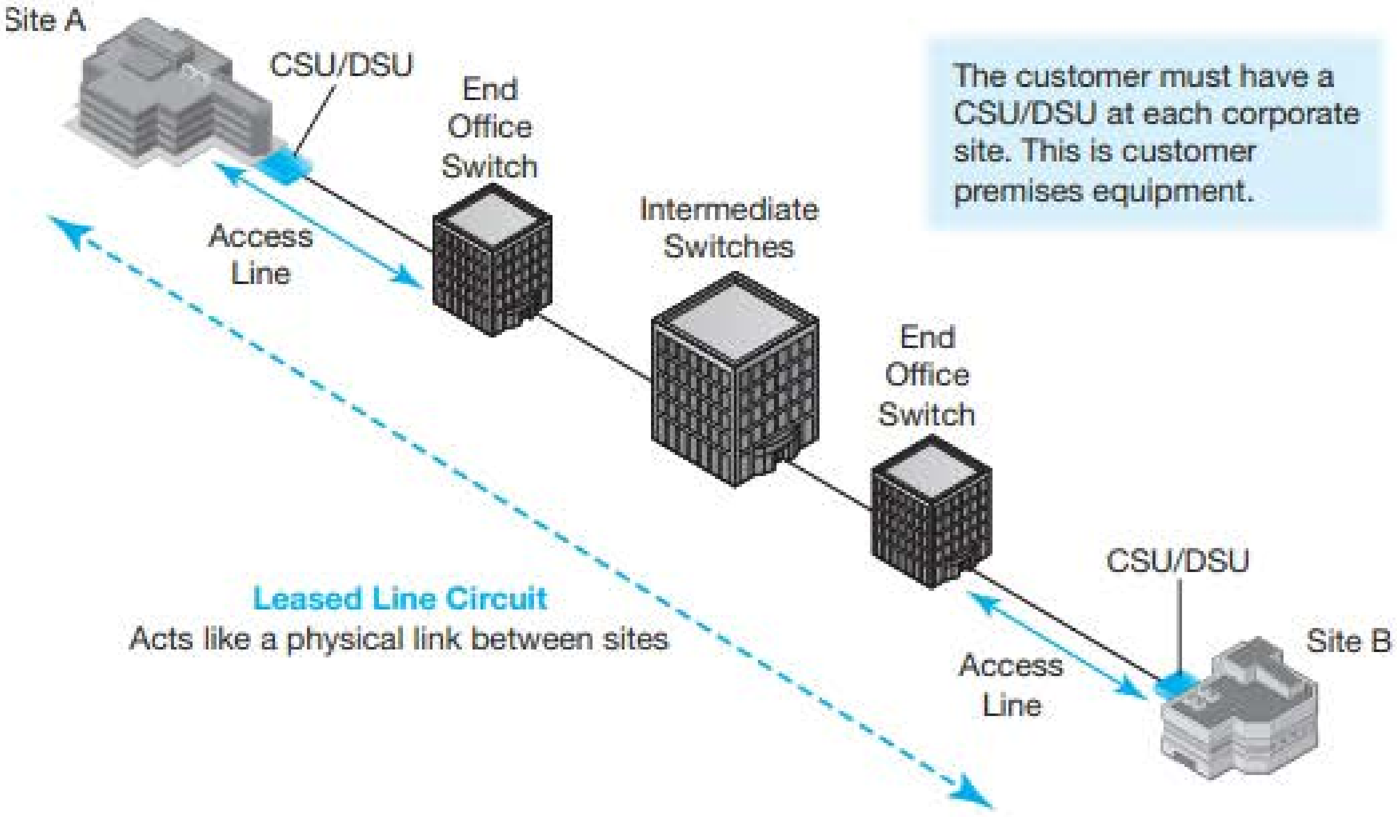
Site A
CSU/DSU
End Office Switch
The customer must have a CSU/DSU at each corporate site. This is customer premises equipment.
Intermediate Switches
Access Line
End Office Switch
CSU/DSU
Site B
Leased Line Circuit
Acts like a physical link between sites
Access Line

FIGURE 10-12 Leased Line, Trunk Lines, and Access Lines

To use a leased line, a company needs a piece of customer premises equipment called a **CSU/DSU**.[6] The purpose of this device is to translate the physical layer signals of network devices on the customer premises into physical layer signals in a format that leased lines require.

Test Your Understanding

13. a) What are the characteristics of leased lines? b) Distinguish between leased lines and access lines. c) What device must a customer have at its site to connect to a leased line?

Reaching the ISP via a Leased Line

A company needs to connect to its ISP. The simplest way to do this is to run a leased line from the company to the ISP's nearest access location. We know that this access line will pass through several transmission lines and switches, but networking professionals usually draw leased lines as they appear to be, namely a point-to-point transmission link. Figure 10-11 illustrates this approach.

Test Your Understanding

14. When a customer uses a leased line to connect to its ISP, what two points does the leased line connect?

Leased Line Private Corporate WANs

Companies need to communicate with their ISPs. If they have multiple sites, they also need to connect these sites into a coherent network for internal communication. Figure 10-13 shows that they can do this by building a leased line network that will create a private internal WAN. Site routers route packets among the sites.

Figure 10-14 shows that leased line speeds vary widely. Under 50 Mbps, leased line speed standards were set regionally. The United States and Canada use the North American Digital Hierarchy Standard. Europe uses the CEPT Hierarchy. Other countries may use different standards. Fortunately, it is possible to translate between different leased line hierarchies, but the diversity of standards does cause minor problems.

Above 50 Mbps, carriers have standardized on a single standard that is called Synchronous Optical Network (SONET)[7] or Synchronous Digital Hierarchy (SDH). SONET and SDH use different naming conventions for their lines. For example, SONET labels its lines with OC (optical carrier) numbers, while SDH uses STM (synchronous transport module) designations. Other than naming differences, their services are identical and compatible.[8]

The line naming conventions and speeds are easier to understand if you understand that all SONET/SDH speeds are multiples of 51.84 Mbps. The slowest OC line

[6] Channel Service Unit/Data Service Unit. Not very informative.

[7] SONET is the terminology used in the United States and Canada. The rest of the world uses the SDH nomenclature.

[8] Apart from a few unimportant differences.

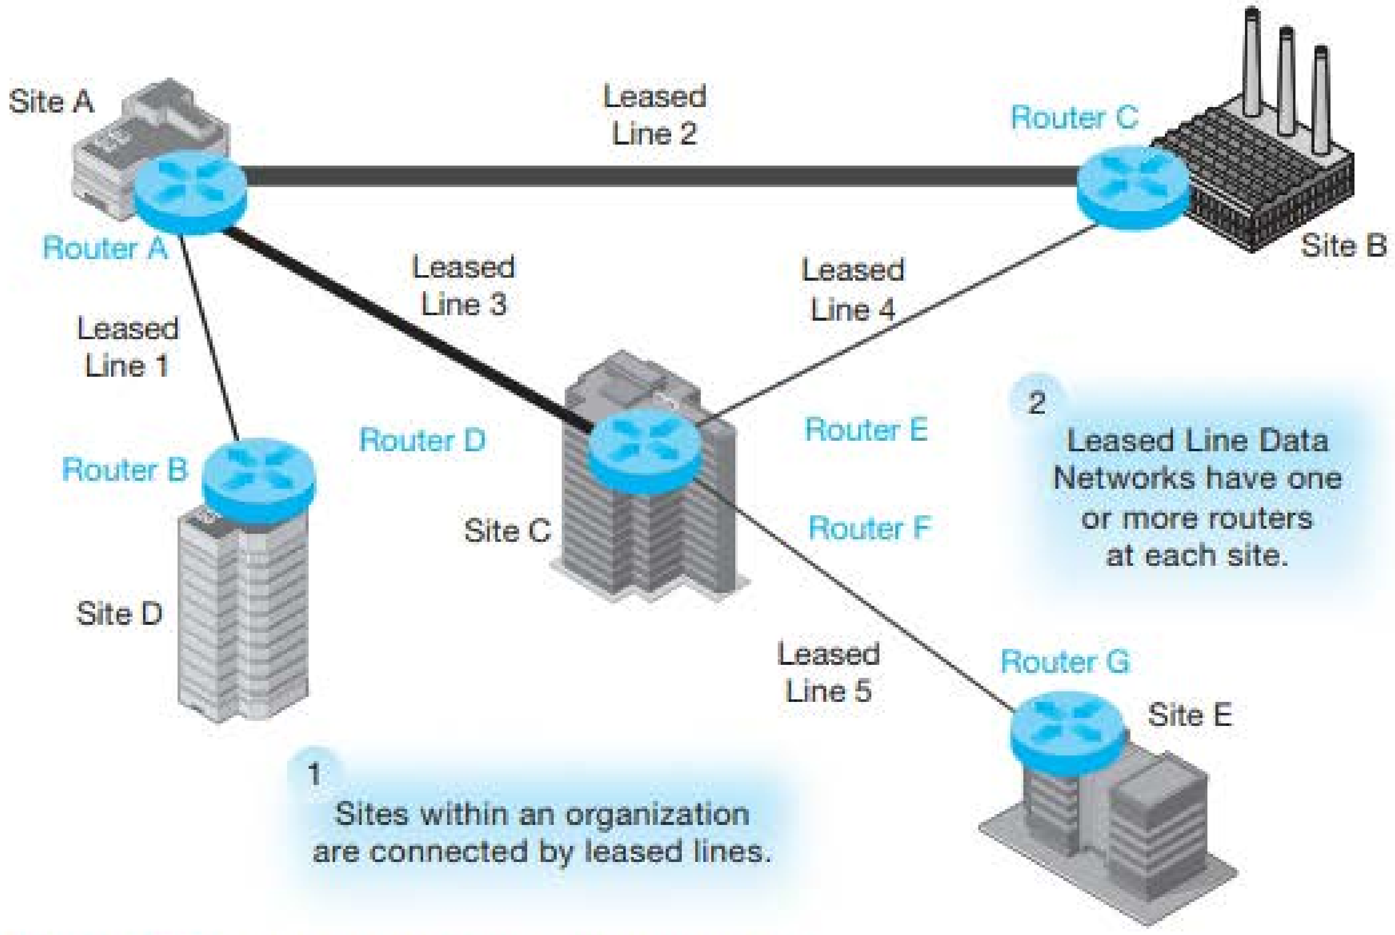

2 Leased Line Data Networks have one or more routers at each site.

1 Sites within an organization are connected by leased lines.

FIGURE 10-13 Leased Line Private Corporate WAN

that carriers offer is OC-3, which is three times the base speed. SDH carriers call this STM-1 because it is the first (slowest) speed they offer.

Applying Figure 10-14 Applying the information in Figure 10-14 is straightforward. If you have a requirement for a particular speed between two points, you select a leased line sufficient for that speed. For example, if you require a speed of 100 Mbps, you select an OC-3 or STM-1 line.

Carriers often offer more choices, predominantly at lower speeds. WAN line speeds traditionally were slow, around one to two megabits per second. This was roughly T1/E1 speed. Given frequent demand for a fraction of a T1 or E1 line, carriers typically offer fractional T1/E1 speeds for a fraction of the cost of a full T1/E1 line. If you need 200 kbps, you could get a fractional T1 line running at 256 kbps, which is 16.5% of a T1 line. As you might suspect, carriers will charge more than 16.5% of what they charge for a full T1 line.

Carriers also allow a customer to bond two or more T1/E1 lines together between a pair of sites. For example, if you need 2.8 Mbps between a pair of sites, you might bond two T1 or E1 lines.

Traditionally, T1/E1 leased lines required running a new 2-pair data-grade UTP line to the customer's premises. This is expensive. In addition, the telephone system already runs 1-pair voice-grade UTP to all premises, including business premises. We saw earlier in this chapter that carriers run asymmetrical digital subscriber line (ADSL) services over these lines. We also saw that ADSL today is much faster than T1/E1 speeds. Consequently, many carriers who offer "T1" and "E1" lines today are really offering DSL service over 1-pair voice-grade UTP.

North American Digital Hierarchy		
Line	Speed	Typical Transmission Medium
T1*	1.544 Mbps	2-Pair Data-Grade UTP
T3	44.736 Mbps	Carrier Optical Fiber
CEPT Hierarchy		
Line	Speed	Typical Transmission Medium
E1*	2.048 Mbps	2-Pair Data-Grade UTP
E3	34.368 Mbps	Carrier Optical Fiber
SONET/SDH Speeds		
Line	Speed (Mbps)	Typical Transmission Medium
OC3/STM1	155.52	Carrier Optical Fiber
OC12/STM4	622.08	Carrier Optical Fiber
OC48/STM16	2,488.32	Carrier Optical Fiber
OC192/STM64	9,953.28	Carrier Optical Fiber
OC768/STM256	39,813.12	Carrier Optical Fiber

*Often offer synchronous DSL over existing 1-pair voice grade UTP rather than offering traditional T1 and E1 service over 2-pair data grade UTP, which must be pulled to the customer's premises.

Fractional T1 speeds are often offered by carriers. These typically include some subset of the speeds 128 kbps, 256 kbps, 384 kbps, 512 kbps, and 768 kbps.

T1 and E1 lines can be bonded to provide double, triple, or quadruple the capacity of a single line.

FIGURE 10-14 Leased Line Speeds

However, carriers do not offer asymmetric DSL service to organizations because businesses need symmetric speed—the same speed in both directions. Consequently, carriers offer **synchronous DSL** services to businesses. Businesses also require quality-of-service (QoS) guarantees, so these synchronous DSL lines come with service level agreements. SLAs mean that the DSL services offered to businesses are more expensive per bit transmitted than residential ADSL service.

Managing the WAN Leased line corporate WANs do not design and operate themselves. A company that uses leased line networks to connect its sites faces substantial labor and customer premises equipment costs.

Test Your Understanding

15. a) If you need a speed of 1.2 Mbps between two points in the United States, what leased line would you specify in the United States and in Europe? b) Repeat for 160 Mbps. c) Repeat for 3 Mbps. d) Why do carriers offer low-speed "leased lines" that are really DSL lines? e) How do business DSL lines differ from residential DSL lines? f) Why is the need to manage the leased line network an issue?

CARRIER WAN SERVICES

It is possible for corporations to do all their internal wide area networking using leased lines. But building and managing complex networks of leased lines is a great deal of work that requires a networking staff with high expertise. Instead, most firms turn to carriers, who offer complete Layer 2 and Layer 3 WAN services. There are large economies of scale in managing such networks, so carriers can offer them at an attractive price per bit transmitted. Two technologies dominate these carrier WAN services today. These are carrier Ethernet and Multiprotocol Label Switching. We now look at critical features of both.

Carrier Ethernet

In the 1980s, there were many LAN technologies. However, Ethernet alone survived, thanks to its low-cost operation and its ability to grow to ever-faster speeds. Today, Ethernet is available in wide area networking. This extension of traditional Ethernet was originally called metropolitan Ethernet. Today, it is called **carrier Ethernet**. Carrier Ethernet is not exactly the same as Ethernet for LANs. Carrier Ethernet requires some extensions, but they are not large. If you know traditional Ethernet, it is straightforward to extend your expertise to carrier Ethernet. Carrier Ethernet services are developed by MEF (formerly the Metropolitan Ethernet Forum), and detailed technical standards are created by the IEEE 802.3 Working Group. Figure 10-16 shows that two of these services have dominated so far.

- **E-Line Service** is a site-to-site service. It competes directly with leased lines but handles frame formatting and other Layer 2 functionality.
- **E-LAN Service** essentially extends the LAN to the wide area. Sites can use Ethernet to communicate back and forth as if the Carrier Ethernet was simply a set of trunk lines between switches.

In both services, and in several other carrier Ethernet services, the leased line carrier service terminates in a CSU/DSU that is connected to an Ethernet switch instead of a router. This means that the service does not require TCP/IP expertise or become involved with the complexities of TCP/IP. E-line service essentially offers a long-distance trunk link between Ethernet switches. E-LAN service, in turn, acts to connect Ethernet site LANs into what is effectively a super LAN.

	Traditional Ethernet	Carrier Ethernet
Use Case	LANs	WAN connections, mostly in metropolitan areas
Operates at Layers	1 and 2	1 and 2
Standards Creation	Standards are created by the IEEE 802.3 Working Group.	Services defined by MEF. Standards extensions created by the IEEE 802.3 Working Group.
Standards Scope	Core Ethernet 802.3 standards	These, plus some extensions developed by the IEEE 802.3 Working Group.

FIGURE 10-15 Traditional Ethernet versus Carrier Ethernet (Study Figure)

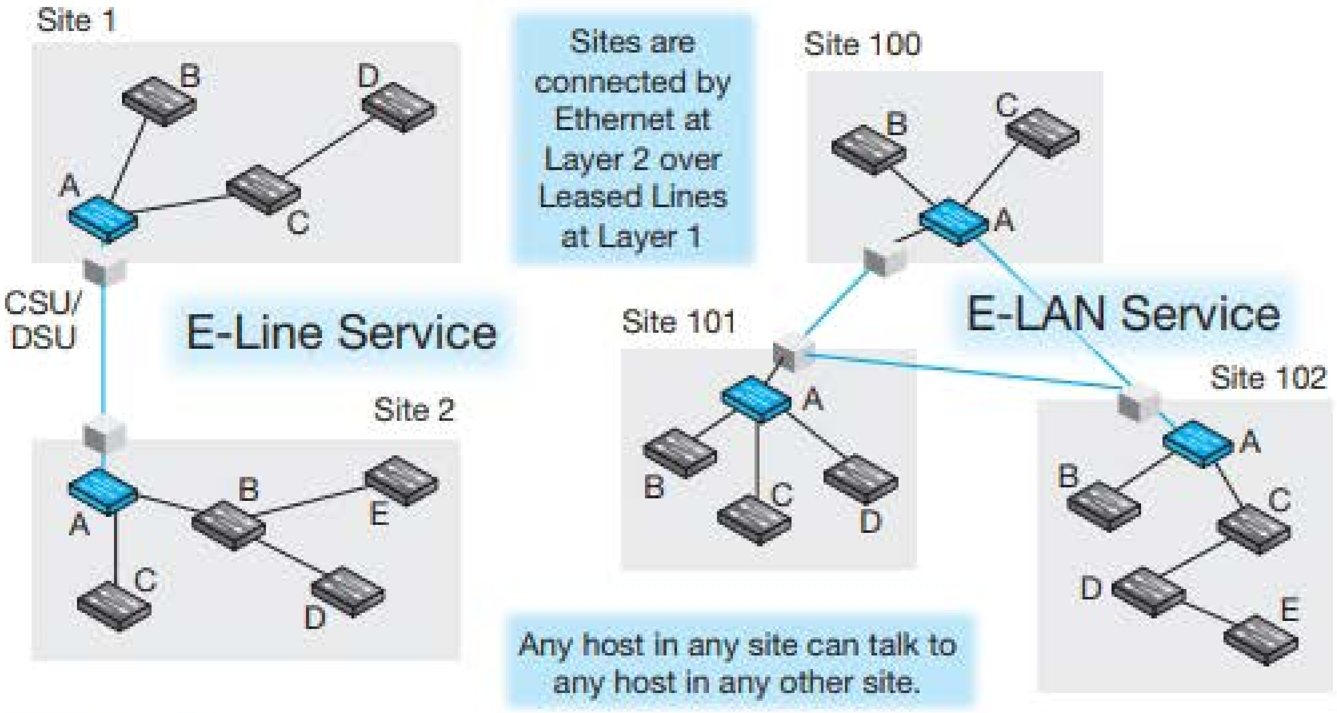

Sites are connected by Ethernet at Layer 2 over Leased Lines at Layer 1

Any host in any site can talk to any host in any other site.

FIGURE 10-16 Using E-Line and E-LAN Carrier Ethernet Services to Extend Ethernet LANs

Carrier Ethernet has a number of attractions.

- *Cost*. Using Ethernet's familiar low-cost MAC layer functionality, carrier Ethernet is inexpensive.
- *Familiarity*. Sites only have to plug carrier termination equipment into an Ethernet switch port. There is no need to learn a new technology.
- *Speed*. Companies that need fast connections can get 100 Mbps, 1 Gbps, or 10 Gbps at attractive cost.
- *Speed Agility*. If companies need extra capacity for a limited period of time, such as a year-end crunch, carrier Ethernet carriers can usually reprovision their services quickly.
- *Quality of Service*. Carrier Ethernet can offer quality of service Ethernet guarantees for speed, availability, frame delay, frame jitter, and frame loss.
- *Security*. Although Carrier Ethernet does not include cryptographic protections, the traffic of different customers is kept separate to prevent eavesdropping. Only some offer cryptographic security beyond this traffic segregation.

Test Your Understanding

16. a) Why is it attractive for companies to use Layer 2 and Layer 3 WAN services offered by carriers? b) How does carrier Ethernet differ from traditional Ethernet? c) What is the distinction between E-Line and E-LAN services? d) What are the attractions of carrier Ethernet for corporations? e) What is speed agility?

Low cost per bit transmitted
Uses familiar Ethernet technology
High speeds available
Speed agility: increases in speed can be provisioned rapidly
Quality of service
Security by segmenting customer traffic but not always cryptographic security

FIGURE 10-17 Carrier Ethernet Attractions (Study Figure)

Multiprotocol Label Switching (MPLS)

Making Routing More Efficient In Chapter 8, we saw that routers look at an incoming packet's destination IP address. They compare that IP address to every row in the routing table, select the best match, and send the packet back out a certain port to a certain IP address. The next packet to arrive gets the same treatment—even if it goes to the same IP address.

Many routers can do decision caching, in which they remember their decisions for certain IP address ranges. We saw that this is dangerous. Fortunately, there is a more robust way to avoid having to look at all rows for all packets. This is **Multiprotocol Label Switching (MPLS)**, which Figure 10-18 illustrates.

Operation When two hosts start to converse, an MPLS network first determines the best path for the packets. This is the **label switched path**. Routers will send all packets along this path rather than making traditional routing decisions for each packet at each router.

As Figure 10-18 shows, after the label switched path is established, the source host transmits packets normally. The first router is a **label switching router**. It inserts a 32-bit **label header** in front of the IP header and after the frame header. The IP packet syntax and the frame syntax are unchanged.

The label header's **label number** identifies the label switched path selected for this conversation. The first label switched router and all others along the label switched path have MPLS lookup tables. These tables allow routers to look up the label number, read the corresponding interface, and send the packet out the indicated interface. For example, if the label number is 47, the router in Figure 10-18 will send it out Interface 1. Table lookups are fast because only one row will match the destination IP address. There is no need to look at all routing table rows to select the best interface to send a packet back out. The hard work was done when the label switched path was created.

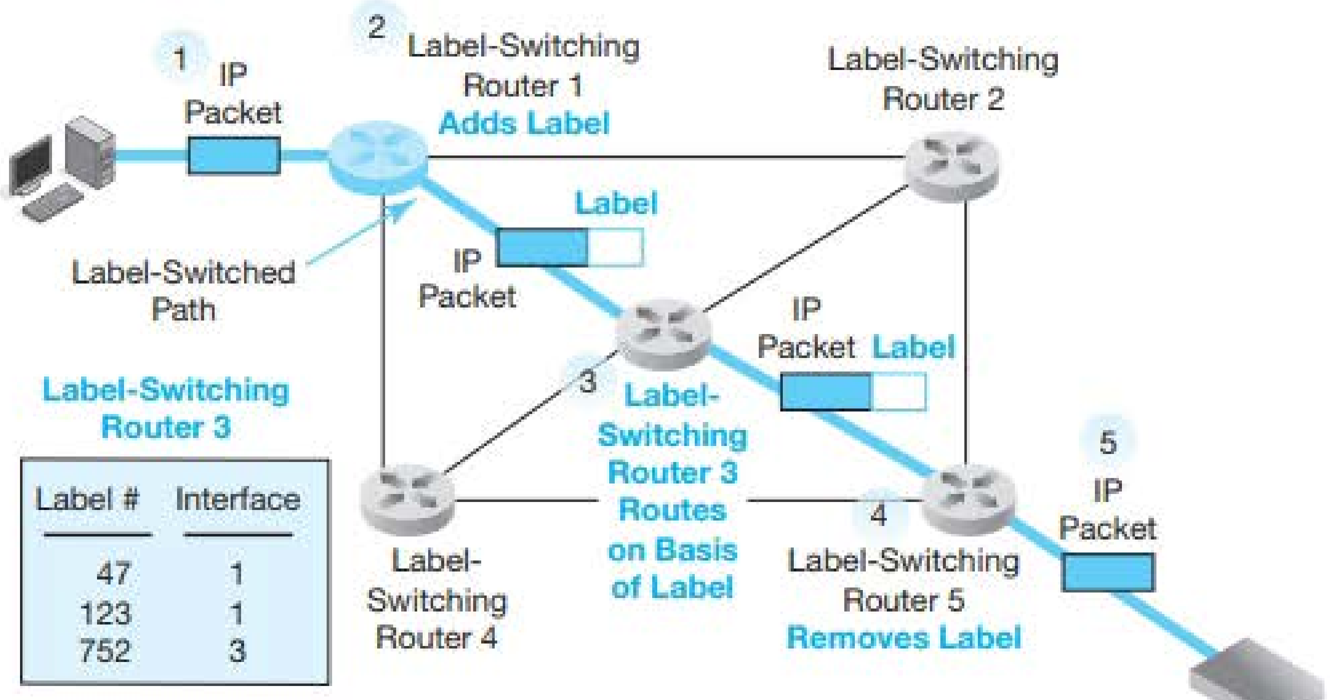

Label #	Interface
47	1
123	1
752	3

MPLS reduces forwarding costs and permits traffic engineering, including quality of service and traffic load balancing.

FIGURE 10-18 Multiprotocol Label Switching (MPLS)

The last label switching router removes the label. Note that neither the source host nor the destination host knows that label switching was done. MPLS operation is transparent to hosts.

Often, all traffic between two sites is assigned a single label number. Or, traffic between two sites might receive one of a handful of label numbers. Different label numbers might correspond to label switched paths with different quality of service characteristics.

Benefits MPLS offers three major benefits.

- First, MPLS slashes the work each router must do and therefore slashes a company's router costs.
- Second, as we just noted, MPLS can be used to assign paths based on the QoS requirements of different packets.
- Third, MPLS can do **traffic engineering**, that is, manage how traffic will travel through the network. One traffic engineering capability is **load balancing**, that is, moving some traffic from a heavily congested link between two routers to an alternative route that uses different and less-congested links.

Carrier MPLS Companies can create their own MPLS networks, but they typically use carriers to provide MPLS for their WAN communication. Many of these carriers are Internet service providers who already use MPLS within their own internets. They extend the benefits of MPLS to their customers. That is impossible for the Internet as a whole because there is no central control organization for the Internet.

An Expensive Service Given that MPLS should reduce router costs when delivering individual packets, you might think that MPLS would be an inexpensive service compared to just using the Internet. Unfortunately, that is not the case. The main attraction of MPLS for corporations is its ability to implement strong QoS service level agreements (SLAs). Corporations cannot live with the "best effort" limitations of the core Internet for many of the critical transmission needs in the core transaction processing applications that drive firms. MPLS allows network administrators to control traffic priorities in a flexible way, ensuring that critical services get the service they require. Companies pay high prices for this manageability, and although they are not happy about these prices, they need MPLS to satisfy their QoS needs.

Extendibility If MPLS is so good, why not use it everywhere on the Internet? The answer is that MPLS requires a single administrator to manage the entire network of label switched routers.

Test Your Understanding

17. a) In MPLS, is selecting the best interface for each packet at each router done when the packet enters the network or before? b) Why is this beneficial? c) What is the name of the path selected for a particular conversation? d) When a source host first transmits to a destination host after a label switched path is established, what will happen? e) Do label switching routers along the MPLS path look at the packet's IP address? The answer is not explicitly in the text. Explain your reasoning. f) On what basis does each label switched router base routing decisions? g) Why is MPLS transparent to the source and destination hosts? h) What are MPLS's attractions? i) What is traffic engineering? j) Can MPLS provide traffic load balancing? k) How does the price of using MPLS compare with the price of simply sending traffic over the Internet? l) Why do firms pay this price difference?

WAN Optimization

Given the high cost of long-distance transmissions, companies need to squeeze out every bit of performance improvement they can find for data over WANs. Figure 10-19 shows that one approach is to install **WAN optimization devices** at each end of important shared lines between sites.

Compression The most important action that WAN optimization devices take is to compress all data being transmitted into the line and decompress the data at the other end. **Compression** is possible because almost all data contains redundancy that can be reduced through encoding. For movies and voice, compression can be substantial. For word processing documents and spreadsheets, compression is less effective. In the figure, the WAN optimization devices can provide an average of 10:1 compression. Source A is transmitting at 3 Gbps, and Source B is transmitting at 5 Gbps. This is a total of 8 Gbps arriving at the WAN optimization device. However, with 10:1 compression, the transmission line only has to carry 0.8 Gbps. This will fit in a 1 Gbps transmission line. Without compression, the company would need a much more expensive 10 Gbps transmission line.

Caching Another way to reduce the number of bits flowing through the transmission line is **caching**. (See Figure 10-20.) Suppose the company produces a large annual report. The server holding the report is in Source A. The annual report is likely to be transmitted multiple times from Source A to recipients in Source C and Source D. With a WAN

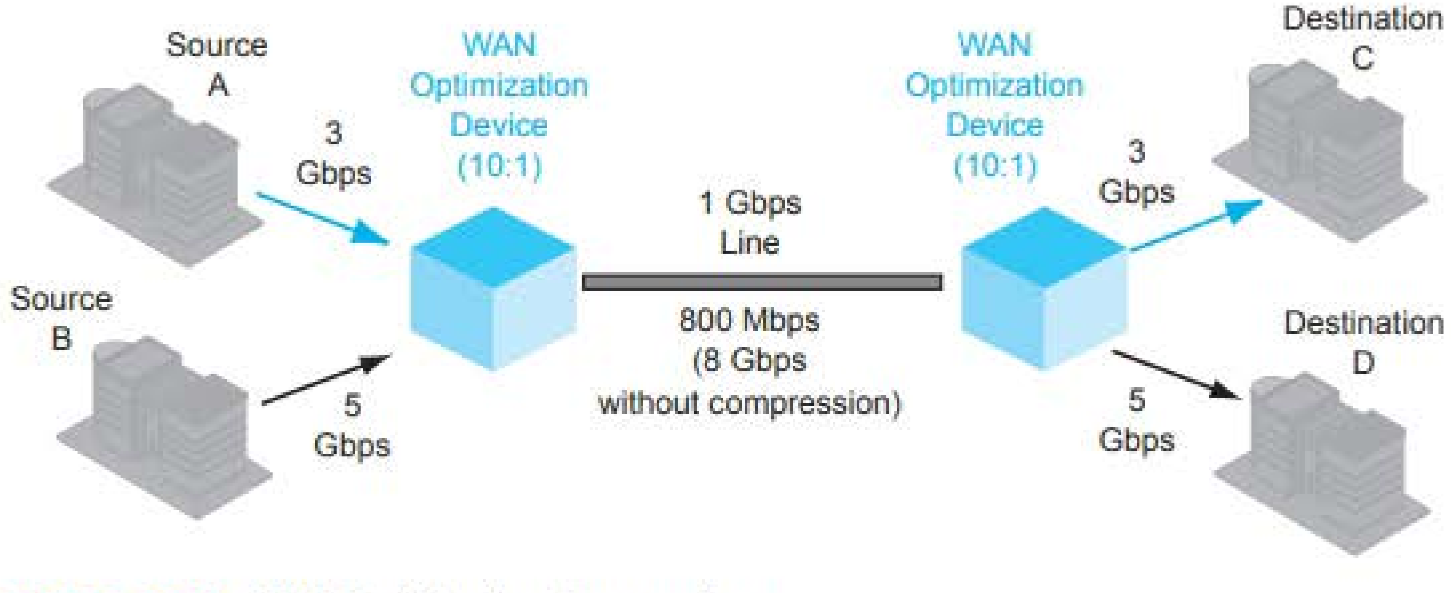
Source A
3 Gbps
WAN Optimization Device (10:1)
1 Gbps Line
800 Mbps (8 Gbps without compression)
WAN Optimization Device (10:1)
3 Gbps
Destination C
Source B
5 Gbps
5 Gbps
Destination D

FIGURE 10-19 WAN Optimization: Compression

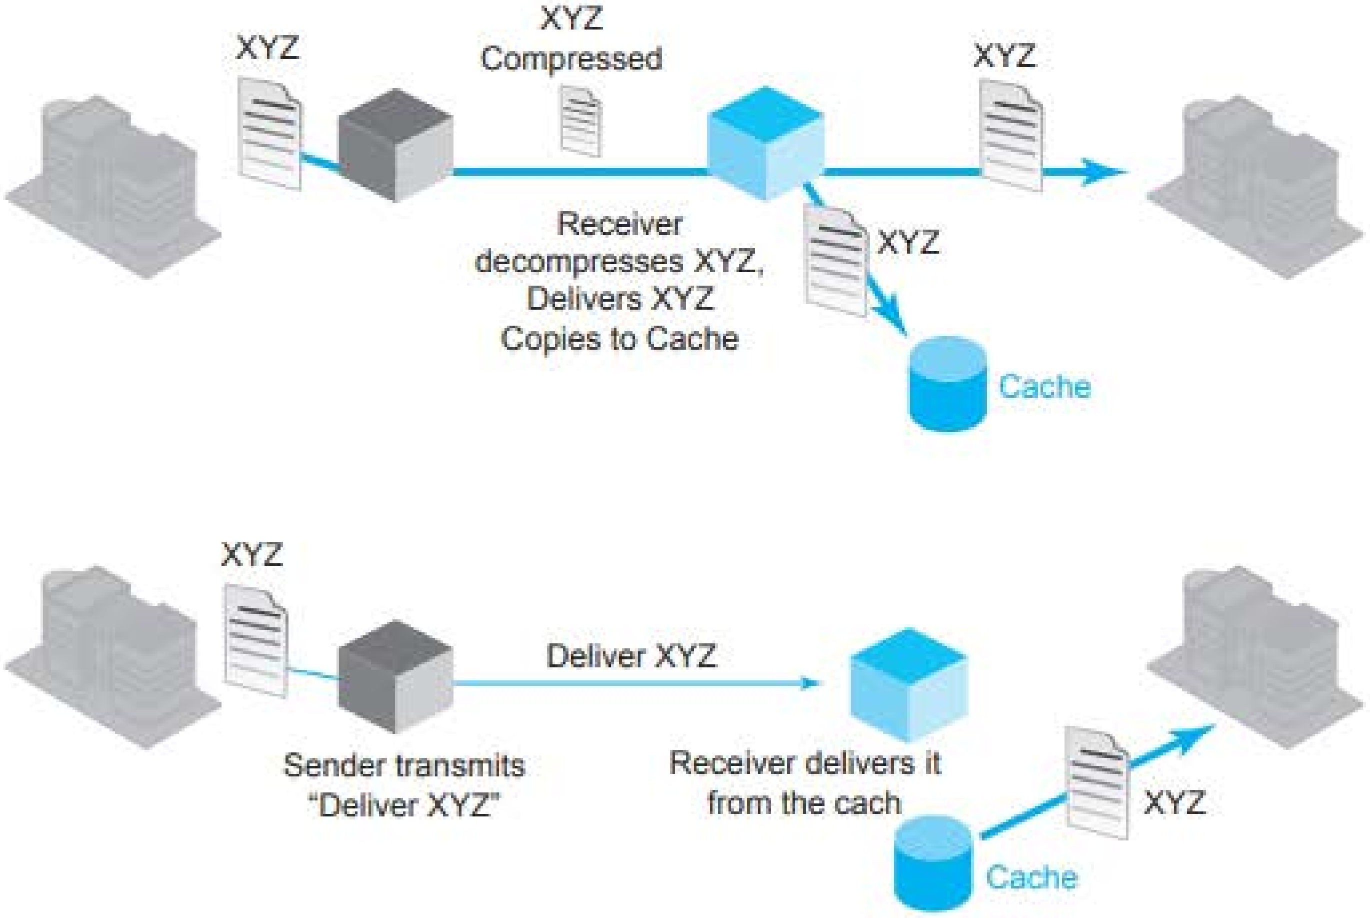
XYZ
XYZ Compressed
XYZ
Receiver decompresses XYZ, Delivers XYZ Copies to Cache
XYZ
Cache
XYZ
Deliver XYZ
Sender transmits "Deliver XYZ"
Receiver delivers it from the cach
XYZ
Cache

FIGURE 10-20 WAN Optimization: Caching

optimization device that has caching, when the annual report is first delivered, it is copied onto the receiving WAN optimization device's disk **cache**. Later, when the annual report is to be transmitted again, the WAN optimization device near Source A and Source B will not transmit the entire file. Instead, it will send a brief message to the WAN optimization device near Source C and Source D. This message asks the WAN optimization device on the right to retrieve the annual report from the cache and send it to the receiver. Avoiding the retransmission of frequently transmitted files can reduce traffic considerably.

Traffic Shaping In many cases, unfavored applications take up too much capacity. Unfavored applications may include YouTube, Netflix, and BitTorrent for file sharing. Some WAN optimization devices do **traffic shaping**. (See Figure 10-21.) When undesired traffic reaches an optimization device, the device may simply prohibit

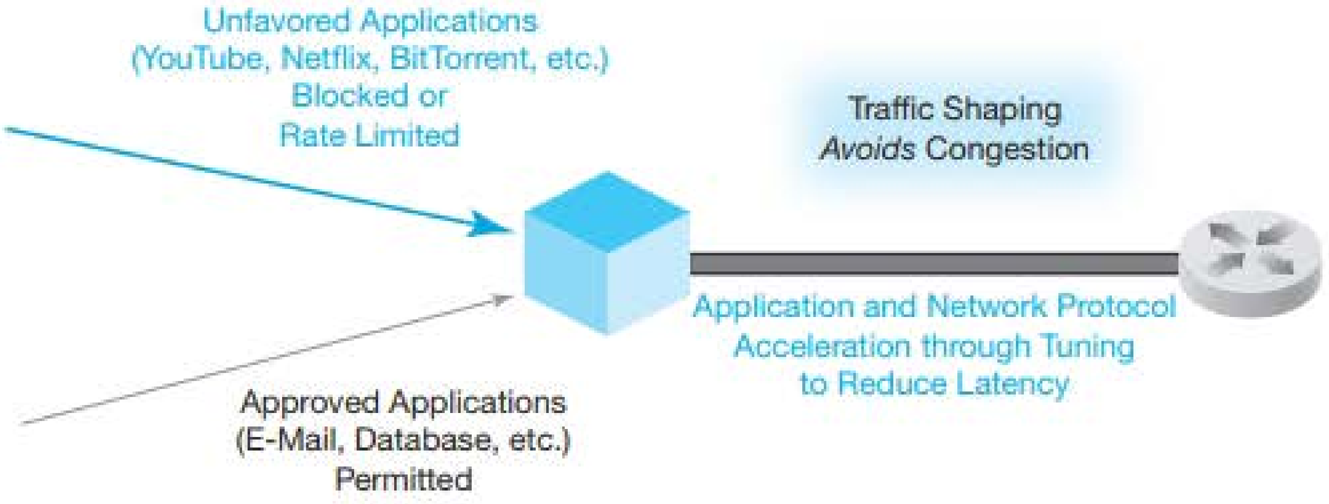
Unfavored Applications (YouTube, Netflix, BitTorrent, etc.) Blocked or Rate Limited
Traffic Shaping *Avoids* Congestion
Application and Network Protocol Acceleration through Tuning to Reduce Latency
Approved Applications (E-Mail, Database, etc.) Permitted

FIGURE 10-21 WAN Optimization: Traffic Shaping and Application and Network Protocol Acceleration

it. The device can also take a less drastic action—limiting the application to a small percentage of the total traffic. Both can dramatically reduce overall traffic, allowing the firm to avoid upgrading its transmission lines.

Application and Network Protocol Acceleration (Tuning) Many applications and network protocols are somewhat inefficient when they transmit over long-distance lines. TCP, for example, tends to have conservative transmission defaults that slow transmission. It may be possible to tune TCP by adjusting such things as time spent waiting for acknowledgments before retransmitting a TCP segment. To give another example, when a WAN optimization device receives a TCP SYN segment, it may send back an ACK even before it passes the segment on to its intended host. Application and network acceleration is a family of tactics the WAN optimization devices can use to reduce *latency*, which tends to be a problem when signals must travel long distances. Although tuning can take place on hosts, WAN optimization devices provide a central point for tuning and tuning tools.

> *Application and network acceleration is a family of tactics the WAN optimization devices can use to reduce* latency, *which tends to be a problem when signals must travel long distances.*

Test Your Understanding

18. a) Where are WAN optimization devices found? b) List the four mechanisms we discussed for optimizing transmission over a transmission link. c) How does compression reduce traffic? d) How does caching reduce traffic? e) Explain traffic shaping. f) How does traffic shaping reduce traffic? g) What is the main benefit of application and network protocol acceleration?

END-OF-CHAPTER QUESTIONS

Thought Questions

10-1. Distinguish between dial-up telephone service you use as a consumer and leased line services used in business. (You will have to extrapolate from your own experience with dial-up lines.)

10-2. If you have a network of leased lines, you have options for how many sites you connect. Sites can communicate directly or through intermediate sites. a) In a full mesh, every pair of sites will be directly linked by a leased line. If there are *N* sites, there will be $N^*(N-1)/2$ connections. In Figure 10-13, how many leased lines would be used in a full mesh? b) In a hub-and-spoke network, there is a central site, and a leased line radiates from it to each other site. In Figure 10-13, how many leased lines would be used in a hub-and-spoke network with the hub located at Site A? c) What is the benefit of full mesh networks over hub-and-spoke networks? d) What is the advantage of hub-and-spoke networks over full mesh networks? e) How would you use this information about advantages to advise a company about what to do when it installs a network of leased lines?

10-3. In ADSL service, there is a single UTP pair running from the end office switch to the individual household. In cable modem service, the thick coaxial cable in the neighborhood is shared by many subscribers. Yet, typically, cable modem service provides faster service to individual customers than ADSL. How can this be? Hint: Draw a picture of the entire situation for both ADSL and cable modem service.

10-4. a) What two wired WAN technologies are growing rapidly? b) Compare their relative attractions and main uses. c) Why will leased lines continue to be important even if networks of leased lines are no longer used?

Hands-On

10-5. If you have a smartphone, download an app to tell your data transmission throughput. What did you find?

Perspective Questions

10-6. What was the most surprising thing you learned in this chapter?

10-7. What was the most difficult part of this chapter for you?

Chapter **11**

Networked Applications

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Explain core concepts in networked applications and application architectures.
- Describe how taking over an application can give an attacker the ability to control the computer.
- Describe how Netflix uses cloud computing and how this illustrates the importance of host technology (and cloud computing specifically) as a driving force for networking.
- Describe the World Wide Web in terms of standards and explain how a webpage with text, graphics, and other elements is downloaded.
- Describe electronic mail standards and security.
- Describe voice over IP (VoIP) operation and standards.
- Explain why peer-to-peer (P2P) computing is both desirable and dangerous.

INTRODUCTION

We finally arrive at Layer 5, the application layer. This is the only layer users care about. (Lower layers come to the attention of users only when they fail.) Of course, improvements in lower layers have a major impact on application design. However, application quality and performance are always the litmus tests for users.

Networked Applications and Application Architectures

The first computers were mainframes that worked with dumb terminals that were remote keyboards plus attached printers or displays. (This was before microprocessors, so terminals were not smart.) All applications used **stand-alone processing** on the central computer. PCs also began as stand-alone devices in which a program ran on a single machine.

In stand-alone processing, all application processing is done on a single machine.

Today, however, most applications are **networked applications,** which are sets of interacting programs running on two or more hosts. When a software developer writes a program, it is usually a program designed to talk to another program on a different host. This radically changes the way programs are written, tested, and deployed. We will focus on networked applications.

Networked applications can use an infinite number of alternative interaction patterns. However, two have dominated, and two more are emerging. We call patterns of interactions between networked applications on different machines **application architectures**. Figure 11-1 lists four application architectures.

Application architectures are patterns of interactions between networked applications on different machines.

- As just noted, the first computers were large mainframes connected to dumb terminals. They did stand-alone processing.

Networked Applications

- Networked applications require networks to operate
- World Wide Web, e-mail, etc.

Application Architectures

- How application layer functions are spread among computers
- Driven in part by growing client processing power, memory, etc.
- Stand-alone computing
- Client/server architecture
- Peer-to-peer architecture
- Distributed computing architecture

Changing Programming and Server Locations

- Programmers now write software on one machine that interacts with software on other machines
- Programmers must understand application architectures and networking
- Falling networking costs are resulting in the consolidation of servers

FIGURE 11-1 Basic Networked Application Concepts

- When PCs began to appear, their processing power, memory, and communication abilities were rudimentary. Programmers again wrote *stand-alone programs*, this time running on desktop PCs and later on laptops. One stand-alone application emulated a dumb terminal, allowing PC users to communicate with mainframes over dial-up telephone modems that usually sent and received at a mere 9,600 bps or less. Terminal emulation software turned expensive PCs into cheap dumb terminals. The mainframe, unaware that the device at the end of the transmission line had some intelligence, continued to run stand-alone programs.
- As PCs grew more powerful, we saw the rise of the programs that were true network applications. These implemented the **client/server architecture** in which the now-smart PC shared processing chores with servers. The World Wide Web was the driving force behind the client/server architecture.
- Later, clients grew as powerful as older servers. This meant that a client could provide server capabilities to another client, eliminating the need for a server. PCs sit idle most of the time, so this **peer-to-peer computing** had the potential to greatly reduce the number of servers a company needed to buy and manage.
- We seem to be moving toward an era of **distributed computing architecture**, in which a program running on one machine calls multiple programs on other machines, which may call programs on yet other machines. After calling other programs, the calling program uses results from the called programs in its own logic flow.

Programming Networked Applications As application architectures grow more complex, the job of writing programs to run on them also becomes more complex. Today's programmers mostly write programs that call other programs on other machines to do their work. In this sense, networking has revolutionized programming. At the same time, more complex architectures create greater challenges for networking professionals because growth in client/server computing and other innovations places demands on networks that are greater in terms of speed and often in other characteristics, such as latency. When you talk to your voice assistant on your phone, tablet, or computer, back-end servers must do sophisticated artificial intelligence processing before your local device replies. This and other application development trends require slashing latency in round-trip response time to a few milliseconds.

Changing Server Locations Today, network transmission costs are falling rapidly. As they do, the economic need to keep interacting computers very close together continues to fade. For several years, companies with dozens of local server rooms have been consolidating in a few regional locations to take advantage of economies of scale and inexpensive rural areas. Taking advantage of this trend, cloud computing involves outsourcing the ownership and management of servers, programs, or both to a cloud service provider such as Amazon Web Services or Microsoft Azure. We look at an example of cloud computing in the next section so you can understand its far-reaching implications for networking. We also look at cloud computing because it is important in and of itself to IT professionals.

Test Your Understanding

1. a) What is a networked application? b) What is an application architecture? c) Which application architecture is dominant today? d) What host innovation brought it about? e) What is the major driving force behind the peer-to-peer (P2P) application architecture? f) How is programming changing because of new application architectures? g) What change is falling network cost driving?

Application Security

In the past, hackers focused primarily on vulnerabilities in the operating system to break into computers. Today, however, hackers primarily attack individual applications running on the computer.

The reason for this is shown in Figure 11-2. If a hacker can take over an application, then he or she receives all the permissions that the operating system gives the application. Many applications run with **root privileges**, which means that they can do anything on the computer. Taking over such an application gives the hacker total control over the computer.

If a hacker can take over an application, then he or she receives all the permissions that the operating system gave the application.

Attackers now agree that finding a vulnerability in the operating system is very difficult today. However, with the many applications running on most computers, and with inconsistent security quality across applications, the probability of finding a vulnerable application on a computer is high. Security vulnerabilities in specific applications are listed in many hacker forums that are readily available to attackers.

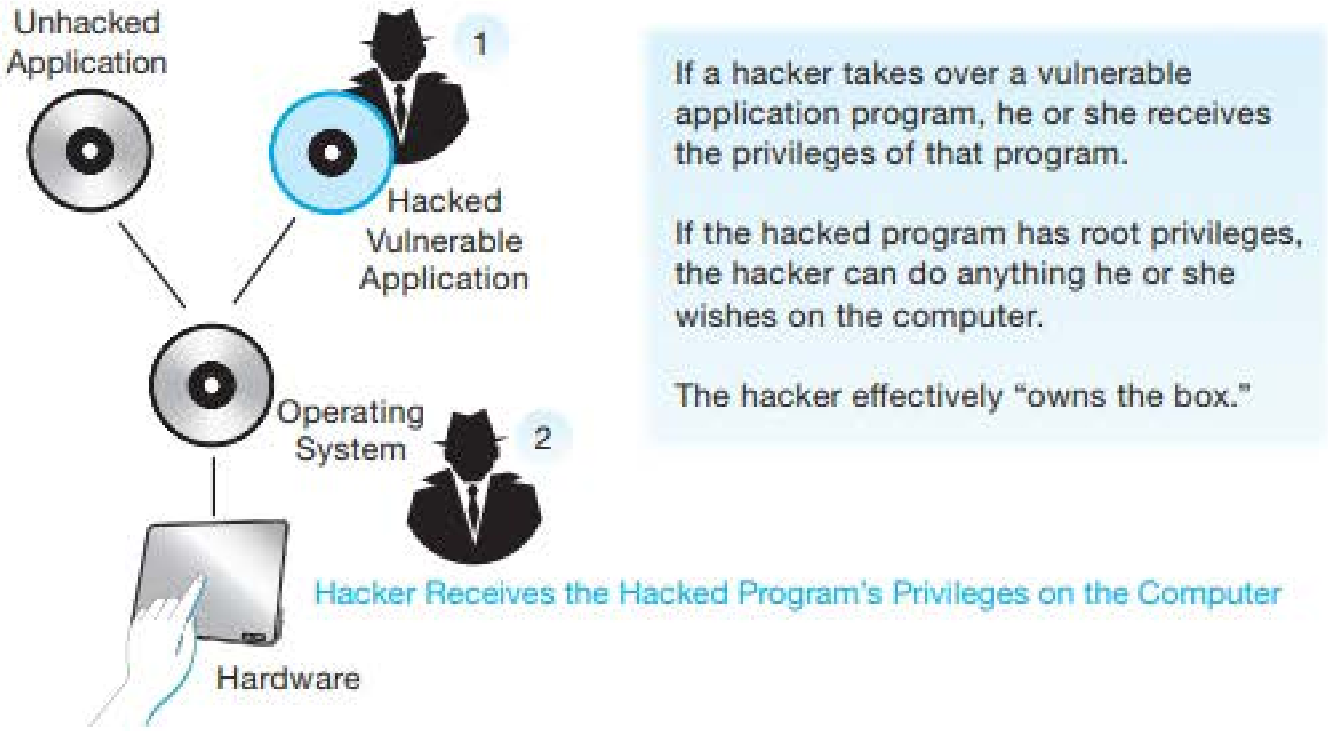

If a hacker takes over a vulnerable application program, he or she receives the privileges of that program.

If the hacked program has root privileges, the hacker can do anything he or she wishes on the computer.

The hacker effectively "owns the box."

Hacker Receives the Hacked Program's Privileges on the Computer

FIGURE 11-2 Application Hacking

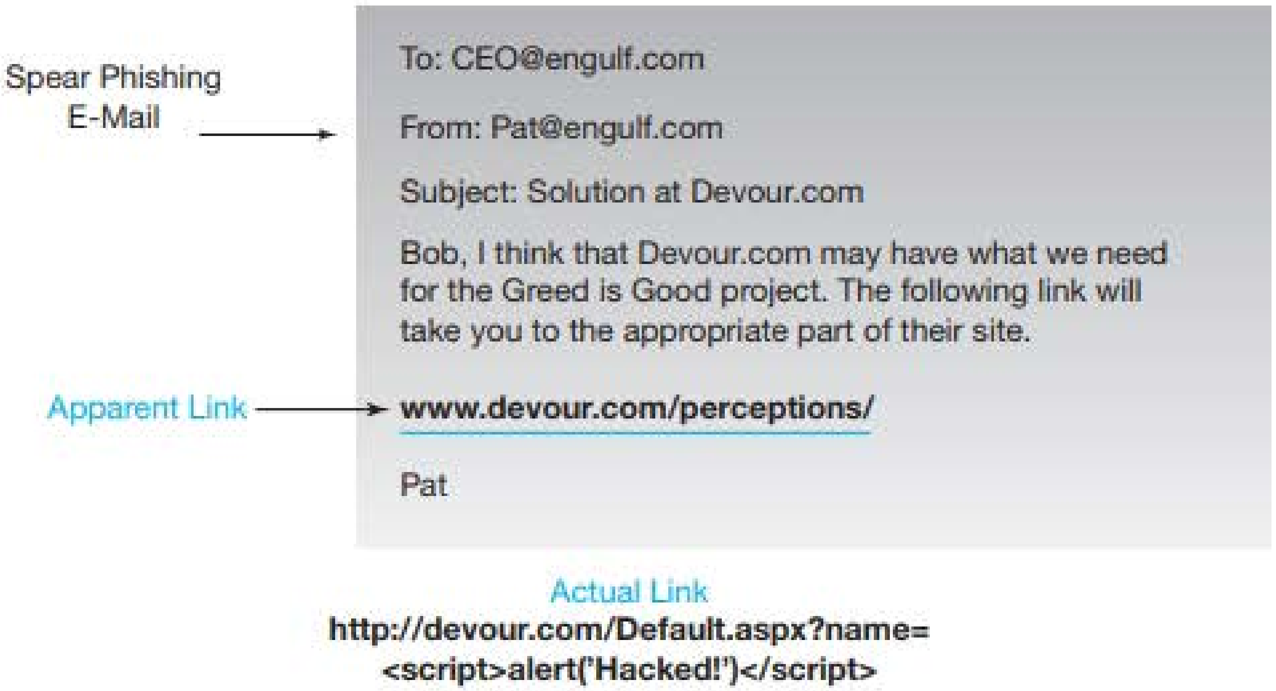

To: CEO@engulf.com

From: Pat@engulf.com

Subject: Solution at Devour.com

Bob, I think that Devour.com may have what we need for the Greed is Good project. The following link will take you to the appropriate part of their site.

www.devour.com/perceptions/

Pat

Spear Phishing E-Mail
Apparent Link
Actual Link
http://devour.com/Default.aspx?name=<script>alert('Hacked!')</script>

FIGURE 11-3 Cross-Site Scripting (XSS) Attack Using Reflection

In particular, we are seeing an explosion in **apps**—small applications created for mobile devices. In addition, we are seeing diversity in mobile operating systems. The newness of mobile operating systems and mobile applications has led many inexperienced developers to create applications with severe vulnerabilities. Coupled with a lack of corporate control over mobile devices, this lack of experience has created a flood of application (and operating system) vulnerabilities.

Cross-Site Scripting (XSS) There are many ways to hack application programs. One popular attack vector is the **cross-site scripting (XSS)** attack. In these attacks, the application asks a user for an input variable such as his or her name. The user may enter the name "Pat." The website then creates a webpage that says something like "Hello Pat." This is called **reflection**. It is dangerous because the webpage will contain whatever the user chooses to input.

Figure 11-3 shows why reflection is dangerous. In this example, the attacker sends the CEO of a corporation an e-mail message that purports to be from a subordinate. The message contains an apparently safe link to devour.com. Presumably, the company uses devour.com extensively, so the CEO sees the site as "safe."

In HTTP, the text that appears for a link may not be the true link. In Figure 11-3, the actual link is *http://www.devour.com/Default.aspx?name=<script>alert('Hacked!')</script>*. The link does take the victim to default.com. However, the problem is that it does more than that.

Most important, it will pass information to a particular program on Devour.com, Default.aspx. Default.aspx expects an input string for its name variable. Not shown in the figure, Default.aspx will reflect this name on a webpage. Probably, it will include something like "Hello *name*" on the webpage.

Given the e-mail message's crafted URL, however, the webpage being visited will reflect the script *<script>alert(Hacked!')</script>* on the webpage. When scripts are placed on a webpage, the user does not see them. However, the script executes when the page is rendered. This script is not damaging. The user will see a pop-up alert box that contains the message "Hacked!"

Most XSS attacks are extremely damaging. For example, the script may steal the user's login cookie and send it to the attacker. This may give the attacker the victim's username and password. XSS attacks can also redirect the victim to another webpage and install malware while making it look like nothing has occurred.

Cross-site scripting attacks do not always use e-mail or websites with deceptive links. For example, suppose a legitimate site allows user comments on webpages. Typically, the user enters text in a dialog box. The website then writes the comments onto the bottom of the webpage. If the comment contains a script, the script will execute every time someone visits the webpage afterward.

How can website designers thwart XSS attacks? At the broadest level, programmers should never trust user input. If information is to be reflected onto a webpage, the programmer must test the user input. It may seem simple to identify <*script*>and</*script*> tags, but scripts can be obfuscated (made less obvious). Also, there are many cross-site scripting attacks that do not use scripts. Thwarting cross-site scripting attacks is a difficult skill. This may explain why XSS vulnerabilities are pandemic on websites.

Programmers should never trust user input.

Test Your Understanding

2. a) Why are hackers now focusing on taking over applications? b) What can hackers do if they take over an application with root privileges? c) Why is the explosion of "apps" on small mobile devices a concern?
3. a) Why is reflecting a user's input dangerous? b) What attitude should programmers have about user input?

NETFLIX DIVES INTO THE AMAZON[1]

As noted in the introduction, cloud computing is one of the main driving forces in networking today. We introduce cloud computing with an illustrative case. This case does not focus directly on networking. As also noted earlier, cloud computing is radically changing the locations of servers, and is doing so in complex ways that radically change network traffic.

[1] Sources for this section include Brandon Butler, "Three Lessons from Netflix on How to Live in a Cloud," *NetworkWorld*, October 9, 2013, http://www.networkworld.com/news/2013/100913-netflixcloud-274647.html; Matt Petronzio, "Meet the Man Who Keeps Netflix Afloat in the Cloud," *mashable.com*, May 13, 2013, http://mashable.com/2013/05/13/netflix-dream-job/; Kevin Purdy, "How Netflix is Revolutionizing Cloud Computing Just So You Can Watch 'Teen Mom' on Your Phone," *www.itworld.com*, May 10, 2013, http://www.itworld.com/cloud-computing/355844/netflix-revolutionizing-computer-just-serve-you-movies; Ashlee Vance, "Netflix, Reed Hastings Survive Missteps to Join Silicon Valley's Elite," *Business Week*, May 9, 2013, http://www.businessweek.com/articles/2013-05-09/netflix-reed-hastings-survive-missteps-to-join-silicon-valleys-elite.

Netflix

You personally know how individuals use the Internet. The corporate experience is very different. We will illustrate this by looking at Netflix's use of the Internet. Netflix is a streaming video service with over 100 million subscribers around the world. Streaming video needs massive network capacity. A two-hour high-definition movie must deliver 5 million bits each second. This is 9 gigabytes for that one movie. On any given night, Netflix accounts for roughly a third of the Internet traffic going into U.S. homes. When subscriber numbers begin to reach the billions for Netflix and other streaming services, network capacity will have to grow massively.

Requirements Users expect high video quality, and they will not tolerate delay or unreliability. The Internet was not designed for these requirements. The Internet is a "best effort" delivery system that often has insufficient speed and reliability and that often has too much delay for Netflix users. Netflix had to overcome these limitations.

The Internet is a "best effort" delivery system.

Video streaming also requires vast amounts of server processing to deliver video, but the need for heavy server capacity extends well beyond streaming.

- Each movie must be **transcoded** into many streaming formats, and when a customer requests a movie, a streaming server must select the best transcoded format for that customer's equipment, network speed, and other matters.
- In addition, at the heart of Netflix's business plan is a **recommendation system** that creates personalized viewing suggestions for individual customers. This requires the analysis of extensive data about the customer's viewing habits and the choices of other customers with similar viewing profiles.

Server Outsourcing In 2008, when Netflix only delivered movies by sending DVDs through the mail, the company suffered a crippling server outage that stopped shipments for several days. That was a wakeup call for Netflix. Management realized that reliability would be critical for online delivery. It also realized that Internet delivery would become its core business, but managing servers would not. Netflix decided to outsource server operation.

Netflix turned to **Amazon Web Services (AWS)**. Amazon had leveraged its expertise in managing vast server farms for its e-commerce needs into a cloud service that customers like Netflix could use. Figure 11-4 shows that AWS's enormous server farms had the capacity that Netflix needed for customer ordering, transcoding, and analyzing viewing patterns with extreme reliability. In all of these cases, Netflix wrote the applications because these supported their core business.

Content Delivery However, for the job of streaming its more than one petabyte of movie content, Netflix realized that it could not outsource server operations because this was the most central part of its business. Figure 11-5 shows how Netflix delivers video content to individual customers via its **content delivery network (CDN)** *Open Connect*.

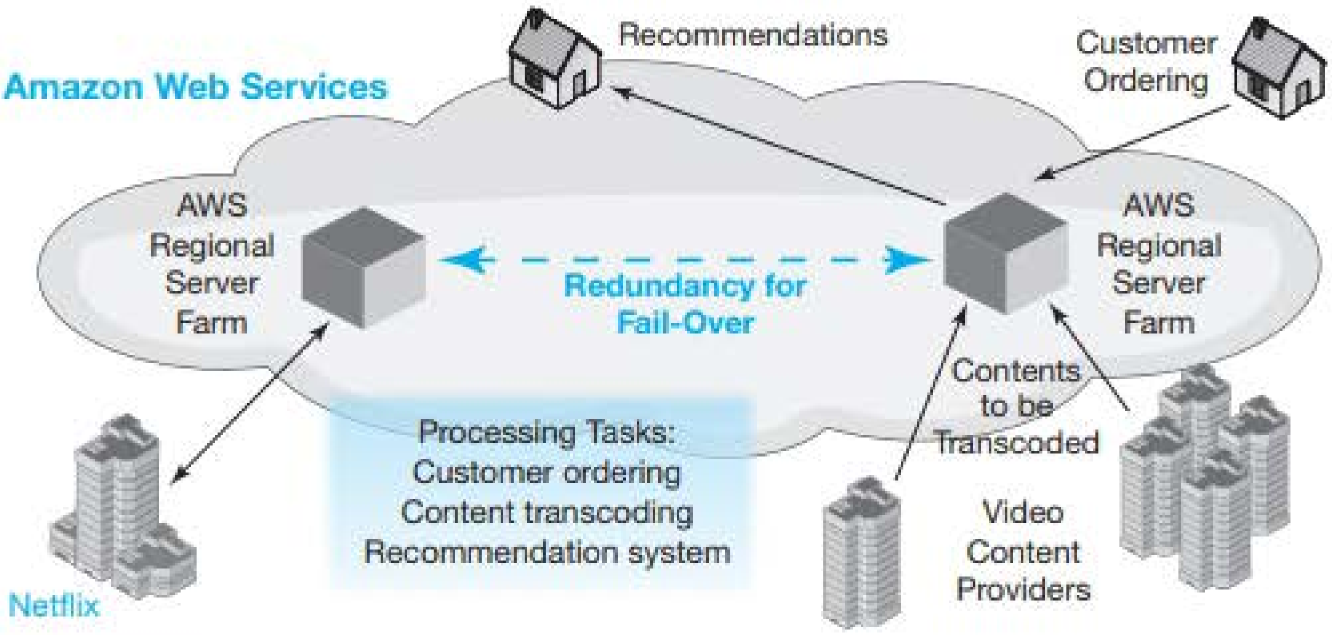

FIGURE 11-4 Netflix and Amazon Web Services (AWS)

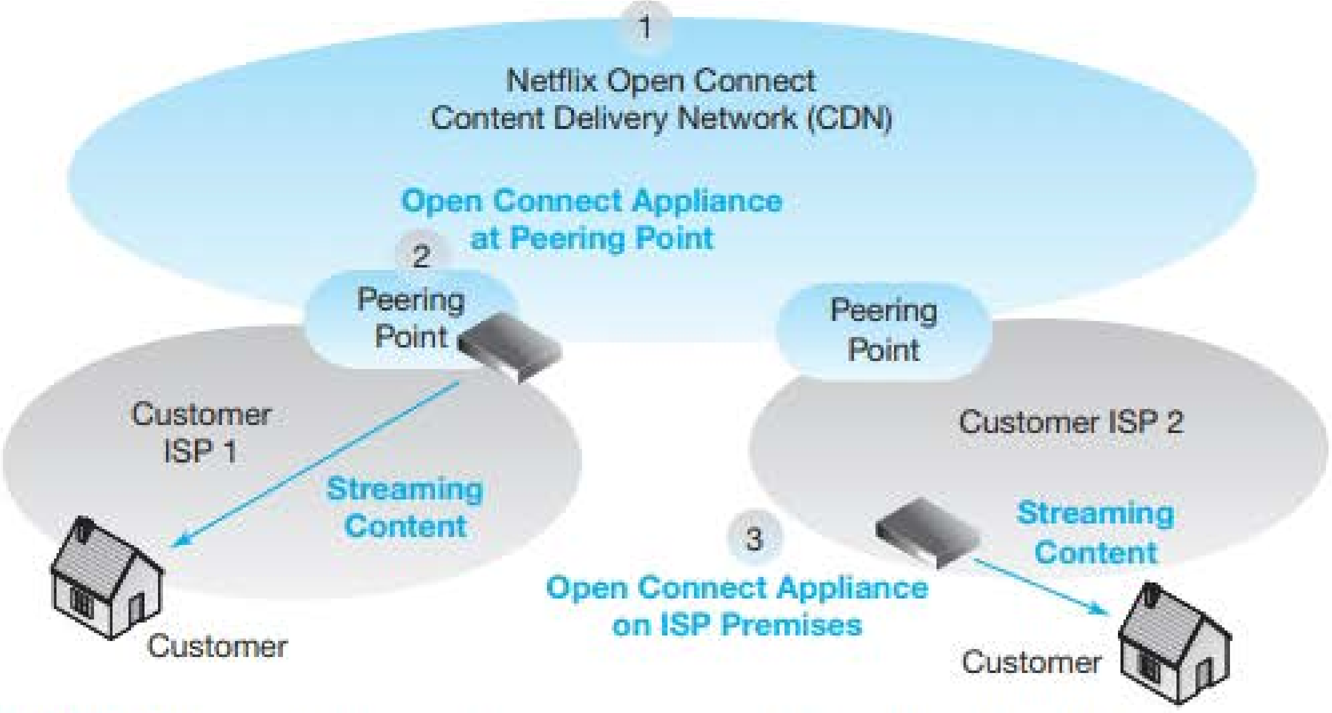

FIGURE 11-5 The Netflix Open Connect Content Delivery Network (CDN)

To stream content, Netflix created its own network on the Internet. Netflix is essentially an ISP, but it is a special one that carries nobody else's content. This Open Connect network is tightly managed by Netflix to ensure very high-quality service.

To deliver streaming content, Netflix created its own webserver appliances. Each is a relatively small box that can fit into a standard equipment rack. These **Open Connect appliances** are about 7 inches (18 cm) high and 2 feet (61 cm) deep. This small powerhouse holds about 100 terabytes of data on 36 hard disk drives. Its microprocessor, in turn, is fast enough to stream movies simultaneously to between 10,000 and 20,000 customers. Netflix updates these CDN servers about once a year with newer hardware to increase their capabilities.

Figure 11-5 shows that the **Open Connect network** works with individual ISPs to deliver content to subscribers. In some cases, Netflix places its Open Connect appliance very near the peering point, where the Open Connect network links to the customer's ISP. This minimizes distance to the customer, thus minimizing delay and reliability risks. Many ISPs let Netscape install Open Connect appliances at locations inside the ISP's network to further reduce distance to the customer.[2]

Test Your Understanding

4. a) Why does Netflix make many transcoded versions of each movie? b) Why is the Netflix recommendation system critical? c) For what applications does Netflix use AWS? d) For what major application does Netflix manage the servers themselves? e) How do content delivery networks reduce streaming delays to customers?

Virtualization and Agility

Figure 11-6 shows that AWS uses virtualization to turn each physical server into several **virtual machines (VMs)**. Each VM is a software program running on the physical server. However, it acts like a real server in its connections with the outside world. It has its own IP address as well as its own data. A VM is even managed like a physical server.

Virtual machines provide **agility**, which is the ability to make changes quickly—even very large changes. For example, Amazon can move VMs quickly from one physical server to another simply by transferring their files. In addition, Netflix can add **VM**

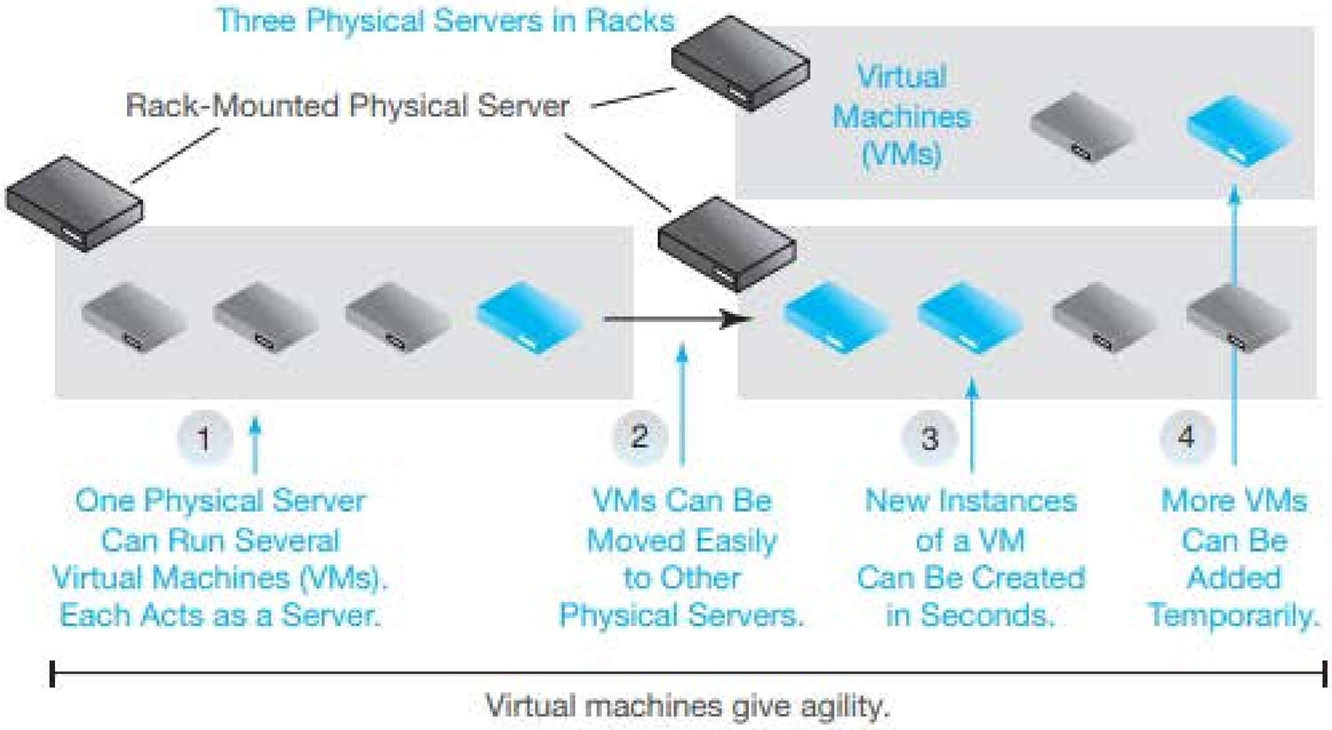
Three Physical Servers in Racks
Rack-Mounted Physical Server
Virtual Machines (VMs)
1
One Physical Server Can Run Several Virtual Machines (VMs). Each Acts as a Server.
2
VMs Can Be Moved Easily to Other Physical Servers.
3
New Instances of a VM Can Be Created in Seconds.
4
More VMs Can Be Added Temporarily.
Virtual machines give agility.

FIGURE 11-6 Physical Servers and Virtual Machines (VMs)

[2] Even with massive storage, Open Connect appliances can only hold a small portion of Netflix's 1 petabyte of content. Consequently, Netflix uses sophisticated analysis to identify the 100 TB of content most likely to be demanded by the customers served by different Open Connect appliances. It installs this content on the individual CDN servers. Of course, customer interests change rapidly, so this content is rebalanced daily. Netflix deletes content declining in popularity and installs content of increasing demand.

instances (specific virtual machines) in seconds. In fact, a company can **spawn (instantiate)** many copies of the same virtual machine at once, in no more time than it takes to spawn a single VM instance. Physical servers offer nothing like this degree of agility. To make virtualization even more attractive to customers such as Netflix, AWS provides a simple self-service application for customers to use to add new instances and to do many other things themselves, in real time.

Transcoding each movie into a hundred or more versions for delivery is an enormous task. Whenever Netflix needs to transcode a movie, it spins up (spawns) multiple VMs, splits the work up among them, processes the data in parallel, and then spins them down.

Providing customized viewing recommendations to subscribers also requires an enormous amount of processing power because it must analyze individual user viewing practices and the viewing practices of people who have viewed similar movies. This recommendation system also requires Netflix to spin up and release large numbers of virtual servers throughout the day as demand increases in the evening hours and declines at other times.

Test Your Understanding

5. a) Distinguish between physical servers and virtual machines. b) What can be done with virtual machines that would be difficult to do with physical servers? c) What is VM instantiation? d) How does Netflix use the agility offered by Amazon Web Services?

Infrastructure as a Service (IaaS) and Software as a Service (SaaS)

We now look in more detail at Amazon Web Services (AWS). Amazon is a **cloud service provider (CSP)**, as Figure 11-7 illustrates. We saw earlier that the Internet and other networks are depicted as clouds. The figure shows that CSPs also operate their services opaquely, forming a second layer of cloud.

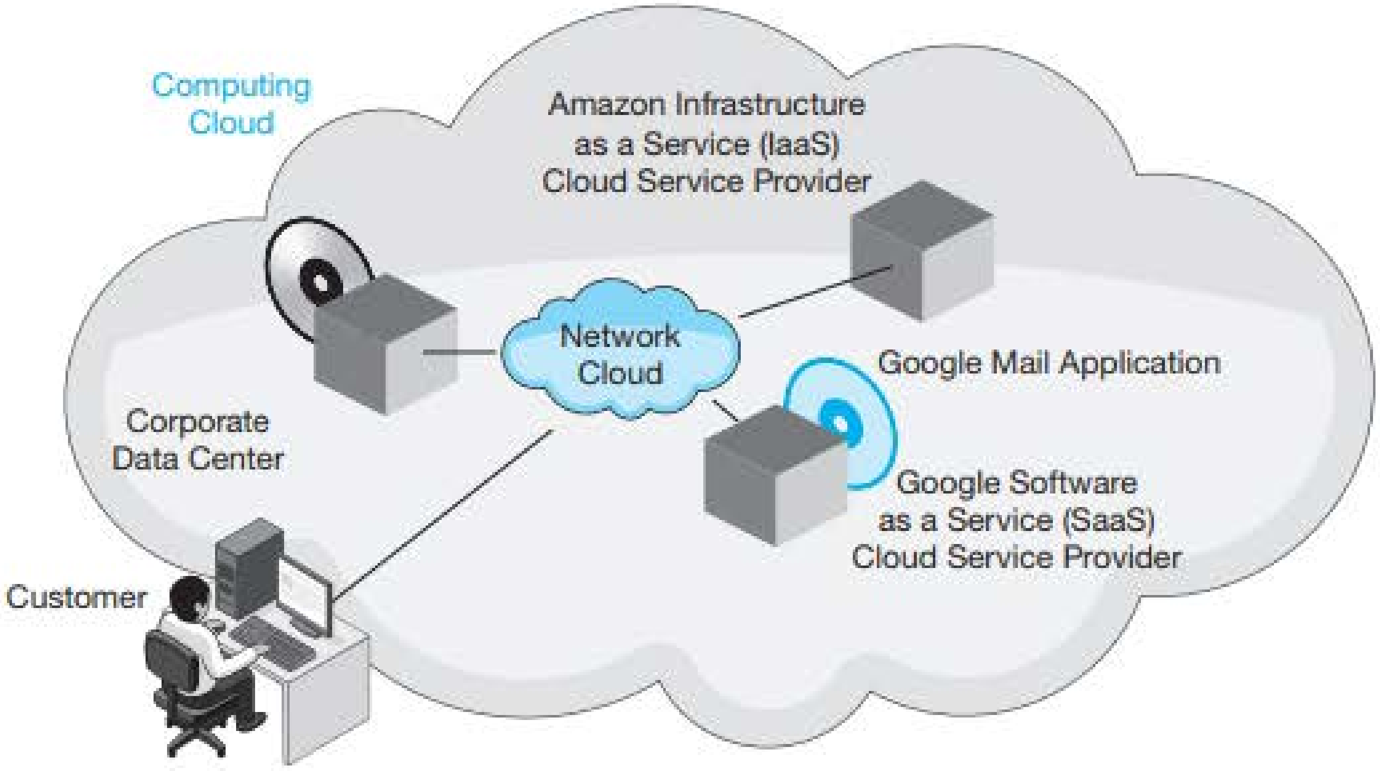
Computing Cloud
Amazon Infrastructure as a Service (IaaS) Cloud Service Provider
Network Cloud
Google Mail Application
Corporate Data Center
Google Software as a Service (SaaS) Cloud Service Provider
Customer

FIGURE 11-7 Cloud Service Providers, IaaS, and SaaS

Infrastructure as a Service The AWS service that Netflix uses is referred to, generically, as **Infrastructure as a Service (IaaS)**. This ungainly name refers to the fact that AWS provides the **computing infrastructure**, which consists of servers and their operation, database management systems, and related services.

Netflix, however, creates and manages its own applications for user ordering, transcoding, personalized viewing suggestions, and other matters. By outsourcing server operation to AWS, Netflix can focus its efforts more fully on developing and extending its applications.

In addition, although Netflix does not manage the servers in AWS, it tests its server/application setups constantly. Netflix has developed a family of programs called the simian army,[3] which it uses to selectively turn off parts of the AWS system to test how well the system responds to outages. When a change is made in an application that runs on many virtual machines, Netflix tries it out on just a few at first, then migrates it to the rest in a smooth manner.

Software as a Service Amazon is not the only cloud service provider that Netflix uses. Another is Google. Netflix uses Google Mail for its internal communication. In contrast to just offering IaaS, Google offers application software as well. This is called **Software as a Service (SaaS)**. Here, *software* refers to application software. SaaS has been popular for many years. For example, many companies use salesforce .com application software for sales force management and customer relationship management.

"As a Product" versus "As a Service" As a Service in IaaS and SaaS refers to pricing. Normally, a company buys servers like other products, such as automobiles and apples. After purchase, the company owns the physical server and manages it.

In contrast, cloud services are sold like electrical service. You pay for the amount you use. This allows customers to avoid the capital expense (CapEx) of purchasing servers. This also avoids the risk of buying too much capacity that would go unused. IaaS appears as an operating expense (OpEx), which can be managed so that money is spent only when it must be. SaaS, in turn, changes application programs from purchased products to per-use services.

Test Your Understanding

6. a) What is a CSP? (Do not just spell out the acronym.) b) Distinguish between IaaS CSPs and SaaS CSPs. c) Is AWS an IaaS provider or an SaaS provider for Netflix? d) Is Google an IaaS provider or an SaaS provider for Netflix? e) Who owns and manages the servers in IaaS and SaaS? f) Who owns and manages the applications in IaaS and SaaS? g) In AWS, what does Netflix manage and not manage? h) For e-mail, what does Netflix manage and not manage?

[3] This name reflects the fact that individual programs have names such as Chaos Monkey and Chaos Gorilla.

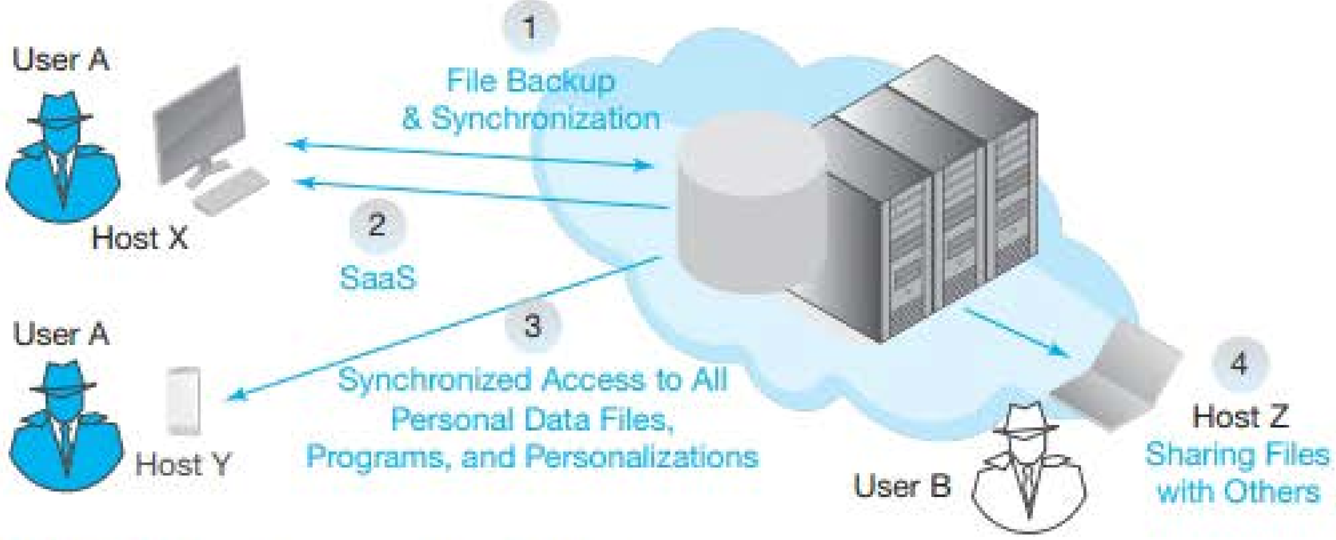

FIGURE 11-8 Client Computing in the Cloud

Clients Move into the Cloud

Netflix embodies how corporations use the cloud. However, as Figure 11-8 shows, many client hosts also use the cloud. Most users today have multiple devices. They want to work on a document on one device, move to another, and pick up exactly where they left off. As Figure 11-8 shows, this requires storing work data and synchronization data in the cloud. As users move between machines, their data is immediately available, and synchronization takes them immediately to the exact spot in their working document they last touched on the previous machine.

Often, the storage content of users' cloud services is offered by itself, through **file storage services** such as Dropbox and iCloud. (Individual users rarely back up their stand-alone devices.) Even if the user accidentally deletes a file on both the client's and the cloud storage server, it is usually possible to retrieve an earlier version.

The fact that data is stored in the cloud also facilities sharing. A user can control who can see a particular folder and what they may do on the files stored there. In addition, some folders "on the user's computer" may be shared folders under a group account with other users.

Often, instead of buying application software and installing it on each machine, the user pays an annual fee that will offer application software on all of his or her machines. This is another case of SaaS. Ideally, there would be no need to download the software, but for large applications such as word processing, a completely cloud-based service would be too slow. Each machine downloads some or all of the application. However, the software vendor frequently updates the files on each computer to the newest version of the software and synchronizes configuration changes and other personalizations.

Test Your Understanding

7. a) With cloud services for clients, what happens when a user moves from one physical client device to another? b) What protections are offered by file storage services? c) In a file storage service, what can you do if you accidentally delete a file on a client and the corresponding file on the server is also deleted? d) In file sharing services, can the user allow others to share some files? e) In SaaS, why is the program or part of the program stored on the client device?

Risk to Corporate Data

- If the cloud service provider fails to protect data, the results can be disastrous
- Customer firm has no control over cloud service provider security

Due Diligence is Necessary

- Must examine cloud service provider protections before using them
- Many companies fail to do this, enticed by low costs and agility

FIGURE 11-9 Cloud Security Concerns (Study Figure)

Rain Clouds: Security

Security must be a concern for every cloud customer. Companies must put critical corporate data on computers owned by other organizations. (In the case of Netflix, Amazon is actually a competitor in the streaming media market.) If cloud service providers fail to protect this data from hackers, the potential damage is enormous.

To deal with security, companies must do extensive due diligence, looking in depth at how cloud service providers handle security. However, there is no way to understand everything about a cloud service provider's security. For the time being, many organizations are crossing their fingers, whistling in the dark, knocking on wood, and yielding to the attraction of cloud computing's low cost and agility.

Test Your Understanding

8. What concerns do customers have about cloud security?

Networks and The Cloud

Networks today must work extremely well, almost perfectly. They must do this while growing at unbelievable rates. And they must do this using standards older than most of today's network engineers.

The demands of cloud computing create enormous stresses on networks. Cloud service providers themselves create massive and fast-changing network transmission

Cloud Computing Stresses Networks

- Traffic is massive and quickly changing
- Massive sudden changes in Internet and local network traffic
- Reliability is critical

Latency Concerns

- Cloud computing may increase latency
- Very low latency is critical for many applications
- Important examples are speech input services (such as Echo) that rely on cloud-based AI to understand user input

FIGURE 11-10 Networks and the Cloud (Study Figure)

loads. Customers of cloud services also find themselves with massive increases in Internet and local network traffic. In addition to growing rapidly, networks are also facing increasing demands for reliability because a company that loses contact with its cloud service providers for even brief periods of time will suffer heavy losses.

Latency is also critical for a growing number of core services. For instance, speech input systems such as Amazon Echo, Cortana, Siri, and Google use cloud servers to do the heavy artificial intelligence (AI) processing needed to understand the user's voice and meaning. For use to be operational, the two-way transmission as well as the server processing must be done in real time, with no appreciable delay.

Test Your Understanding

9. a) How is cloud computing affecting networking? b) Why is latency a problem for artificial intelligence?

THE WORLD WIDE WEB

HTTP and HTML Standards

Having looked at core networked application concepts and cloud computing, we turn to a series of key applications. Given its dominance, we discuss the World Wide Web first. Figure 11-11 shows that the Web is based on two primary standards.

- For **file format standards**, webpages themselves are created using the Hypertext Markup Language (HTML). Once downloaded, tags in the HTML document are used to download related files.
- Second, the transfer of requests and responses uses the Hypertext Transfer Protocol (HTTP) to specify files to be retrieved and to describe file types for delivered files (HTML, JPEG, etc.).

Test Your Understanding

10. a) What are the two major sets of standards for the World Wide Web? b) How do they differ?

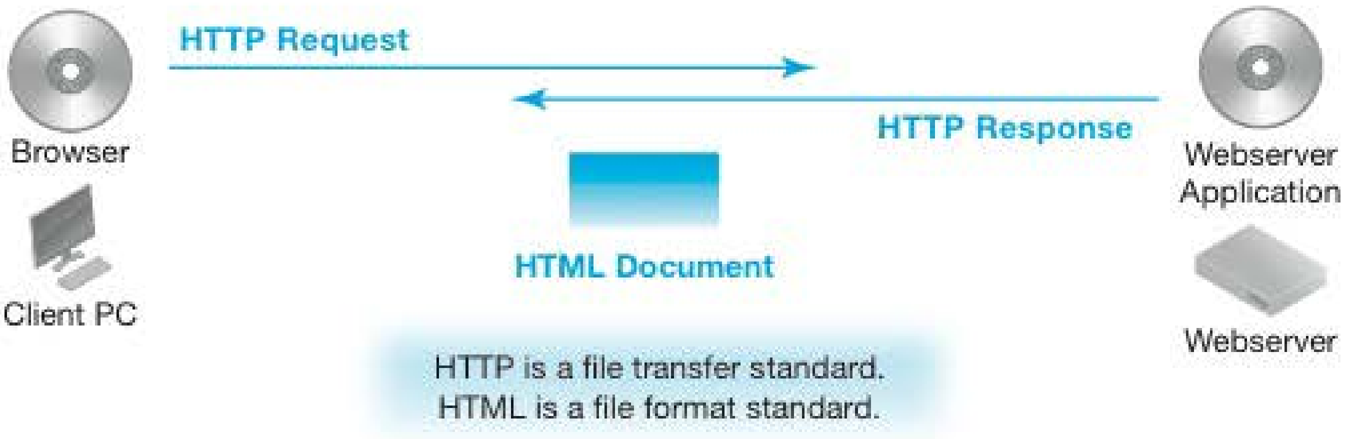

FIGURE 11-11 World Wide Web (WWW) Standards

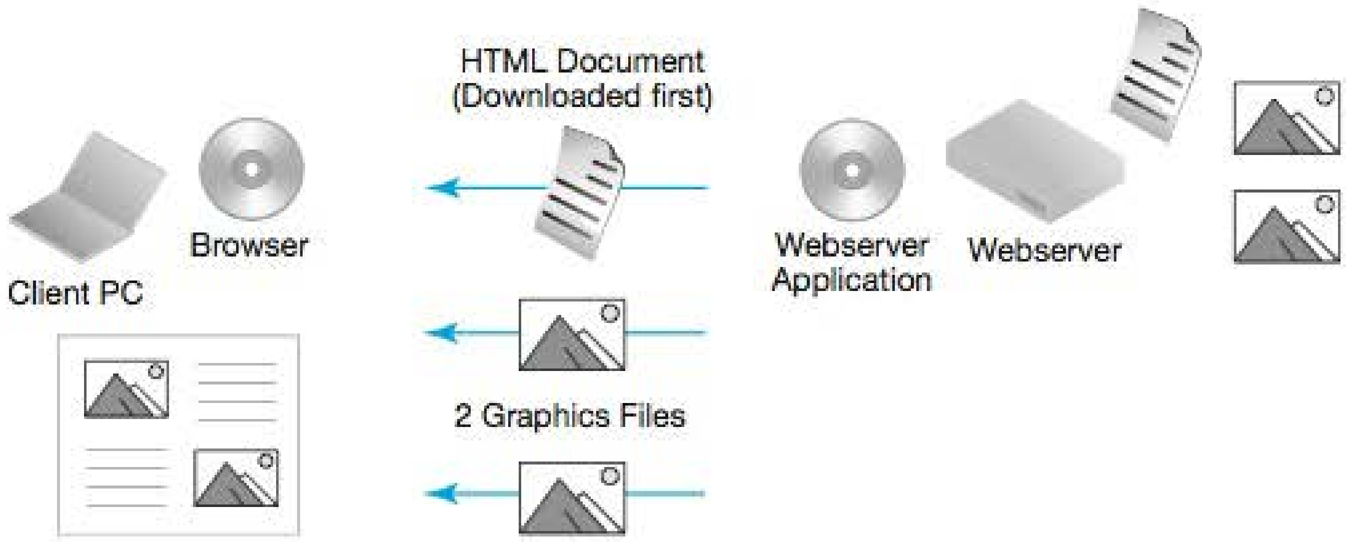

FIGURE 11-12 Downloading a Webpage with Two Graphics Files

Complex Webpages

Nearly all "webpages" really consist of several files—a master text-only HTML file plus graphics files, audio files, and other types of files. Figure 11-12 illustrates the downloading of a webpage with two graphics files.

The HTML file consists merely of the page's text, plus **tags** to show where the browser should render graphics files, when it should play audio files, and so forth.[4] The HTML file is downloaded first because the browser needs the tags to know what other files should be downloaded.

Consequently, several HTTP request–response cycles may be needed to download a single webpage. Three request–response cycles are needed in the example shown in the figure.

To provide an analogy, when you download an e-mail message with attachments, you must read the message first. Then you must click on the attachments to download them.

Test Your Understanding

11. a) You are downloading a webpage that has six graphics and two sound clips. How many request–response cycles will be needed? b) Which file will be downloaded first?

The Hypertext Transfer Protocol (HTTP)

HTTP, again, standardizes interactions between the browser and the webserver to ask for and deliver files.

[4] For graphics files, the IMG tag is used. The keyword *IMG* indicates that an image file is to be downloaded. The SRC (source) parameter in this tag gives the target file's directory and file name on the webserver. If the HTML document was not downloaded first, the browser would have no tags to determine what other files to download.

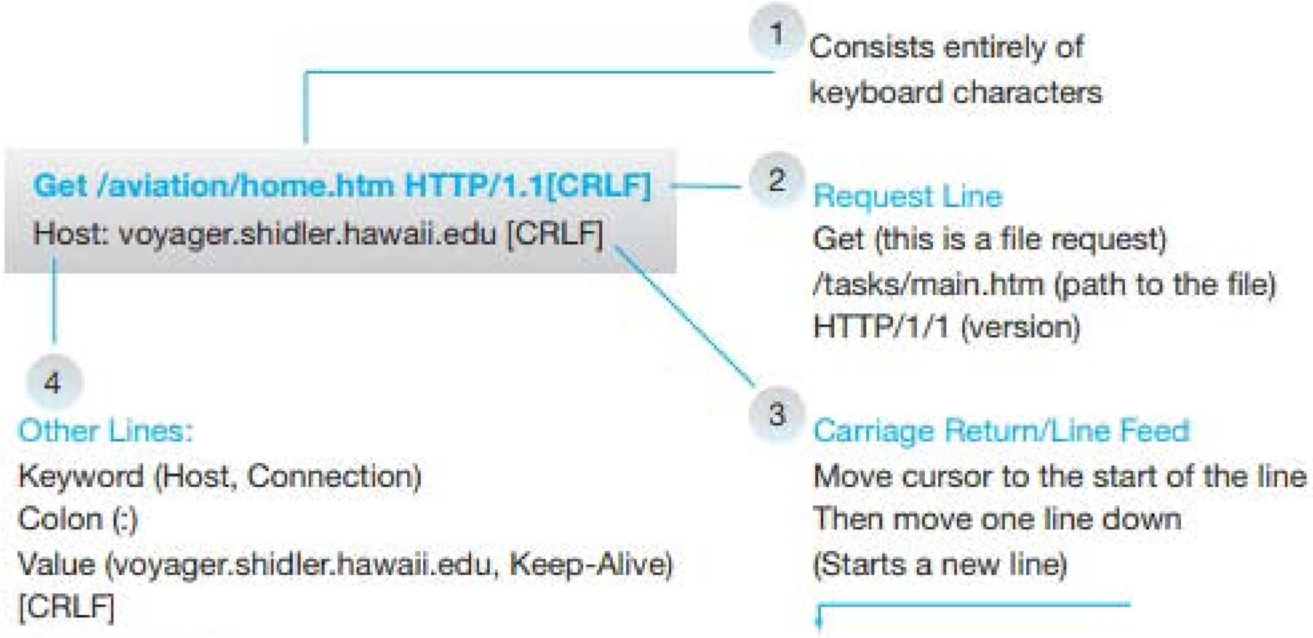
1
Consists entirely of keyboard characters
Get /aviation/home.htm HTTP/1.1[CRLF]
Host: voyager.shidler.hawaii.edu [CRLF]
2
Request Line
Get (this is a file request)
/tasks/main.htm (path to the file)
HTTP/1/1 (version)
3
Carriage Return/Line Feed
Move cursor to the start of the line
Then move one line down
(Starts a new line)
4
Other Lines:
Keyword (Host, Connection)
Colon (:)
Value (voyager.shidler.hawaii.edu, Keep-Alive)
[CRLF]

FIGURE 11-13 HTTP Request Message

HTTP Request Messages Figure 11-13 shows the syntax of an **HTTP request message**. As we saw in Chapter 2, most other content transmission standards consist of bit strings that are not designed for people to read. In contrast, HTTP messages consist entirely of alphanumeric symbols that can be typed on a keyboard. The characters are encoded into bit streams before being passed to the transport layer, but before that they are clearly readable. They are also readable in plain text on the receiver.

In fact, they look like traditional e-mail messages. Most lines begin with a keyword, followed by a colon, then a value, and finally a carriage return/line feed (CRLF). **Carriage return** takes the cursor back to the start of the current line. **Line feed** moves the cursor one character down. This combination starts a new line. On old typewriters, the carriage return handle on the left side of the machine combined these actions.

Carriage return and line feed (CRLF) together start a new line.

The first line in the header has a different format than other header lines. This is the *Request Line*.

- It begins with a method to indicate what the sender wishes to be done. This is usually *Get*, indicating that the requestor wishes to get a file.
- This is followed by the location of the file. This tells the receiver to begin at its web root directory, go one directory down to tasks, and retrieve the file main.htm.
- Finally, the request line tells the webserver program that the client is the 1/1 version of HTTP.

HTTP Response Message Figure 11-14 shows an HTTP response message. This is more complex than the HTTP request message, but it has almost the same basic structure. It starts with a *status line*, followed by multiple lines in the *Keyword: value[CRLF] format*.

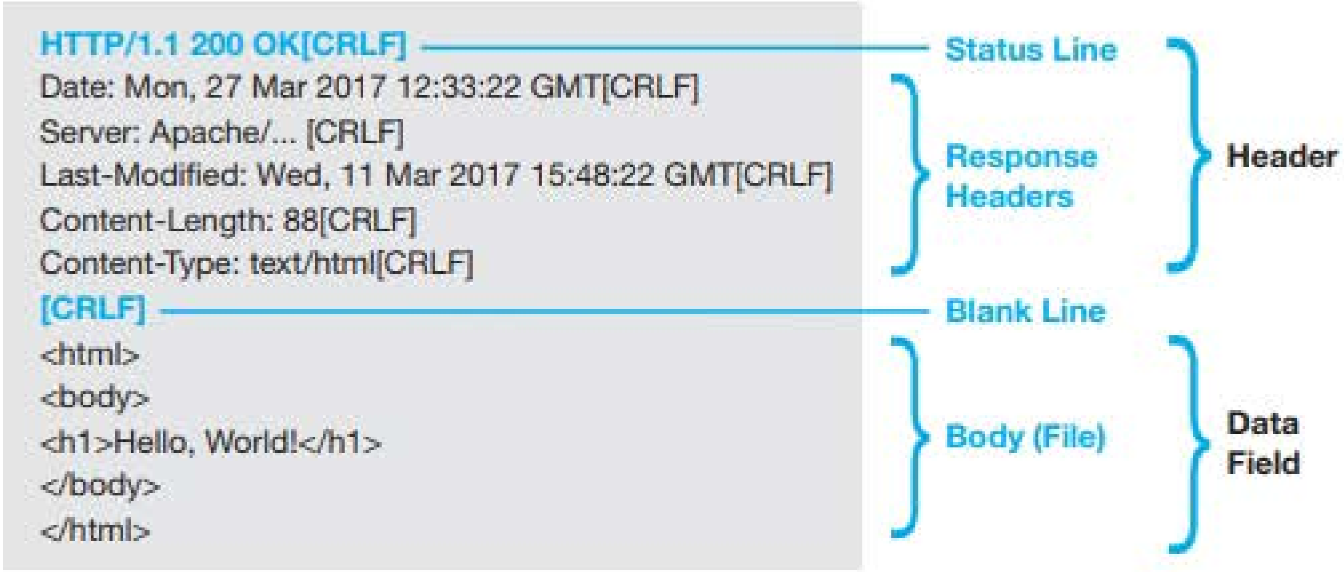

```
HTTP/1.1 200 OK[CRLF]
Date: Mon, 27 Mar 2017 12:33:22 GMT[CRLF]
Server: Apache/... [CRLF]
Last-Modified: Wed, 11 Mar 2017 15:48:22 GMT[CRLF]
Content-Length: 88[CRLF]
Content-Type: text/html[CRLF]
[CRLF]
<html>
<body>
<h1>Hello, World!</h1>
</body>
</html>
```

FIGURE 11-14 HTTP Response Message

The status line alerts the receiver to how the server has responded.

- It begins with the *HTTP version* the webserver will use to talk to the browser. This is HTTP 1/1, the same version the browser signaled.
- It then gives a *status code* to indicate how it has responded to the request. The code is 200, which indicates that the request has been accepted and executed. There are many status codes, such as the famous 404 status code: Page Not Found.
- The status line continues with a *reason phrase*, which is an expression to help a user understand what the status code means. In this case, the reason code is simply "OK."[5]

Following the status line are HTML **response headers**. These give the date the response message was sent, the operating system of the webserver (. . . indicates missing content), the date and time the file was last modified, the data field length in bytes, and the type of data in the data field.

For example, the *Content-Length field* gives the length of the data field in bytes. Next, the *Content-Type field* indicates that the data field consists of HTML text. (HTML documents consist entirely of keyboard characters.)

Next comes a line with a single CRLF. This is a blank line. It separates the header from the data field. This is a crude separator, but it works.

Finally comes the data field. This is the HTML file the response message is delivering.

Of course, HTTP does not only deliver HTML files. What if this response message was delivering a jpeg graphic file? In that case, the Content-Type field would say image/jpeg, and the Content-Length field would give the size of the jpeg image in bytes.

[5] HTTP is designed to be humanly readable.

Test Your Understanding

12. a) Which tends to have a data field, HTTP request or response messages? b) What is the first line called in an HTTP request message? c) What is the first line called in an HTTP response message? d) What is the format for header lines after the first line? e) What would the Content-Type field be for an HTTP response message delivering a GIF graphics file?

ELECTRONIC MAIL (E-MAIL)

We now turn to electronic mail (e-mail), which was one of the earliest applications on wide area networks, and it is still growing rapidly today.

Delivery Standards

As in the case of the World Wide Web, e-mail uses two sets of standards—one for message delivery and one for file formats. We begin with delivery standards.

The Delivery Process Figure 11-15 illustrates how e-mail messages are delivered to receivers.

- Most fundamentally, the sender does not send the message directly to the receiver. Instead, each party has a mail host. When the sender transmits a message, it sends it to its own mail host (1).
- The sender's mail host notes the destination e-mail address and looks up the IP address of the receiver's mail host.
- The sender's mail host then sends the message to the receiver's mail host (2).
- The receiver's mail host holds the e-mail until the receiver downloads it (3).

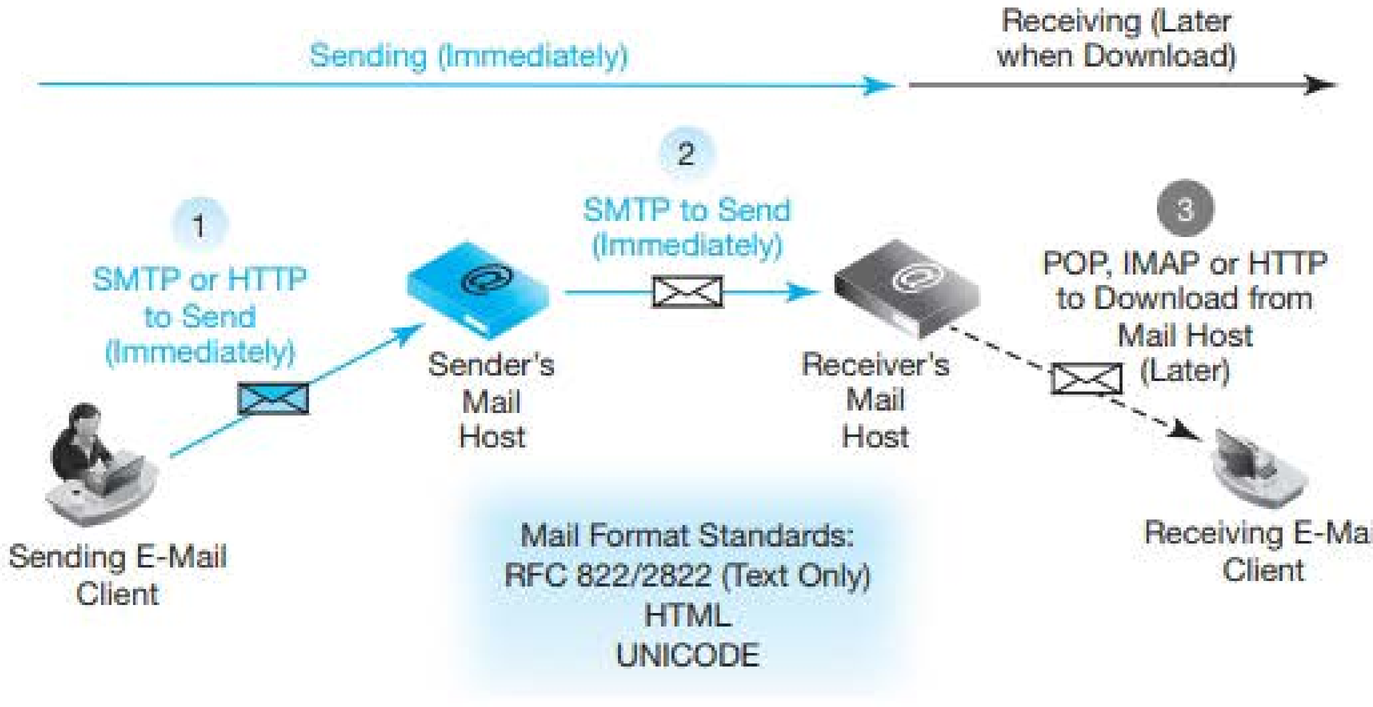

FIGURE 11-15 E-Mail Delivery Standards

Having intermediate mail hosts might seem cumbersome, but people do not read the mail immediately. Mail hosts have mail boxes for each user to store mail until the user checks for them. This arrangement allows e-mail users to pick up their mail whenever they feel like it.[6]

Transmission to the Sender's Mail Host Using the Simple Mail Transfer Protocol (SMTP) When the sender transmits the mail to his or her mail host, the traditional transmission standard was the **Simple Mail Transfer Protocol (SMTP)**. When you set up an e-mail account on your smartphone or other device, you may be asked for the host name of your SMTP host—the host to which you will send mail.

Web Mail and HTTP People increasingly use Web mail, which allows you to send mail through your browser. In this case, most communication uses HTTP. This includes sending mail.

SMTP between Mail Hosts However, the sender communicates with his or her mail host, and mail hosts communicate via SMTP.

Immediate Delivery When you hit *send*, your mail is uploaded to your mail host immediately. Your mail host does an IP address lookup and then immediately sends the mail to the user's mail host. The delay from the time you hit send to the time the receiver's mail host gets the message is rarely more than a second or two.

Test Your Understanding

13. a) In traditional e-mail, when a client sends a message to its mail host, what standard does it use? b) Which standard is used for this in Web-based e-mail? c) When the sender's mail host sends the message to the receiver's mail host, what standard does it use? d) What do you think are the advantages of a Web-enabled e-mail system? (The answer is not explicitly in the text.)

Receiving Standards

Sometime after mail is delivered to the receiver's mail host, the receiver will retrieve it from his or her mail box on the receiver's mail host. Receiving is a more complex process than sending because users want a great deal of flexibility in how they read their mail. Therefore, receiving standards are different from sending standards. Figure 11-15 notes that the two most common traditional receiving standards are the **Post Office Protocol (POP)** and the more sophisticated **Internet Message Access Protocol (IMAP)**. When you set up your e-mail on a device, you may be asked for the host name of your

[6] Also, locating the receiver's mail host is easier than locating an individual receiver. A mail user's IP address is likely to change every time he or she boots up. In contrast, mail servers have static IP addresses. Suppose that you are sending to ray@panko.com. When your mail server sees ray@panko.com, it realizes that panko .com is the domain name. So it does a DNS lookup on panko.com. Instead of asking for the A or AAAA record, which would return the IP address of the panko.com webserver, it asks for the MX record for panko. com. This returns the IP address of panko.com. The sender's mail server then sends the mail to the receiver's mail server.

POP or IMAP host. In some cases, your POP or IMAP server can be different from the name of your SMTP host. Of course, in Web-based e-mail, HTTP is used for downloading as well as sending.

Test Your Understanding

14. a) In traditional e-mail, when the receiver's e-mail client downloads new mail from its mail host, what standard is it likely to use? b) What standard is used for downloading e-mail in Web-based e-mail? c) Why is there usually a time difference in transmission from the sending client to the receiver's mail host and the time when the message is downloaded?

E-Mail File Format Standards

The World Wide Web, like many applications, has two sets of standards. HTTP governs message delivery. We have seen that e-mail, in contrast, has several standards for transmitting messages, including SMTP, POP, IMAP, and HTTP. HTML governs the file format of the Web's main file type, which is always the first file downloaded in webservice. E-mail also has several standards for file formats.

ASCII and Searchable Header Fields The earliest e-mail messages were limited to the characters you can type on a standard American keyboard. They were called **text standards**, although they also include digits, punctuation marks, and various other characters. This was unexciting visually, but it placed little burden on displays and transmission lines.

These early text standards did bring one major facility, **searchable fields**. The header consisted of several fields of the keyword–colon–context format. These include To:, From:, Date:, and other fields. This gives structure to e-mail messages. It allows us to display e-mails with the most recent ones first (which is the norm), by sender, and by other fields. We also can do searches by character strings in specific fields. Without this, e-mail would be far less useful.

Graphics in E-Mail Messages Two developments created today's e-mail file formats.

- One is attachments, which allow even text messages to deliver files in the formats of specific applications such as Microsoft Excel.
- The other is the gradual and growing addition of graphics into the document body. Many e-mail programs already show graphics files (png, jpg, and gif) in the body of messages, and some go much further. The ability to have HTML bodies has brought extremely rich content to e-mail, even when Web mail is not being used. With Web mail, of course, headers and even the body can be as rich as desired.

UNICODE Another trend in e-mail headers and bodies is the growing support for non-English characters. Originally, the searchable header and body text was limited to characters from the **American Standard Code for Information Interchange (ASCII)**. ASCII cannot represent diacritical marks such as German umlauts (except with awkward extensions). Nor can it represent Japanese, Sanskrit, Cyrillic, or the world's other languages with entirely different symbol systems. To address these limitations, most mail systems now support bodies in **UNICODE**, which can represent nearly all language symbol systems.

The use of UNICODE is good, but it creates problems for message filtering to identify spam and phishing attacks, cross-site scripting attacks, and several other common e-mail attacks. Searching for string patterns that are the signatures of attacks becomes extremely challenging because different languages have very different codes for some grammatical marks such as slashes.

Test Your Understanding

15. a) Text messages are limiting, but they introduced an important innovation. What was it? b) How can e-mail deliver content suitable for specific applications, such as word processing programs? c) What is the state of graphics in e-mail today? d) Why is UNICODE good? e) What security issue does it create?

Cryptographic E-Mail Protections

Given e-mail's importance and potential for security failures, one might assume that encryption and other cryptographic protections are used almost all the time. In fact, they are not. In addition, even when cryptographic protections are used, they may be very limited.

Link Encryption Figure 11-16 shows how encryption is done in e-mail. The top part of the figure shows **link encryption**, which protects a message over a single hop between devices. All links must be encrypted to give comprehensive encryption for e-mail messages.

- When you transmit messages to your e-mail host, you use either SMTP or HTML. Both can protect your transmission with SSL/TLS. However, Figure 11-16 shows that this only protects transmission to your mail host.

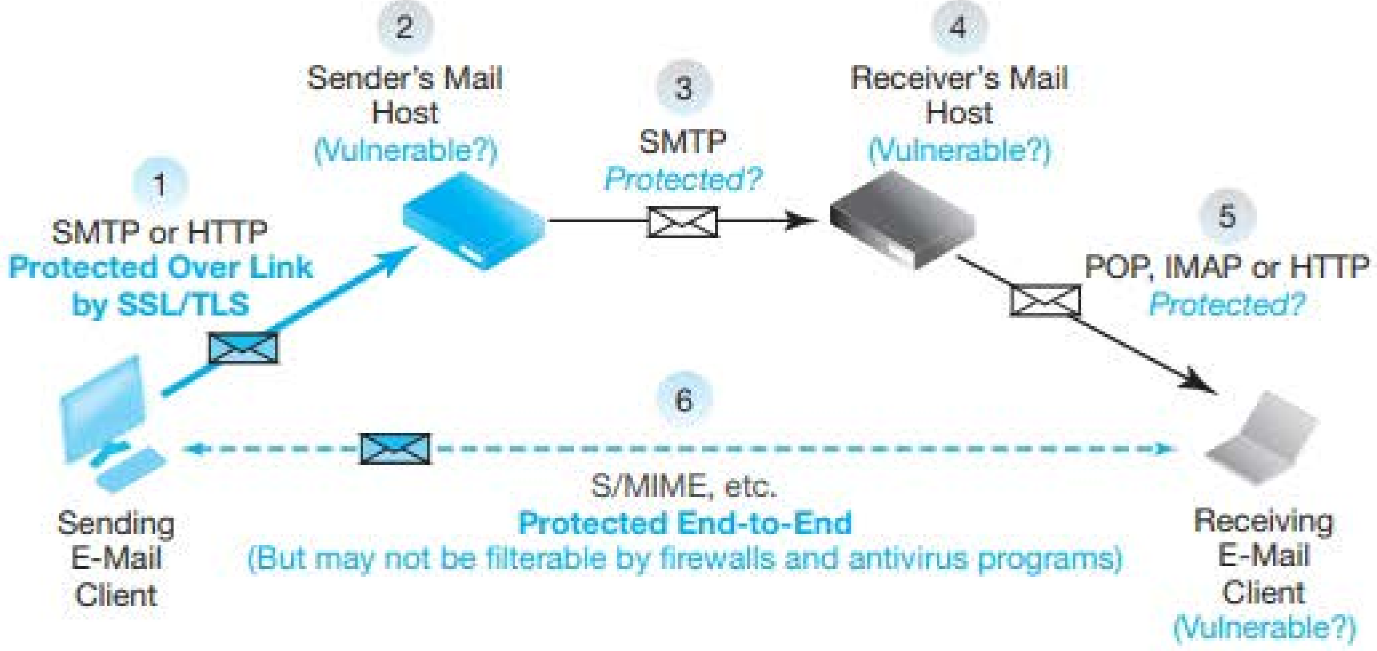

FIGURE 11-16 Cryptographic Protections for E-Mail

- Are the transmissions between the two mail hosts protected? The answer is, "Maybe." To use SSL/TLS for communication between mail hosts, both hosts must agree to do so. Today, data from Google indicates that this is done nearly all the time, but there are still a few percent of mail hosts that fail to accept SSL/TLS connections. If there is no encryption between mail hosts, this is a vulnerability.
- Finally, when the receiver downloads the message, do they use SSL/TLS protection? Again, the answer is, "Maybe." Also again, although the use of SSL/TLS is very common, it is not universal. This final link is another potential failure point for encryption protection as a message travels over the Internet.

End-to-End Encryption The bottom part of Figure 11-16 shows **end-to-end encryption**, in which the sender encrypts a message and the receiver decrypts it. This ensures that the message is encrypted throughout its journey. Unfortunately, there are multiple standards for end-to-end e-mail encryption, and most of these standards, including the popular **S/MIME protocol**, require both parties to have digital certificates. (In contrast, SSL/TLS only requires the mail host to have a digital certificate.) Companies need to employ corporate-wide digital certificates to use encrypted end-to-end e-mail. Transmissions between organizations require both companies to do so, and their certificates must be acceptable to each other.

Link Encryption

- Between the sending client and its mail host
- Between the two mail hosts
- Between the receiver's mail host and the receiver
- All links must be protected for fully encrypted communication

End-to-End Encryption

- Between the two clients
- Requires choosing the same encryption method
- Usually requires digital certificates for both parties
- Firewalls and antivirus programs cannot filer content unless the same extended encryption method is used

Encryption on Mail Hosts and Clients

- E-mail stored on the client's and mail hosts must be encrypted
- The hosts must be hardened with good security to prevent decryption of encrypted files
- Social engineering can bypass these protections

Internal Corporate Communication

- E-mail security is possible for all transmissions internally in a corporation
- A strong standardized set of protections can be enforced and enabled
- This is not possible for general communication over the Internet

FIGURE 11-17 Issues in Corporate E-Mail Protection (Study Figure)

In addition, although end-to-end encryption enhances confidentiality, integrity, and authentication, it lowers security by making it impossible for firewalls along the way to read packets or antivirus servers along the way to scan for malware. Sometimes it is possible to provide keys to such devices to decrypt, process, and then reencrypt messages. However, this raises security issues. In addition, until all parties use e-mail programs modified to provide keys for temporary decryption and use the same algorithms to do so, this process will not work beyond single firms.

File Encryption on Mail Hosts and Clients Of course, end-to-end transmission security means nothing unless the four hosts are also secure. If an attacker can compromise a client or an e-mail host, the attacker will be able to read all messages on the host.[7] Mail clients and the mail host should encrypt all mail in their protection and provide broader host security protections to prevent takeover, which might lead to being able to decrypt encrypted messages. A social engineering attack, furthermore, can defeat the strongest technical protections.

Internal versus External Transmission Creating strong and effective e-mail encryption is not feasible for general e-mail. However, inside individual corporations, strong security policies and implementation can make this possible. Corporations can even standardize on e-mail clients with the built-in security functionality that corporations need for handling e-mail with protection and that can be governed by corporate e-mail policy server requirements. Corporate communication that uses the Internet is another matter.

Test Your Understanding

16. a) If a message sender uses SSL/TLS when it sends a message, how is protection likely to be limited? b) Distinguish between link encryption and end-to-end encryption for confidentiality. c) Why is link-by-link encryption for confidentiality not fully secure even if there is encryption for confidentiality in all links along the way? d) What is the remedy for the limitations of link-by-link encryption? e) Why is end-to-end encryption uncommon?

VOICE OVER IP (VoIP)

Voice over IP (VoIP) has traditionally meant sending digitized voice data in IP packets. The use of IP is important because it means that telephony can share a company's IP data network. This can slash the cost of long-distance telephone service among a company's sites by taking advantage of economies of scale in networking. In addition, VoIP compresses the voice signal, allowing it to consume relatively little IP capacity. Today, anyone who uses Skype knows that VoIP can also stand for **Video over IP**.

VoIP (Voice over IP and Video over IP) is the transmission of voice and video information over IP networks. It permits a company to slash voice and video transmission costs.

[7] There is also a very short period between when an e-mail host decrypts an incoming message and reencrypts it for outgoing transmission. Owning the mail host may allow this to be exploited.

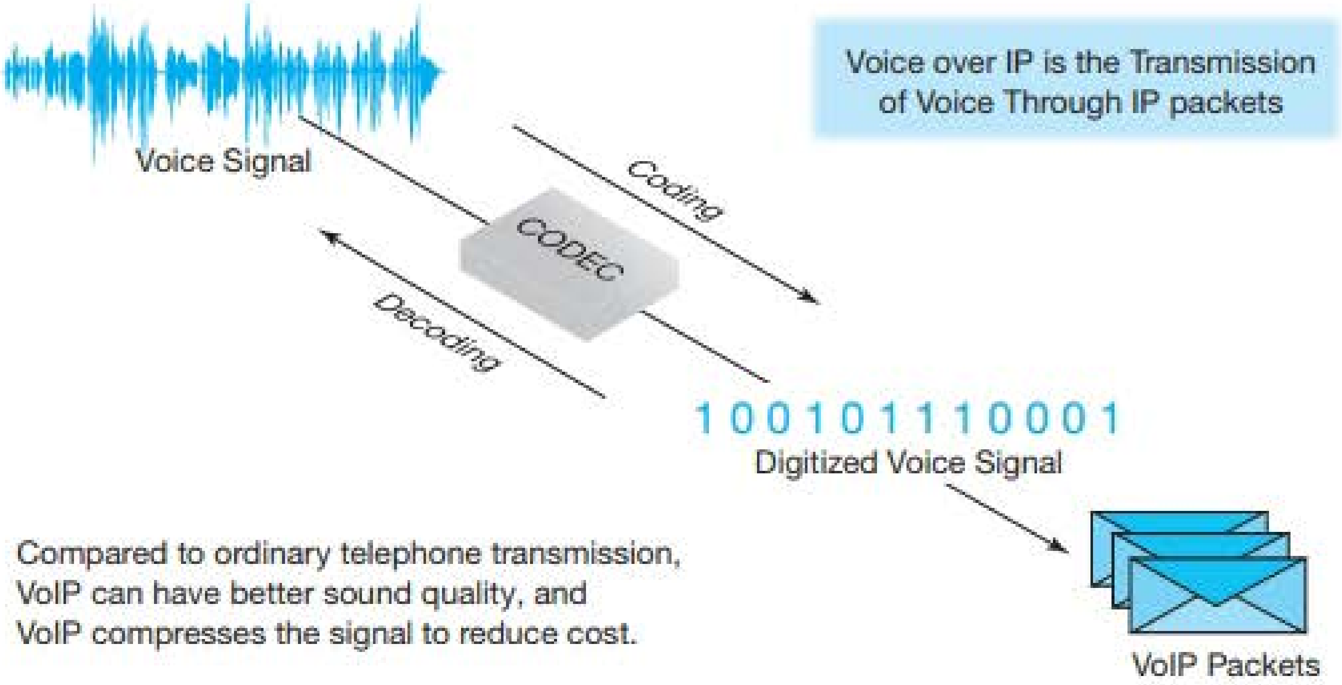
Voice over IP is the Transmission of Voice Through IP packets
Voice Signal
Coding
CODEC
Decoding
1 0 0 1 0 1 1 1 0 0 0 1
Digitized Voice Signal
Compared to ordinary telephone transmission, VoIP can have better sound quality, and VoIP compresses the signal to reduce cost.
VoIP Packets

FIGURE 11-18 VoIP Transmission Using CODECs to Digitize Voice Signals

Test Your Understanding

17. For what two things is VoIP an acronym?

CODEC

The human voice rises and falls in amplitude thousands of times per second. These rises and falls appear to be sudden, but at the detailed level, these changes are continuous rises and falls in intensity. These voice signals must be sampled and encoded into 1s and 0s to be transmitted over a network. At the other end, they must be decoded back into voice signals. The circuit that provides these two functions is called a CODEC (Figure 11-18).

Compared to ordinary telephony, digital transmission can provide higher voice quality than traditional voice telephony. However, to achieve voice quality equal to that of the telephone system, encoding must generate 64 kbps of digital traffic. As Figure 11-19 shows, however, most CODEC standards do more compression, trading off voice quality against transmission costs by sending fewer bits.

Codec Standard	Bits Transmitted per Second
G.711	64 kbps
G.722	48, 56, or 64 kbps
G.721	32 kbps
G.722.1	24, 32 kbps
G.726	16, 24, 32, 40 kbps
G.728	16 kbps
G.729AB	8 kbps
G.723	5.33 6.4 kbps
G.7231A	5.3 6.3 kbps

FIGURE 11-19 CODEC Encoding Standards

Test Your Understanding

18. a) What two things do CODECs do? b) What is reduced when CODEC transmission is compressed?

External Components

Figure 11-20 shows the three external components in VoIP. First, there are the VoIP client devices.

- Businesses typically use dedicated VoIP telephones, which contain CODECs and TCP/IP networking functionality.
- Residential users and an increasing number of business employees now use PCs with a software CODEC and TCP/IP functionality. Using a PC is especially desirable for videoconferencing.

To connect the VoIP system to the public switched telephone network (PSTN), a device called a **media gateway** handles the translation between digital and voice communication.

Test Your Understanding

19. a) What are the two options for VoIP clients? b) What are the functions of media gateways?

VoIP Signaling

In telecommunications, there is a fundamental distinction between signaling and transport.

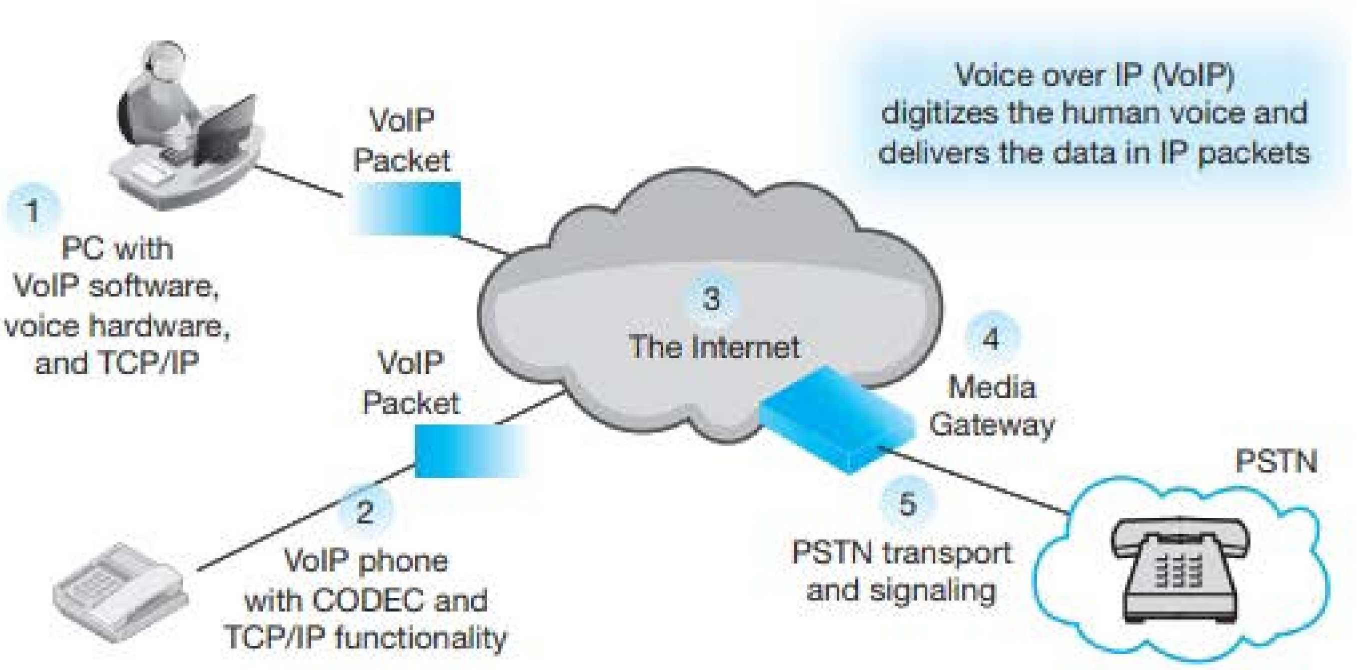
Voice over IP (VoIP) digitizes the human voice and delivers the data in IP packets
VoIP Packet
1
PC with VoIP software, voice hardware, and TCP/IP
3
The Internet
4
Media Gateway
VoIP Packet
PSTN
5
PSTN transport and signaling
2
VoIP phone with CODEC and TCP/IP functionality

FIGURE 11-20 Voice over IP (VoIP) Components

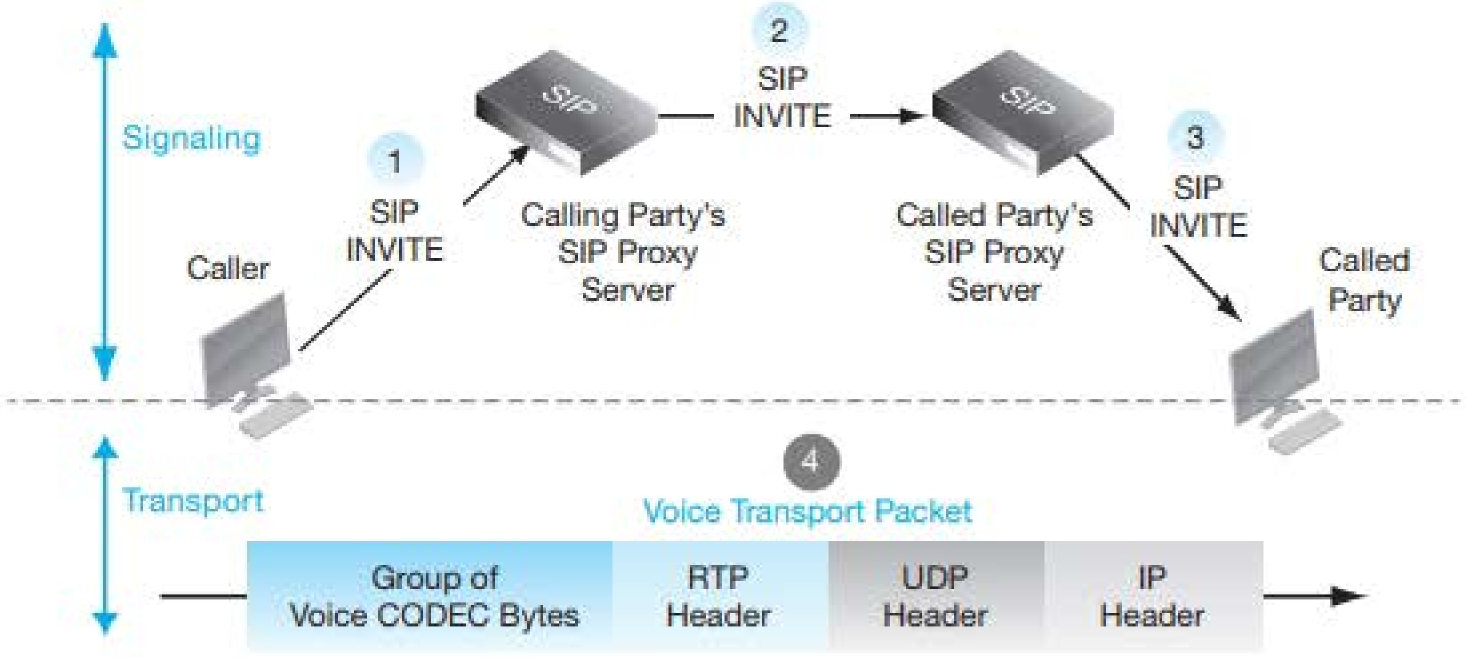

FIGURE 11-21 VoIP Signaling (SIP) and Transport Packet

- Signaling consists of the communication needed to set up circuits, tear down circuits, handle billing information, and do other supervisory chores.
- Transport is the actual carriage of voice.

Figure 11-21 illustrates the **Session Initiation Protocol (SIP)**, which is the main signaling protocol for VoIP. Each subscriber has a SIP proxy server. The calling VoIP telephone sends a SIP INVITE message to its SIP proxy server. This message gives the IP address of the receiver. The caller's SIP proxy server then sends the SIP INVITE message to the called party's SIP proxy server. The called party's proxy server sends the SIP INVITE message to the called party's VoIP telephone or multimedia PC.

After SIP creates a connection, the two VoIP clients begin communicating directly. This is the beginning of transport, which is the transmission of voice between callers. VoIP, as its name suggests, operates over routed IP networks. Therefore, digitized voice must be carried from the sender to the receiver in packets.

Test Your Understanding

20. a) Is SIP a signaling protocol or a transport protocol? b) Describe how SIP initiates a communication session.

The VoIP Transport Packet

Signaling includes session setup, breakdown, and other supervisory communication, whereas transport, again, is the transmission of packets containing fragments of voice or video between the two users.

VoIP Transport Packets As noted in Chapter 1, long application messages are fragmented into smaller pieces that can be carried in individual packets. Each packet carries a small part of the application message. Figure 11-21 shows a VoIP transport packet. Here, the application "message" is a stream of voice CODEC bytes. Each packet carries a few bytes of the conversation.

UDP with RTP at the Transport Layer TCP allows reliable application message delivery. However, the retransmission of lost or damaged TCP segments can take a second or two—far too long for voice conversations. Voice needs to be transmitted in real time. Consequently, VoIP transport uses UDP at the transport layer. UDP reduces the processing load on the VoIP telephones, and it also limits the high network traffic that VoIP generates. If packets are lost, the receiver creates fake noise for the lost CODEC bytes. It does this by extrapolating between the content of the preceding and following packets.

Although UDP must be used instead of TCP, UDP has two serious limitations for VoIP. Consequently, VoIP adds an additional header, a **Real Time Protocol (RTP)** header, to make up for these two deficiencies.

- First, UDP does not guarantee that packets will be delivered in order. RTP adds a sequence number so that the application layer can put packets in the proper sequence.
- Second, VoIP is highly sensitive to jitter, which is variable latency in packet delivery. Jitter literally makes the voice sound jittery. RTP contains a time stamp for when its package of octets should be played relative to the octets in the previous packet. This allows the receiver to provide smooth playback.

The final VoIP packet, then, consists of an IP header, a UDP header, an RTP header, and a snippet of the voice conversation.

Test Your Understanding

21. a) In a VoIP traffic transport packet, what does the data field contain? b) What standard is used at the transport layer? c) What two limitations of UDP does the RTP address?

PEER-TO-PEER (P2P) APPLICATIONS

A major driving force for applications, as we noted in the introduction, is increasing client processing power, including processor speed, memory, storage, networking speed, and other matters. The first PCs, which arrived in the 1970s, were expensive toys, and they continued to have minimal processing power for many years. However, chip density has been doubling every 18 to 24 months, allowing processing chips to add much more functionality and allowing memory chips to hold more data. This doubling of chip density in about two years is known as Moore's Law. In addition, chip speeds have also been increasing at about the same rate, growing from megahertz cycle speeds in the 1980s to gigahertz processing speeds today. Combined, the exponential growth in chip density and speed means that processing power has doubled roughly every year. In the last few years, speed increases have been more modest as energy consumption has become more critical, but increasing chip density has permitted more parallel processing in software, making up for much of the impact of reduced cycle speeds.

This power growth dynamic has permitted us to have ever-smaller devices with impressive processing power. Smartphones are a lot smarter than they were just five years ago, and many IoT devices will be tiny but fairly capable. Much of this increasing power has been absorbed by every-more-capable user interface capabilities, but even

Client Processing Power Increases

Moore's Law: capacity of chips doubling about every two years

Speed also increases rapidly

Today, clients have more processing power than servers did just a few years ago

Yet this processing power and storage goes largely unused during work hours

Peer-to-Peer (P2P) Applications

In peer-to-peer (P2P) applications, one client provides services to another client

Peers are client computers that provide services to other client computers

This can save a great deal of money by not buying servers

Problems with Clients

Not on all the time so not always available

Do not get the same IP address each time

Users fear that P2P applications will use too many of their computer's resources

IT departments are concerned about the lack of central control

FIGURE 11-22 Peer-to-Peer Evolution (Study Figure)

smart watches are beginning to do impressive work. Over the next decade or more, the maturity of small devices should accelerate.

Although small devices have taken the spotlight away from traditional desktop and laptop PCs, both have turned into seriously powerful computers that match servers of a few years ago. Their transmission pipe to the Internet has also grown enormously in speed.

Yet nearly all the time, our desktops and laptops are idle. Even when we actively work, we only use a fraction of the device's power. These realities have caused many to wonder why we still use servers as much as we do. Why not have client PCs provide service to other client PCs? This insight has led to a growing number of **peer-to-peer (P2P) applications** that do exactly that. When one client computer provides P2P services to others, it seems odd to call it a client. We will follow the common practice of referring to computers that provide P2P services as **peers**.

Peers are client computers that provide services to other client computers.

A major attraction for users is the appeal of generously making their unused resources available to others. This creates a grassroots cooperative spirit among those who allow their computers to be peers for P2P applications.

P2P applications need to address some nonprocessing limits of desktops and laptops, however. One is simply that when they are not on, they are not available. Another is that clients get a different dynamic IP address each time they boot up. How can one peer find another to use?

In addition, P2P applications run in the background all the time on the clients that provide services to other clients. Of course, ways must be found to prevent P2P applications from being too "greedy" in using resources. If they reduce the performance of user machines enough to be noticeable to users, they are likely to be deleted.

A concern that many companies have is that when clients provide services to other clients, the central IT department loses some of its control. IT departments are increasingly wary of the security issues raised by "shadow IT" of all types. At the same time, IT departments are intrigued by the possibility of buying or using less expensive server time by taking advantage of P2P processing to use idle IT resources. Users, in contrast to IT departments, tend to view reduced IT control as a benefit rather than a problem. It leads to less red tape and more freedom to act.

Test Your Understanding

22. a) What is the promise of P2P applications? b) What issues do P2P applications create for users?

Skype

We only look at two P2P applications in this chapter. The first is Skype. This is a peer-to-peer Voice over IP (and Video over IP) application. We already saw traditional VoIP, so we chose Skype to illustrate differences between traditional and P2P applications in the same category. We also chose it because it illustrates how P2P applications usually deal with transient IP addresses by requiring each client to log into the system.

Skype is a P2P VoIP service that currently offers free calling among Skype customers over the Internet and reduced-cost calling to and from Public Switched Telephone Network customers. Skype offers a range of features, from phone calls to instant messaging and video calling. At the time of this writing, Skype is the most popular P2P VoIP service. Skype's free calls from computer to computer have greatly contributed to this popularity. Figure 11-23 illustrates how Skype operates.

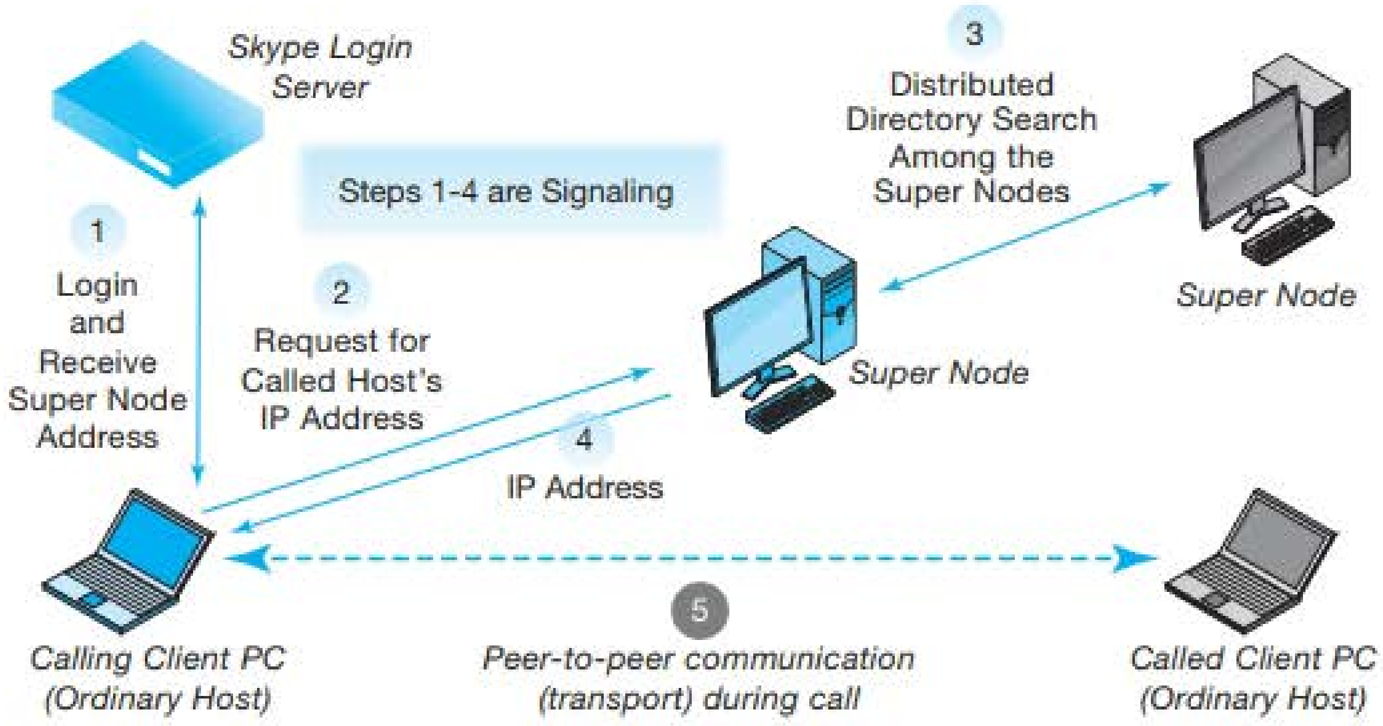
Skype Login Server
Steps 1-4 are Signaling
3
Distributed Directory Search Among the Super Nodes
Super Node
1
Login and Receive Super Node Address
2
Request for Called Host's IP Address
Super Node
4
IP Address
5
Calling Client PC (Ordinary Host)
Peer-to-peer communication (transport) during call
Called Client PC (Ordinary Host)

FIGURE 11-23 Skype P2P VoIP Operation

There are three main elements in the Skype network: the Skype login server, ordinary host nodes, and super nodes.

- The login server is a central server managed directly by Skype. It is the only centralized component in the Skype network.
- A host node is a Skype application that runs on a user's computer.
- A super node is a host node that takes on the work of signaling. Any regular host node may be made a super node if it has enough memory, network bandwidth, and CPU.

These elements are involved in the three steps that must occur for a user to place a call with Skype.

- *Step 1 Login.* First, a user must log in to the Skype login server. In this step, the username and password are authenticated. The Skype server also notes the user's IP address, which will be needed later in the directory search process. Login is the only step that involves a central server; the rest of the call process is done peer-to-peer using host nodes and super nodes. This step is like the login process in traditional voice over IP, where each client must log in to its own proxy server.
- *Step 2 Signaling/Directory Search.* After login, the user can place calls. His or her host will begin the signaling process. One of the main aspects of Skype signaling is the **directory search,** the process in which a Skype application looks up the username and IP address of the party it wants to contact. A Skype directory search is a completely P2P process that is done using the super nodes. This is a major difference from traditional voice over IP, where signaling uses servers (proxy servers).
- *Step 3 Transport.* Figure 11-24 compares Skype with traditional VoIP. Skype's super nodes handle signaling, but transport is done entirely by the two host nodes involved in the call. In transport, the voice packets are routed completely P2P, from caller to called party and vice versa. This is like traditional voice over IP transport, where the two clients also communicate directly.

Because the signaling and transport are done by peers rather than going through a central server, Skype only carries the burden of managing a login server. This greatly reduces Skype's operational costs, resulting in its low-cost calls.

	Traditional VoIP	Skype	Comparison
Login	Server: user logs into his or her proxy server	Server: User logs into the Skype login server	Similar
Signaling	Server: proxy server manages signaling	Peer-to-Peer: Super nodes manage signaling, using P2P searching	Major difference
Transport	P2P between the two hosts	P2P between the two hosts	Similar

FIGURE 11-24 Traditional VoIP versus Skype

Test Your Understanding

23. a) What service or services does Skype provide? b) List and define Skype's three main elements. c) Why is Skype login necessary? (This is a common problem in P2P processing.) d) What is a directory search in Skype? e) Which element of the Skype network is in charge of signaling? f) Which element of the Skype network is in charge of transport? g) Which of Skype's three steps is done P2P? h) Compare Skype and traditional voice over IP in terms of whether login, signaling, and transport are P2P or whether they use servers.

Tor

Another popular P2P application is Tor. **Tor** has a unique purpose—to permit anonymous IP transmission in which the IP address of the original sender is unknowable to the receiver. This goal is controversial because it is used by cybercriminals such as distributed denial-of-service (DDoS) attackers and crimeware purchasers. However, it is also used by those wishing to send tips to law enforcement agencies anonymously and to provide assurance of anonymity to ordinary private citizens concerned with indiscriminate government data collection.

Tor Routing Figure 11-25 illustrates how Tor works in simplified (but hopefully useful) fashion. The Tor network is a large collection of peer computers acting voluntarily as Tor routers.[8] Host X is the host wishing to transmit anonymously. Host X encrypts the message three times, then sends it to a selected Tor router, Tor Router 10.

- Tor Router 10 looks at the message. The message contains a key for decrypting its encrypted content. The router uses the key to decrypt the encrypted content. Note

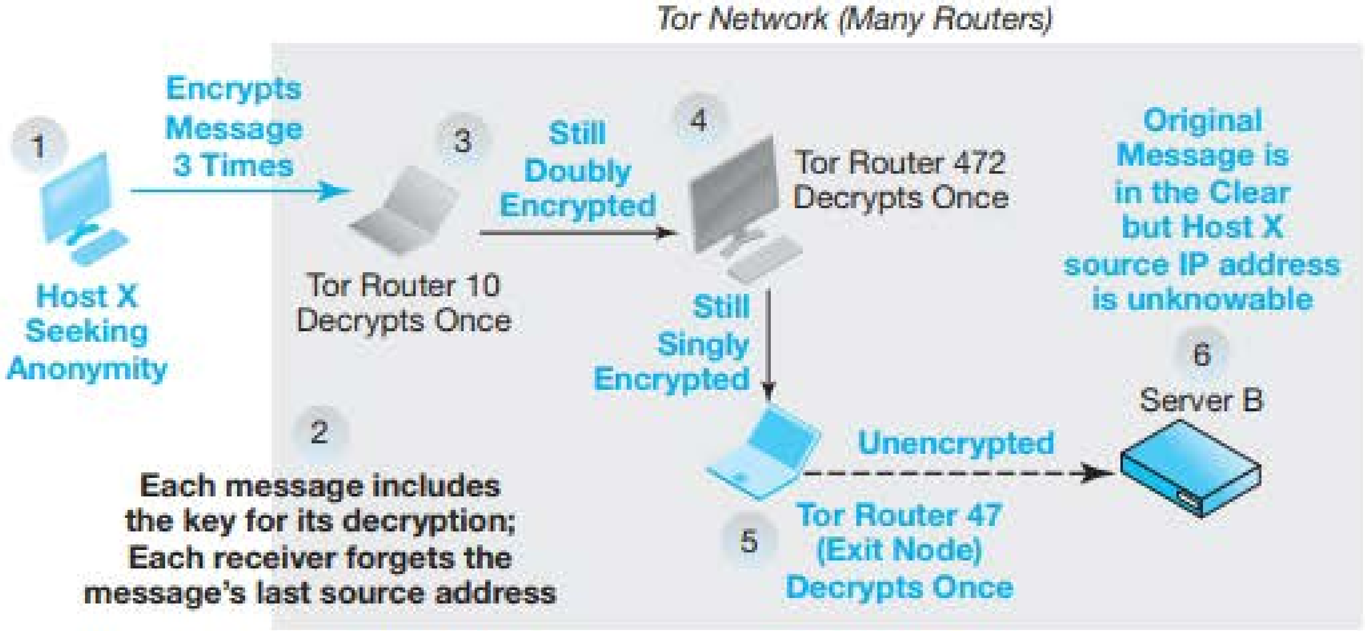
Tor Network (Many Routers)
1
Encrypts Message 3 Times
Host X Seeking Anonymity
3
Tor Router 10 Decrypts Once
Still Doubly Encrypted
4
Tor Router 472 Decrypts Once
Still Singly Encrypted
Original Message is in the Clear but Host X source IP address is unknowable
6
Server B
Unencrypted
5
Tor Router 47 (Exit Node) Decrypts Once
2
Each message includes the key for its decryption; Each receiver forgets the message's last source address

FIGURE 11-25 Tor Anonymous Transmission Network: A Simplified View

[8] These are not Layer 3 routers. They operate at the application layer. However, they do send messages across a group of peer nodes, so the name "router" is certainly evocative of what the peers do.

that Tor Router 10 knows the source IP address of the packet and the decryption key. However, it deliberately forgets them, wiping knowledge of them from its memory and storage. Note also that Tor Router 10 cannot read the original message. That message is still doubly decrypted.

- Tor Router 10 then forwards the result to Tor Router 472. This Tor router repeats the process, decrypting the now-doubly-encrypted message with the key it receives, forgetting the key and source IP address, and passing the now singly encrypted message on to Tor Router 47.
- Tor Router 47 does another decryption. This time, however, when the message is decrypted, it is in the clear, readable to anyone. Tor Router 47 then sends this message on to Server B.

Anonymity, not Confidentiality Server B can read the message. In addition, it knows that the packet's source IP address was Tor Router 47. However, this knowledge does it little good. If it can find out what Tor Router 47 knows, all it will learn is that Tor Router 47 is a Tor router. It cannot even learn the source IP address of the Tor router that sent the message to Tor Router 47—much less the IP address of Host X. Anonymity has been achieved.[9]

What Does the Exit Node Know? The original message is fully decrypted by Tor Router 47. This Tor Router is called a **Tor exit node** because it is the point at which the message leaves the Tor network. Tor Router 47 is slightly dangerous because it can read the unencrypted message before sending it on. It can then deliver this message to someone trying to break the Tor network's cryptographic protection. However, this is merely a violation of confidentiality, and the Tor network does not aim to offer confidentiality. Its only promise is anonymity, which it fulfills very well.[10]

Test Your Understanding

24. a) Does Tor try to achieve confidentiality for the original message? b) Does Tor try to achieve anonymity for the original message? c) What does each Tor router do when a message arrives? d) How is the risk created by the exit node different from the risk created by intermediate Tor routers?

[9] It might help you to understand Tor if you knew that the abbreviation originally stood for The Onion Router. Each Tor Router along the way peels away one layer of encryption from the "onion" message. However, fingerprints are gone from the original outer layer. This analogy is limited because fingerprints are still on the previous outer layers, which may still be lying around. Each stage in the Tor process erases all digital fingerprints that could identify the original sender.

[10] However, its anonymity protection, although strong, is not quite absolute. When the FBI took down the crime site Pirate Bay, it succeeded in breaking anonymity. A short time later, Pirate Bay came back up on a Tor network, but many criminal hackers avoided it, believing that it was an FBI front to identify IP addresses.

END-OF-CHAPTER QUESTIONS

Thought Questions

1-1. The sender uses HTTP to transmit mail. What standard or standards will the receiver use to download the message?

1-2. In VoIP, which of the following is transport or signaling: a) SIP b) The delivery of voice between users c) RTP d) Call setup e) CODEC data?

1-3. Skype uses super nodes, which do more work than ordinary P2P nodes. How do you think nodes become super nodes? There is nothing in the text that will help you with this question. Think broadly in terms of what costs P2P computing imposes on network peers.

Appendix

Managing the Security Process

Security is a process, not a product.[1]

LEARNING OBJECTIVES

By the end of this appendix, you should be able to:

- Discuss failures to stop the attack in the Target breach.
- Explain why security is about management far more than it is about technology.
- Explain the Plan-Protect-Respond-Cycle that governs defensive thinking in security.
- Describe and apply major security planning principles, including risk analysis thinking, comprehensive security, defense in depth, weakest links, single points of takeover, least permissions, comprehensive identity management, segmenting networks into different security domains, and organizational system security.
- Describe and apply policy-based security management.
- Describe how to respond to successful break-ins, including the use of Computer Security Incident Response Teams (CSIRTs), real-time fail-over, and intrusion detection systems.

[1] Ben Schneier, "Computer Security: Will We Ever Learn?" *Crypto-Gram Newsletter*, May 15, 2000, https://www.schneier.com/crypto-gram-0005.html.

THIS APPENDIX

Most teachers use this textbook for the core networking course. The 11 chapters provide that coverage. There is a lot of security in those chapters, especially in Chapter 4. This reflects today's networking job, which for better or worse deals heavily with security.

Some teachers use this book for a combined core course in networking and security. That requires some higher-level security concepts. This appendix provides them. For a networking and security course, I like to teach this appendix after Chapter 4, the main security chapter. That way, students can use and reinforce these security concepts in the rest of the course. However, it is also a nice way to end a semester because there are a lot of cool concepts in it.

Many teachers use this appendix in their network courses despite having a separate core security course. This reinforces the knowledge and skills they learn in the core security course. Today's students can never get too much security.

FAILURES IN THE TARGET BREACH

After every breach, companies should pause to learn from the experience. If this type of reflection leads to appropriate changes, it may prevent similar breaches in the future. It may even warn the company that its overall security is in trouble.

The Security of Business Partners One lesson from the Target breach is that you cannot trust external business partners to have good security. In the case of Fazio Mechanical Services, an employee fell for a spear phishing attack. This could happen in any company. However, Fazio made it more likely. It used the free consumer version of an antivirus program, Malwarebytes Anti-Malware.[2] This free version did not assess arriving e-mail messages and attachments. It only looked for malware already on the computer and then only occasionally. If Fazio had used a commercial antivirus program for their e-mail, the employee probably would have seen a warning that opening an attachment was a bad idea or even that a specific threat existed in the attachment.

Inadequate Network Segregation The breach taught several lessons about Target's security. After the attackers gained a foothold on the vendor's server, they moved into more sensitive parts of the network to download malware onto the point-of-sale (POS) terminals, compromise a server to create a holding server, and compromise another server to act as an extrusion server. The low-security and highly sensitive parts of the network should have been segregated. (Banks do not let customers walk around in the vault.)

Not Following Up on Specific Warnings An even worse issue is that Target received explicit alerts when the attackers were setting up the extrusion server. The thieves had to download malware onto the extrusion server before using it. Target used the FireEye intrusion detection program. FireEye's intrusion detection team notified Target's Minneapolis security staff that this downloading had occurred in

[2] Brian Krebs, "Email Attack on Vendor Set Up Breach at Target," February 14, 2014. http://krebsonsecurity.com/2014/02/email-attack-on-vendor-set-up-breach-at-target/.

a high-priority alert on November 30, 2013.[3] In addition, the thieves had trouble with the initial malware. They had to make additional updates on December 1 and December 3. These resulted in additional FireEye warnings being sent to Target's Minneapolis security group. Had Target followed up on these warnings, they could have stopped or at least reduced the data extrusion, which began on December 2.[4]

Keeping Up with the Threat Environment for POS Systems Target may have been lax in reacting to the danger of POS attacks. In April and August of 2013, VISA sent Target and other companies warnings about new dangers regarding POS data theft.[5] It appears that Target's own security staff expressed concern for the company's exposure to charge card data theft.[6] Target did not respond to this risk aggressively, another serious lapse.

Kill Chain Analysis for the Target Attack Overall, Figure A-1 shows that the thieves had to succeed at every step in a complex series of actions. Lockheed

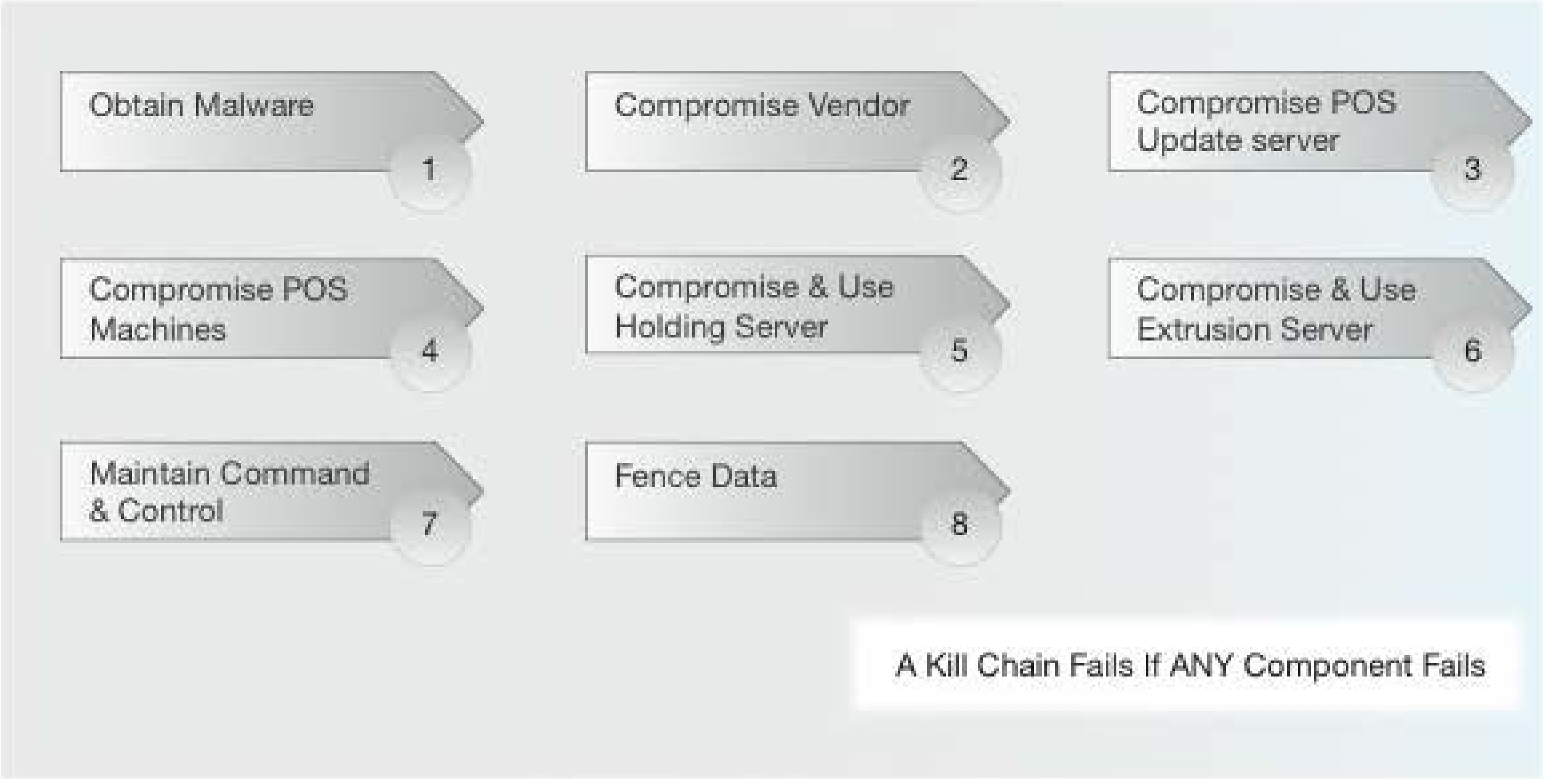
Obtain Malware
1
Compromise Vendor
2
Compromise POS Update server
3
Compromise POS Machines
4
Compromise & Use Holding Server
5
Compromise & Use Extrusion Server
6
Maintain Command & Control
7
Fence Data
8
A Kill Chain Fails If ANY Component Fails

FIGURE A-1 Kill Chain Analysis: Breaking Any Link Stops the Attack

[3] Michael Riley, Ben Elgin, Dune Lawrence, and Carol Matlack, "Missed Alarms and 40 Million Stolen Credit Card Numbers: How Target Blew It," Bloomberg Businessweek, March 13, 2014, http://www.businessweek.com/articles/2014-03-13/target-missed-alarms-in-epic-hack-of-credit-card-data.

[4] Aviv Raff, "PoS Malware Targeted Target," Seculert, January 16, 2014, http://www.seculert.com/blog/2014/01/pos-malware-targeted-target.html.

[5] Jim Finkle and Mark Hosenball, "Exclusive: More Well-Known U.S. Retailers Victims of Cyber Attacks Sources," Reuters, January 12, 2014, http://www.reuters.com/article/2014/01/12/us-target-databreachretailers-idUSBREA0B01720140112.

[6] Danny Yadron, Paul Ziobro, and Devlin Barrett, "Target Warned of Vulnerabilities Before Data Breach," *Wall Street Journal*, February 14, 2014, http://online.wsj.com/news/articles/SB10001424052702304703804579381520736715690.

Martin's Computer Incident Response Team[7] staff calls this a **kill chain**, which is a term borrowed from the military. The kill chain concept was designed to visualize all the manufacturing, handling, and tactical steps needed for a weapon to destroy its target. Failure in any step in a kill chain will create overall failure.

Lockheed Martin has urged companies to implement security kill chain analysis and look for evidence that one of the steps is occurring. Success in identifying an operating kill chain may allow the company to terminate it or at least disrupt or degrade it. The warnings when malware was put on the extrusion server should have done exactly that.

Until one understands likely kill chains in depth, however, it is difficult to realize that particular events are related, how they are related, and what type of attack they are part of. Conversely, understanding a kill chain can allow the company to act before an attack fitting that pattern even begins. For example, even cursory thinking about charge card data theft would lead the company to realize that thieves would probably use FTP transfers from unusual servers, that command communication would probably use certain ports in firewalls, and so forth.

Security Is a Process, not a Product Even well-defended companies suffer security compromises. However, when strategic planning is not done well, if protections are not put into place, or if the security staff is not aggressive in doing the work required for the protections to succeed, the risk of compromises becomes a near certainty. Security expert Ben Schneier has often said that "Security is a process, not a product." Boxes and software are not magic talismans. They must be backed by highly effective management and implementation processes. Schneier has also said, "If you think technology can solve your security problems, then you don't understand the problems and you don't understand the technology."[8] Target failed to understand its security problems, and it failed to develop processes that were effective against the threats it faced. However, Target is merely an object lesson. Many firms have inadequate security processes, and few have uniform excellence in how they manage and implement security.

Test Your Understanding

1. a) What security mistake did Fazio Mechanical Services make? b) Why do you think it did this? (This requires you to speculate.) c) How might segregation of the network have stopped the breach? d) Why do you think the Minneapolis security staff did not heed the FireEye warning? (This also requires you to speculate.) e) What warnings did Target not responded to adequately? f) What happens in a kill chain if a single action fails anywhere in the chain? g) How can kill chain analysis allow companies to identify security actions it should take? h) Explain why security is a process, not a product.

[7] Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," Lockheed Martin, 2011, http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf.

[8] Bruce Schneier, *Secrets and Lies*, 15th Anniversary ed. (Indianapolis, Ind.: Wiley, 2000).

THE PLAN–PROTECT–RESPOND CYCLE

Figure A-2 shows the overall process that companies should follow to deal with threats. On the left is the threat environment, which consists of the attackers and attacks the company faces. We looked at the threat environment in Chapter 4.

The right side of the figure illustrates how companies mount defenses against the threats they face. The figure shows that companies constantly cycle through three phases of security management. This is the **plan–protect–respond cycle**.

Planning The company must **plan** how it will protect its assets. New assets appear, and existing assets change in value. Plans must change accordingly. Adversaries constantly change their attacks, and companies must change their plans to meet the changing threat environment.

Protecting Next comes **protection**, in which companies provide actual protections on a day-to-day basis. We looked at protections such as firewalls in Chapter 4. In Figure A-2, the protect phase bubble is larger than the other two. This emphasizes the fact that the protect phase is much larger than the other two phases in terms of time and resource expenditures. However, without extensive and insightful planning, it is possible to spend a great deal of time and effort mounting protections without being effective.

Responding Finally, a company must **respond** when it suffers a successful security attack. Security breakdowns are inevitable, and it is professional malpractice not to have effective plans for what to do when they occur. Target did not.

Not Exactly in Sequence Logically, planning comes before protection, and response comes afterward. Reality mocks this logic. Security failures constantly require

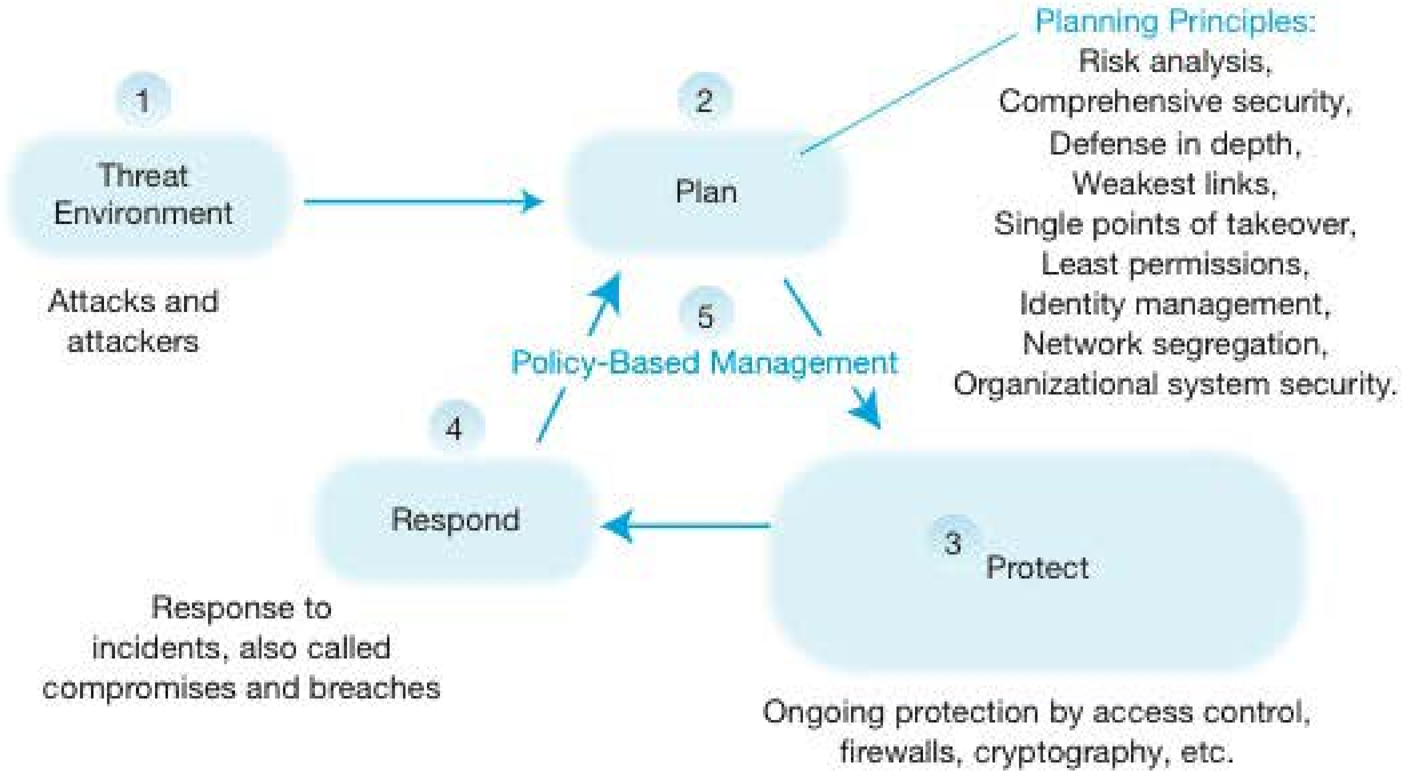

FIGURE A-2 The Threat Environment and the Plan-Protect-Respond Cycle in Security Management changes in planning and protection. Planning, protection, and response are simultaneous processes in the real world. Companies make comprehensive plans once a year or more, but defenders quickly learn the military axiom, "No plan survives first contact with the enemy." In war, the other side tries to learn your plan and use it against you. Security adversaries do the same. At the same time, trying to react without a core plan to improvise around is absurd.

Test Your Understanding

2. a) What happens in each stage of the Plan–Protect–Respond cycle? b) Which stage consumes the most time?

SECURITY PLANNING PRINCIPLES

Perhaps more than any other aspect of IT, effective security depends on effective planning. Security planning is a complex process that we discuss only briefly. We focus on some key planning principles that must be observed in all security thinking. These principles are shown in Figure A-2.

Risk Analysis

Many believe that the goal of security is to stop all threats to the corporation. Surprisingly, that is not true. Fundamentally, stopping all attacks is impossible. Despite strong security efforts, there will always be some risk of a compromise. There has always been crime in society, and there always will be. The same is true of security incidents. No matter how much money a company spends on security, it never stops all threats. Rather, the goal of security is to reduce the risk of attacks to the extent that is **economically feasible**, that is, to the extent that the benefits outweigh the costs.

The goal of security is to reduce the risk of attacks to the extent that is economically feasible, that is, to the extent that the benefits outweigh the costs.

Risk analysis is the process of balancing risks and protection costs. Corporate security planners must ask whether every countermeasure is economically justified. For example, if the probable annual loss from a threat is $10,000 but the security measures needed to thwart the threat will cost $200,000 per year, the firm obviously should not spend the money. Instead, it should accept the probable loss if there is no other available countermeasure.

Risk analysis is the process of balancing risks and protection costs.

A Risk Analysis Example Figure A-3 gives an example of a risk analysis. Without a countermeasure, the damage per successful attack is expected to be $1,000,000, and the annual probability of a successful attack is 20%. Therefore, the annual probable annual damage is $200,000 without a countermeasure ($1,000,000 times 20%). This is the **base case** for the analysis—doing nothing.

Countermeasure	None (Base Case)	A	B
Total Cost of Incident (TCI), per occurrence	$1,000,000	$500,000	$1,000,000
Annual probability of a successful attack	20%	20%	15%
Annual probable damage	$200,000	$100,000	$150,000
Annual cost of countermeasure	$0	$20,000	$60,000
Net annual probable outlay	$200,000	$120,000	$210,000
Annual saving compared to no countermeasure	NA	$80,000	($10,000)

FIGURE A-3 Risk Analysis Calculation

Countermeasure A The first countermeasure will cut the damage of a successful attack in half. So the damage per successful attack is expected to be $500,000 instead of a million dollars. The countermeasure will not reduce the probability of a successful attack, so that continues to be 20%. With Countermeasure A, then, the annual probable damage will be $100,000 ($500,000 times 20%). This seems attractive compared with having no countermeasure. However, a countermeasure is never free. This one will cost $20,000 per year. Therefore, the net annual probable cost is $120,000 with Countermeasure A—$100,000 in probable damage and $20,000 for the countermeasure.

Countermeasure A, then, will reduce the net annual probable outlay from $200,000 to $120,000. The countermeasure therefore gives an annual saving of $80,000 per year compared to the base case. Countermeasure A is cost effective.

Countermeasure B The second countermeasure does nothing to reduce the total cost of an incident. However, it reduces the probability of attack from 20% to 15%. Therefore, it reduces the annual probable damage from $200,000 to $150,000. Unfortunately, the countermeasure's annual cost is $60,000 per year. The net annual probable outlay from incidents plus the countermeasure is therefore $210,000. This is $10,000 more than the base case's $200,000. It does not make sense economically to implement this countermeasure at all.

The Decision For this situation, the final choice is simple. The company should implement Countermeasure A. If it does, it can expect to reduce its annual outlay for this resource by $80,000. Countermeasure B is obviously a bad choice. It would actually increase the firm's probable annual cost outlay. Of course, every situation is different. Sometimes, multiple countermeasures will be able to save money, and in some cases, none will. If none will reduce costs, the choice should be to do nothing.

Countermeasure Costs Security professionals may be tempted to think of countermeasure costs in terms of hardware and software. However, most countermeasures require extensive security labor. In fact, labor is often the biggest cost. More broadly, security often increases labor costs for the users they are defending. If users spend even a few extra minutes each time they use a resource, this can lead to substantial

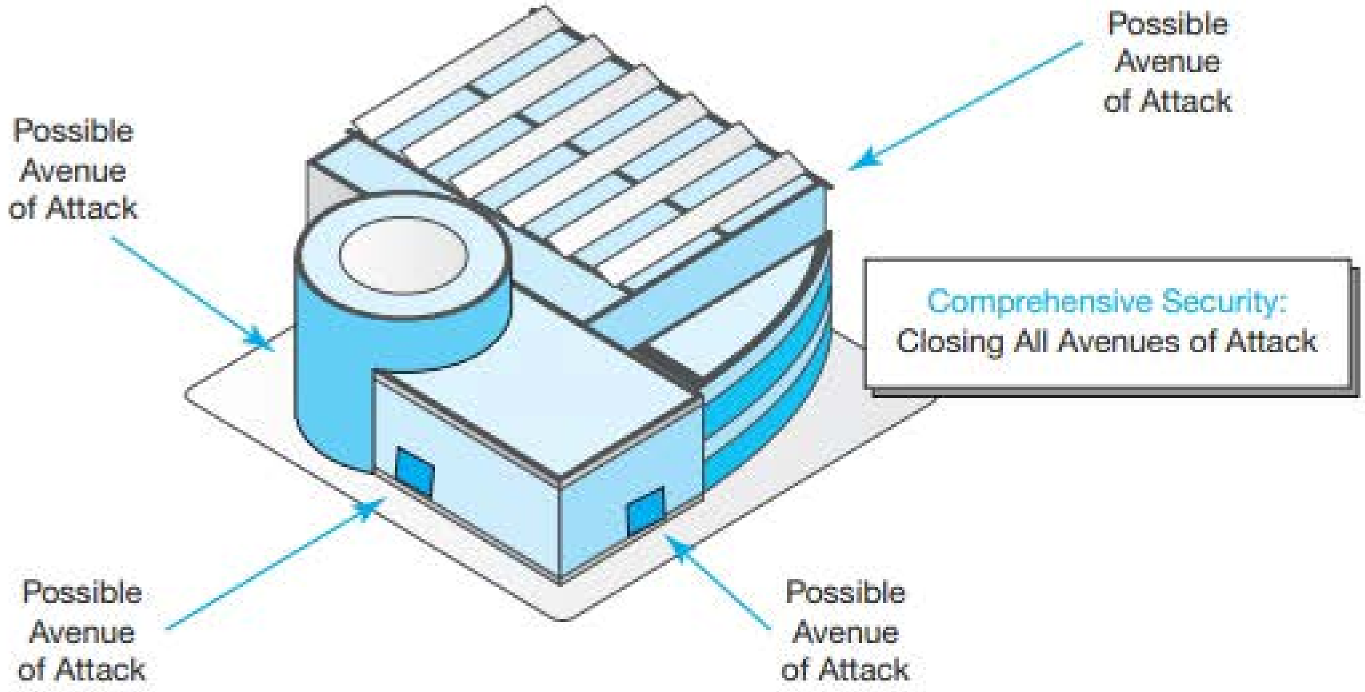
Possible Avenue of Attack
Possible Avenue of Attack
Comprehensive Security: Closing All Avenues of Attack
Possible Avenue of Attack
Possible Avenue of Attack

FIGURE A-4 Comprehensive Security

user labor cost. It could tip the scales against installing the countermeasure. The **total cost of a countermeasure** must include all factors.

Test Your Understanding

3. a) Comment on the statement, "The goal of security is to eliminate risk." b) What is risk analysis? c) Repeat the risk analysis in Figure A-3, this time with Countermeasure C reducing damage severity by a quarter and the likelihood of an attack by 75%. The annual cost of Countermeasure C is $175,000. Show your calculations like the table does. d) What do you conclude? Justify your answer.

Comprehensive Security

To be safe from attack, a company must close off all avenues of attack. Figure A-4 illustrates this simple but fundamental principle. There are four avenues of attack, and the defender must protect all four. In contrast, an attacker only needs to find *one* unprotected avenue to succeed. Although it is difficult to achieve **comprehensive security**, it is essential to come as close as possible.

Comprehensive security is closing off all avenues of attack.

Test Your Understanding

4. a) What is comprehensive security? b) Why is it important?

Defense in Depth and Weakest Links

Defense in Depth Another critical planning principle is defense in depth. Every protection will break down occasionally. If attackers must break through only one line of defense, they will succeed during these vulnerable periods. However, if an

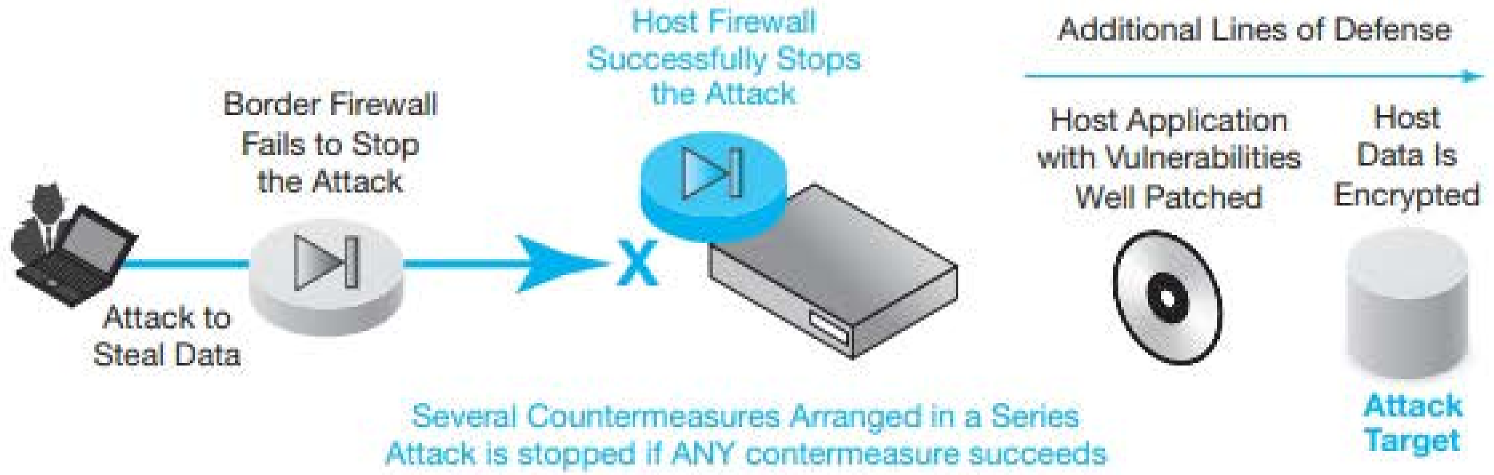

FIGURE A-5 Provide Defense in Depth for Resources

attacker must break through two or more lines of defense, the breakdown of any single defense technology will not allow the attacker to succeed. Having successive lines of defense that must *all* be breached for an attacker to succeed is called **defense in depth**.

Having several lines of defense that must all be breached for an attacker to succeed is called defense in depth.

Figure A-5 illustrates defense in depth. In the figure, there are four protections in succession. The first is a border firewall. The second is a host firewall on a server targeted by the attacker. The third is the use of good practice in patching application vulnerabilities on the server. The fourth is encrypting all server data so the attacker cannot learn sensitive information even if all other defenses fail.

The figure shows what happens if the border firewall fails to stop an attack. The host firewall will stop it. The company should fix the border firewall quickly, so that it becomes part of the effective defense, but attack packets will not get through to the target data while the border firewall is being fixed.

Identify and Manage Weakest Links Defense in depth increases security with a series of protections. In contrast, many individual protections consist of a series of internal parts that must *all* work if protection is to succeed. If one fails, the countermeasure fails. For example, an antivirus program may protect a user by identifying a malicious attachment. However, if the user fails to use good judgment and opens the attachment anyway, there is no protection. A **weakest link** exists when the failure of a single part of a countermeasure can cause the entire countermeasure to fail to stop an attack. It is like the weakest link in a physical chain.

A weakest link exists when the failure of a single part of a countermeasure will cause the entire countermeasure to fail to stop an attack.

Figure A-6 shows how weakest links can compromise a countermeasure. Here the countermeasure is a firewall. The firewall has five components, all of which must be effective for the firewall to be effective. These are the firewall hardware, firewall

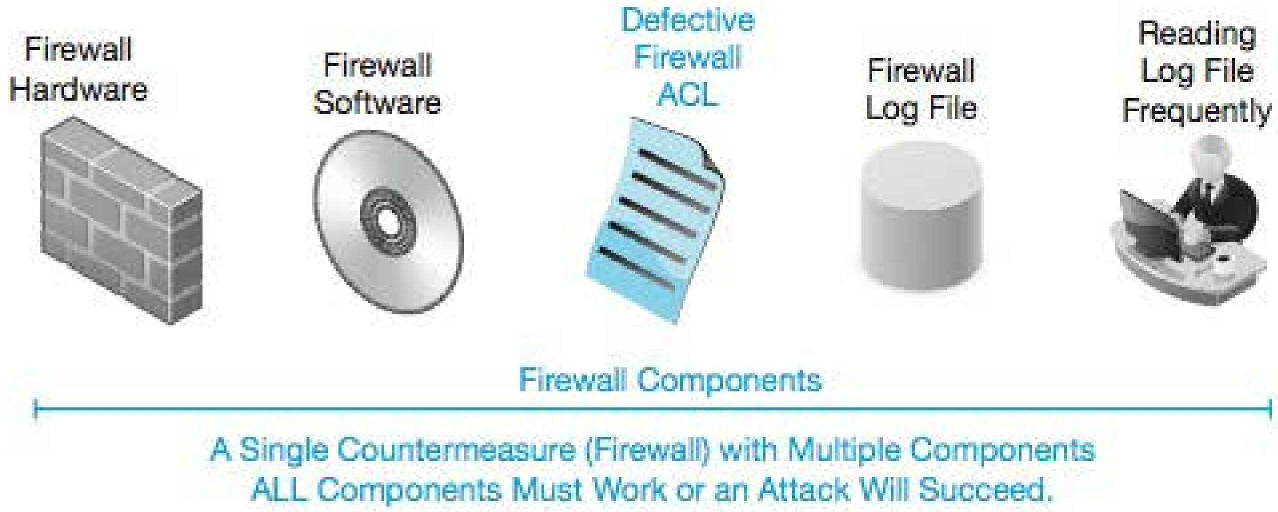
Firewall Hardware
Firewall Software
Defective Firewall ACL
Firewall Log File
Reading Log File Frequently
Firewall Components
A Single Countermeasure (Firewall) with Multiple Components
ALL Components Must Work or an Attack Will Succeed.

FIGURE A-6 Identify and Manage Weakest Links

software, a firewall access control list (ACL), the firewall log file, and the practice of reading the log file frequently. Even if all the other elements are fully effective, if the ACL is defective, the firewall will fail to stop an attack. Similarly, if the company fails to read the firewall log file regularly, it will fail to keep the ACL up to date, and this will also cause the firewall to fail.

Defense in Depth versus Identifying and Managing Weakest Links It is easy to confuse defense in depth and weakest link analysis because both have a series of elements. Figure A-7 compares them in terms of their nature, the number of countermeasures involved, whether the countermeasure has multiple components, and criteria for success or failure. The key point is that weakest link analysis involves a single countermeasure while defense in depth involves multiple countermeasures.

Test Your Understanding

5. a) What is defense in depth? b) Why is defense in depth necessary? b) When does a weakest link exist? c) Distinguish between defense in depth and weakest links.

	Defense in Depth	Weakest Link
Nature	Protection	Weakness
Number of Countermeasures	Multiple	One
Components per Countermeasure	NA	Multiple
Outcome	Protection if *any* countermeasure succeeds	Failure if *any* component fails

FIGURE A-7 Defense in Depth versus Weakest Links

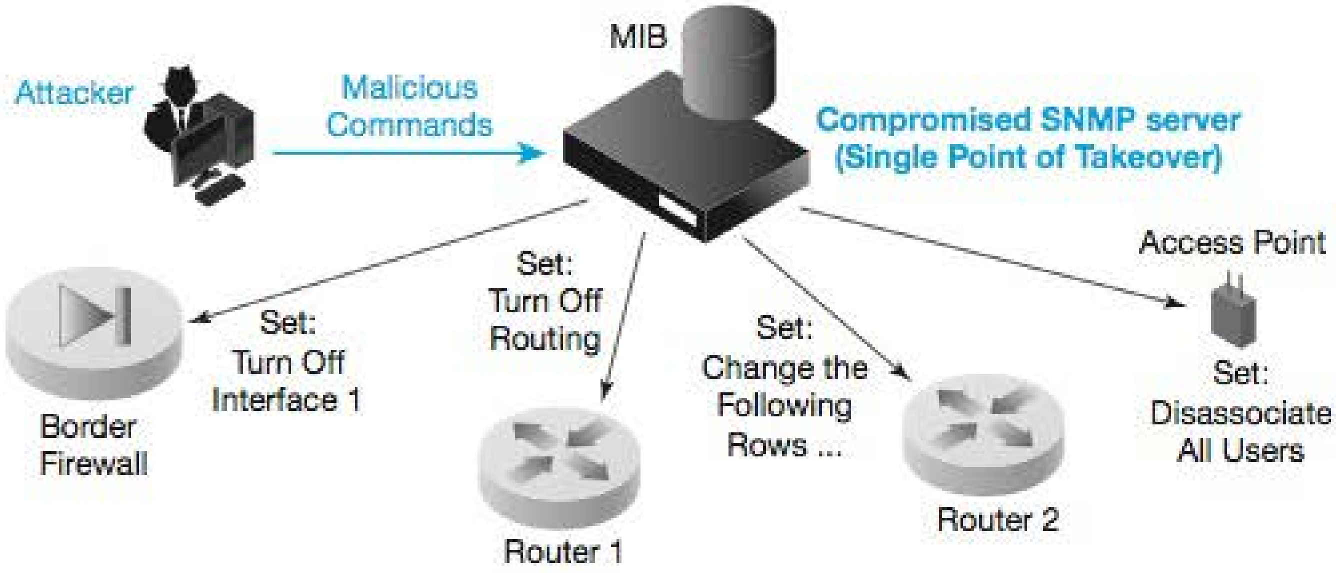
MIB
Attacker
Malicious Commands
Compromised SNMP server (Single Point of Takeover)
Access Point
Set: Turn Off Interface 1
Set: Turn Off Routing
Set: Change the Following Rows ...
Set: Disassociate All Users
Border Firewall
Router 1
Router 2

FIGURE A-8 Identifying and Protecting a Single Point of Takeover (An SNMP Server)

Identify and Manage Single Points of Takeover

Another principle is focusing on potential **single points of takeover**. We saw the Simple Network Management Protocol (SNMP) in Chapters 3 and 9 (see Figure A-8). If an attacker takes over the SNMP server, there is no end to the damage that he or she can do. The SNMP server is a potential single point of takeover, which means that if an attacker can take it over, the attacker gains control over a significant portion of your network. As noted in Chapter 9, companies with weak security should not use the SNMP Set command. However, if security is strong, companies should use SNMP Set commands to manage remote devices.

Companies usually cannot and do not want to eliminate single points of takeover. Having a firewall policy server greatly improves a company's control over its firewalls, eliminating inconsistencies and reducing management costs. It is critical for companies to identify all single points of takeover and harden them very well against attacks.

Test Your Understanding

6. a) Why must companies identify single points of takeover? b) What must companies do about potential single points of takeover?

Least Permissions

Security planners constantly worry about protecting access to resources. People with access to resources can damage them. Those without access cannot. Not surprisingly, companies work very hard to control resource access. **Access control** is limiting who may have access to each resource and limiting what he or she can do to the resource.

Access control is limiting who may have access to each resource and limiting what he or she can do to the resource.

Authorizations One aspect of access control is authentication, which is requiring users requesting access to prove their identities. However, just because you know

Authentication

Test the credentials of a supplicant

Accept or deny decision

Authorizations (Permissions)

After a supplicant is accepted as the True Party . . .

Decide what authorizations or permissions the True Party has

That is, what can the True Party do with the resource

Acceptance should rarely result in getting all permissions

FIGURE A-9 Authentication versus Authorizations (Study Figure)

who someone is does not mean that he or she should have unfettered access to your resources. (There undoubtedly are people you know who you would not let drive your car.)

Authorizations (Permissions) **Authorizations** or **permissions** are the actions that an authenticated user is allowed to take on the resource. For example, although you are permitted to view the U.S. Declaration of Independence, you may not add your signature at the bottom.

Authorizations or permissions are the actions that an authenticated user is allowed to take on the resource.

Authorizations and Authentication Servers We saw in the main chapters that when people wish to use a server or other resource, they must authenticate themselves. The server passes their credentials to an authentication server that vets credentials and sends an accept or deny message back to the server. As (Figure A-10) shows, it also returns a list of permissions. This employee can list files, that is, see what files are in a project directory. He or she may also read the files. However, he is not authorized to do anything else, such as edit files or delete them.

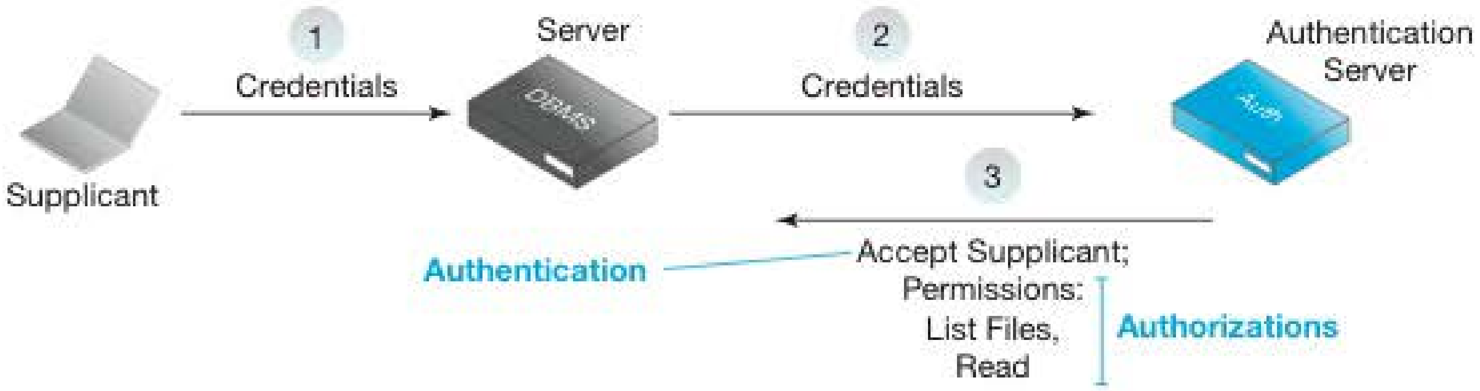

FIGURE A-10 Authentication and Authorizations when an Authentication Server is Used

Permission	Meaning	Team Member	Team Editors	Team Leader	Not Team Member
List files	See list of files in folder	x	x	x	
Read	Read a file	x	x	x	
Write	Edit a file		x	x	
Create/Delete File	Create new files, delete files from directory			x	
Change Permissions	Change permissions for others			x	

FIGURE A-11 Least Permissions for a Project Team's File Folder

Least Permissions An important principle in assigning permissions is to give each person the **least permissions** that he or she needs to accomplish his or her job. For example, if you give an employee permission to enter his or her building in a multi-building complex, you might not give that person permission to enter other buildings in the complex. Even if the person can enter a building, he or she may not have permission to enter special areas such as a bank vault or the corporate planning department. You may even restrict the floors the person can reach using the building's elevators.

Least permissions are the minimum permissions that the user needs to accomplish his or her job.

Example of Least Permissions Suppose, for example, the target asset is a file folder on a server. Figure A-11 shows some permissions that can be applied in a folder. They are shown in increasing order of authorization, meaning increasing order of risk. The figure also shows permissions to be given to various team members.

- All team members can see all files in the folder—but only their names, types, sizes, and other basic information. All team members also have the Read permission so they can read every file in the folder. This permission is sometimes called Read-Only because the person reading the file cannot edit it. These are the least permissions that make sense for everybody in the team.
- Some of the team members will be editors. The write permission allows them to edit files in the folder. These members therefore have more permissions, but they are still not allowed to do riskier things such as delete files. That would allow a rogue team member or an attacker who compromises an editor's computer to eliminate all files in the directory.
- The team leader exclusively retains further permissions. Only he or she can create new files, delete existing files, or change permissions for files.
- Finally, everyone else in the world receives no permissions in the directory. They cannot even see the names of files in the directory. They may not even be able to see that the directory exists.

Test Your Understanding

7. a) Distinguish between authentication and authorizations. b) What is another term for authorizations? c) What is the principle of least permissions? d) Why is it important?

Identity Management

Identity management means having comprehensive visibility and control over individual employee access to and permissions on all resources. The principle of least permissions is part of comprehensive identity management, but there is a major problem remaining.

Many compliance regimes require a company to have strong identity management.

Identity management means having comprehensive visibility and control over individual employee access to and permissions on all resources.

A problem in authentication is that big companies have many authentication servers, not just one. This creates the danger that different authentication servers will have incompatible information. This might, for example, allow someone who had just been fired to try logging into systems that use different authentication servers in hopes of finding one that had not been updated to remove his or her access credentials.

Identity Management Mandates in Compliance Regulation

This sort of thing has always been a problem, but compliance regulations increasingly require firms to have comprehensive identity management to be in full compliance. A major goal of comprehensive identity management is ensuring that there is uniformity in vetting credentials throughout the firm. There must be a way to synchronize data on all of these authentication servers.

Directory Servers

Figure A-12 shows that companies accomplish uniformity by storing credential information not on the authentication servers individually but on a directory server. **Directory servers** contain a great deal of general information about employees and computer resources, including telephone numbers, e-mail addresses, and, yes, security information.

When the database server on the left receives a request from a person (Joe) to log in, the DBMS server will pass Joe's credentials to an authentication server. If there is no directory server, the authentication server will make the accept/deny decision. If it accepts Joe, it will let Joe log in to have access to his resources and permissions.

However, if there is a directory server, as there is in this case, the authentication server will pass the request to the directory server.[9] The directory server will

[9] In smaller firms, directory servers allow user identity information to be managed on a single device, the directory server. Larger firms, however, usually have multiple directory servers, and often from multiple vendors. Most larger firms need metadirectory servers that synchronize identity data across the individual directory servers regardless of their different manufacturers. This adds another layer of complexity.

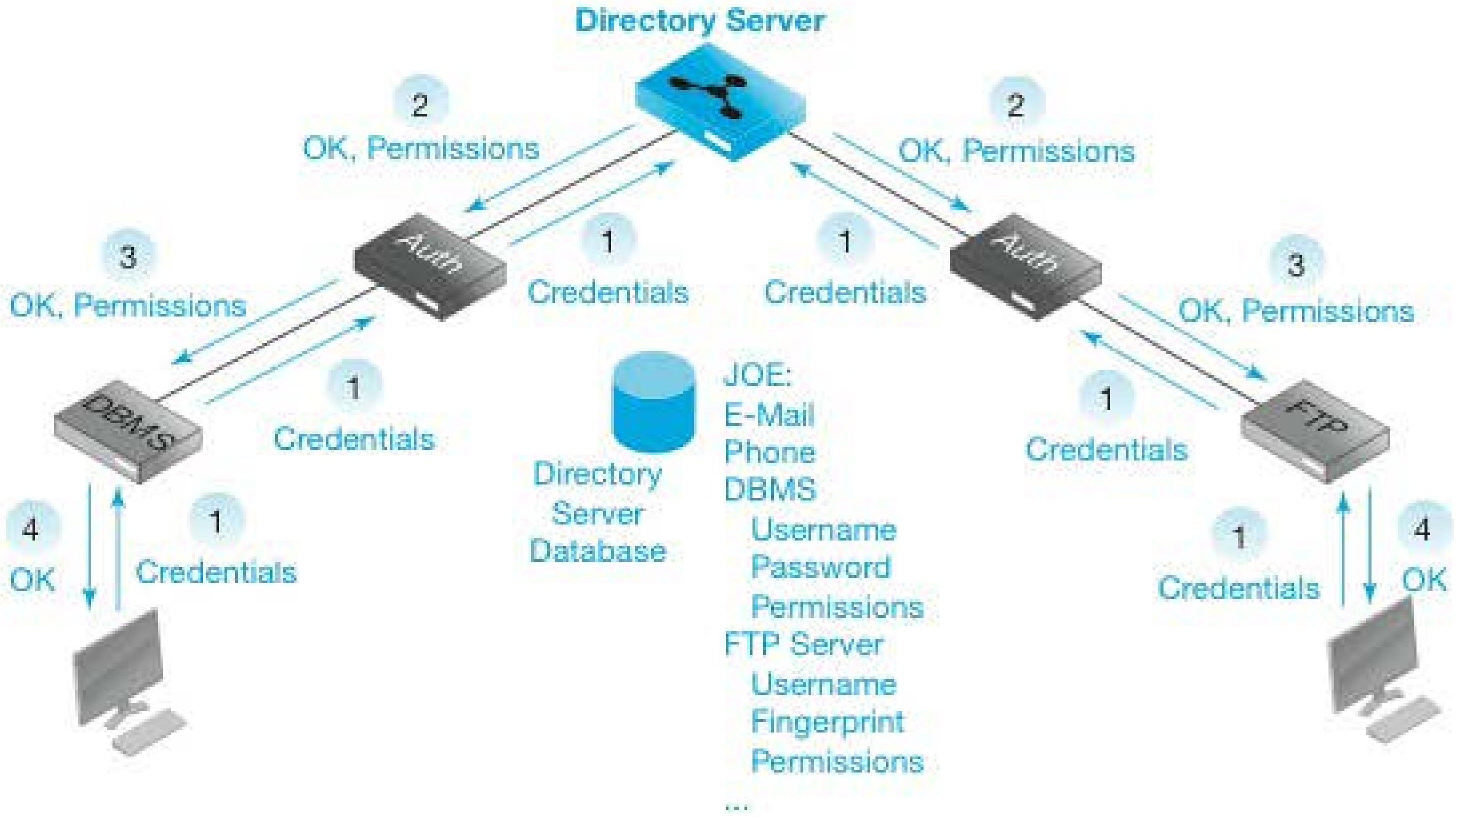

FIGURE A-12 Directory Server with Multiple Authentication Servers

make the decision and pass the acceptance decision plus permissions to the authentication server. This ensures that the vetting data is current. If an employee leaves the company, the security administrator merely disables his or her records on the directory server. That person is instantly cut off from access using any authentication server.

Role-Based Access Control The traditional way to manage identity data has been to decide what level of access to give to each user. This requires creating many access profiles, increasing the likelihood of an error. In addition, it is difficult to ascertain what permissions a particular authenticated user will need on a device. A common way to reduce this complexity is to use **role-based access control (RBAC)**, in which permissions are assigned to roles within each team. Using the example of documents for a project team we saw in Figure A-11, these roles might be leader, editor, general team

Traditional Individual-Based Access Control

- Many access permissions will have to be defined on different devices
- It is difficult to define what permissions an individual will need on a device
- This is error-prone

Role-Based Access Control (RBAC)

- Assign individuals to roles
- They receive role permissions
- If add Joe as an editor, no need to think about specific permissions
- Less likelihood of errors

FIGURE A-13 Role-Based Access Control (RBAC) (Study Figure) member, and outsider. It is much easier to understand what permissions a role will need than what an individual will need in the abstract. The person controlling access merely assigns each member to an assigned role. This reduces both cost and errors.

Test Your Understanding

8. a) How does a company address the problem of changing login credentials and permissions on various servers for an employee when he or she moves to a different department? b) What does a company do when an employee is terminated to ensure that no authentication server will give that person access to any resource? c) What is role-based access control? d) What are its advantages compared to assigning permissions based on individuals?

Segment the Network

Buildings in which confidential projects are worked on are divided to control communication between employees in different departments. Often, different parts of the building will have access controls to prevent unauthorized personnel from entering. A few parts of the building may have particularly strong access controls. For example, they may be on separate floors that cannot be reached by elevators without the correct access credentials. There may even be a guard on duty to check credentials for people arriving and to collect cellphones so they cannot be used internally. The guard may also check employees who are leaving to ensure that they have not taken trade secrets with them on paper or electronically.

In the Target breach, network segmentation was either not used or was defeated. The attackers first gained entrance to the Target network through a vendor access server. From there, they were able to move to the part of the network holding the POS download server. This should not have been possible. In **network segmentation**, the network is divided into different security domains, each with security controls that are appropriate to it. Communication between security domains is restricted, especially between security zones with different levels of security risks.

In network segmentation, the network is divided into different security domains, each with security controls that are appropriate to it. Communication between security domains is restricted, especially between security zones with different levels of security risks.

A good example of segmentation is the use of a **demilitarized zone (DMZ)**[10] for servers **public-facing servers**, that is, servers that must be accessed by the outside world. These servers are likely to be under constant attack, yet they must operate

[10] This name is confusing. In Korea and Vietnam, the demilitarized zone was the boundary area that lay between the two antagonists. It was called the demilitarized zone because it was supposed to be left unoccupied by either side. In practice, the DMZ was the bloodiest battle ground in the wars. Each side knew that they would face the enemy constantly. They had to be completely prepared for battle. By analogy, the security DMZ is where company servers that face the Internet and will certainly be attacked through the Internet must be placed so that they can be accessible to Internet users.

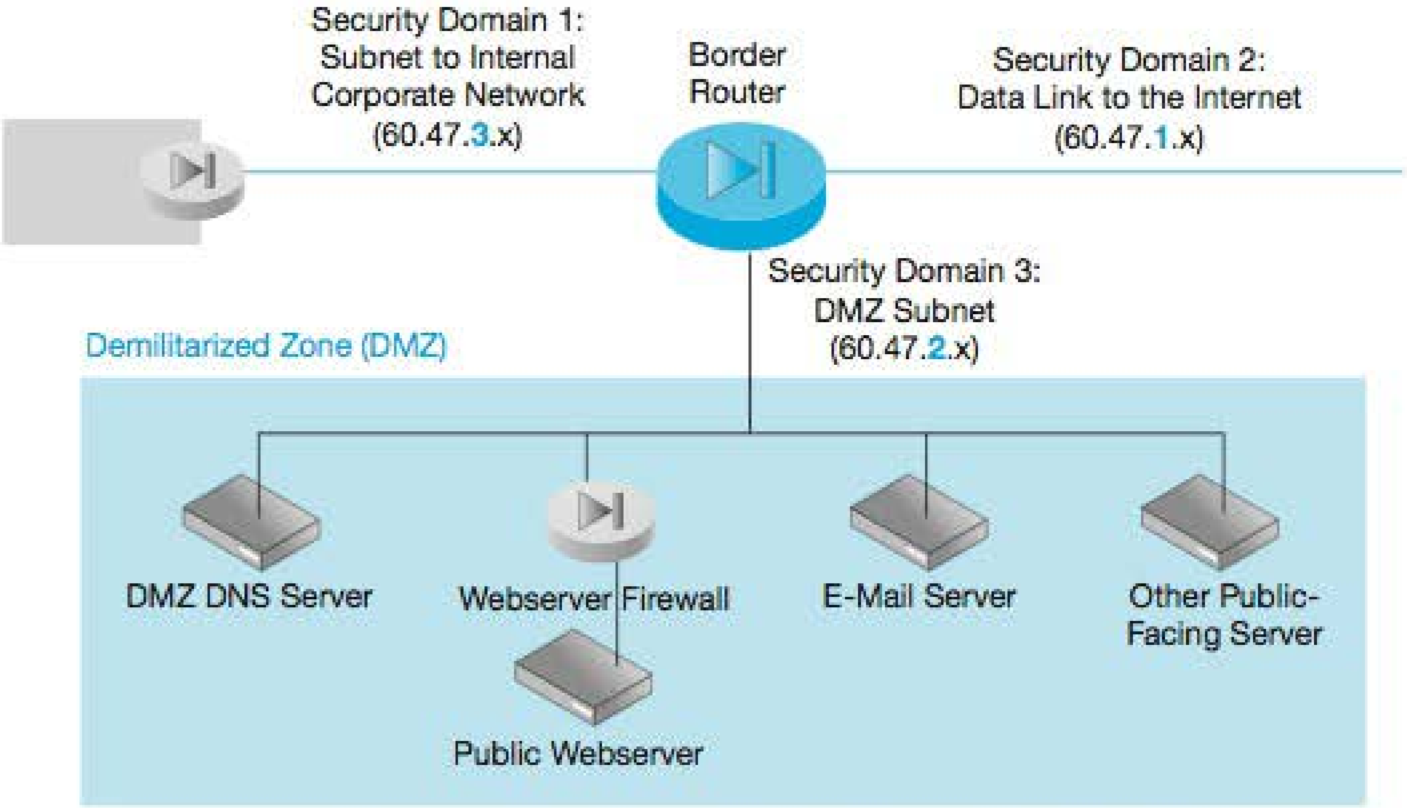

FIGURE A-14 Network Segregation Using a Demilitarized Zone (DMZ) for Public-Facing Servers

publicly to serve a firm's customers. Figure A-14 shows that many firms have an access router that connects three security domains. Like firewalls, routers often have access control lists that determine what traffic should be able to travel from one subnet to another.

The demilitarized zone (DMZ) is for public-facing servers that must be accessible to the outside world.

- First, there is the public Internet, over which firms have no control whatsoever.
- Second, there is the internal network, which should rarely be accessed directly from the Internet. Access controls should be very restrictive.
- Third, servers that must be made accessible to the public are placed in a third network subnet. This is called the demilitarized zone (DMZ). Access from the Internet must be very permissive, but communication from the DMZ to the internal corporate network should be minimized and rigidly controlled.

Because access to servers in the DMZ by Internet users must be easy, the servers themselves must be strongly hardened to survive the inevitable attacks they will face. To give an example, the webserver has a web firewall that focuses entirely on web-based attacks. The fact that the entire firewall server only needs to protect a single webserver means that it will use all its processing power to protect that firewall. This permits it to do very sophisticated filtering. In addition, this server can permit only web traffic to pass. Beyond that, the firewall will have detailed rules for threats specific to firewalls. DMZ servers should also cause minimal damage if they are compromised. For example, the DMZ directory server only knows the IP addresses of hosts in the DMZ. Compromising the directory server will not compromise the IP addresses of internal servers.

Test Your Understanding

9. a) Why is it important to segment networks into multiple security domains? b) How is this done for public-facing servers? c) How restricted should access be from the Internet to the DMZ? d) How restricted should access be from the DMZ to the internal network? e) How restricted should access be from the Internet to the internal network?

Organizational System Security

We often speak of information systems, which are combinations of people, procedures, and technology that produce useful information when needed. However, every information system is embedded in a larger organization system. This organizational system may be a functional department, a task force, a new product identification and development effort, or any other workflow in a firm. It is important to manage security in this broader organizational system because failure to achieve **organizational system security** will make technical protections useless. Figure A-15 shows that managing organizational system security is a complex job. It has many components, and any single failure can lead to a compromise.

Goals Every organizational system has goals. Security must be built one of these goals. It must also be realized that organization systems can never achieve all goals, so there is a danger that the security goal will be sacrificed for other goals.

People and Training The system's people must be capable of doing the required work and must understand pertinent security requirements. For that to happen, they must be trained to work securely. Implementing security well always requires training, not simply announcing security rules. If someone must add a new

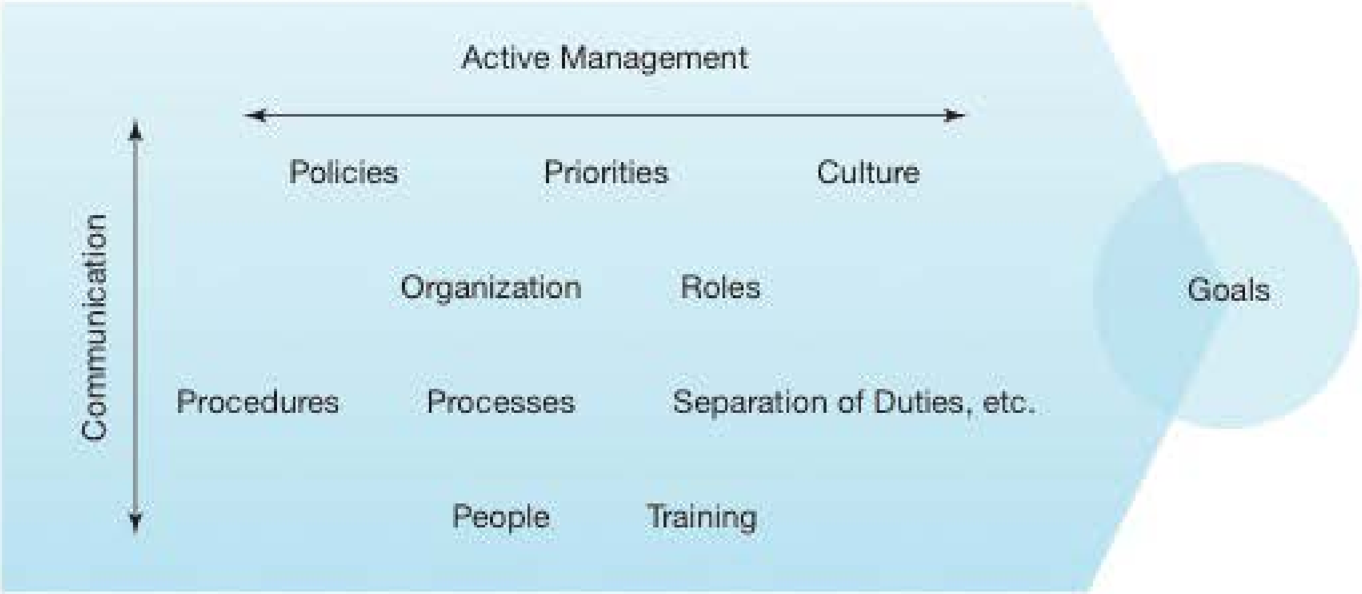

FIGURE A-15 Organizational System Security employee to a team, this will require specialized knowledge for giving them appropriate permissions.

Procedures, Processes, and the Separation of Duties The organizational system will have recurring work patterns. In a few cases, these patterns are **procedures**, which are so simple and well-defined that they can be automated. This has been the traditional focus of information systems. However, most recurring work patterns in organizations are more complex and less well-defined. They are **processes**. A typical process is new product development, which takes place across multiple organizational units over a considerable period of time. Other common processes are employee performance reviews and the creation of annual plans. The fact that processes are not completely defined means that they are difficult both to support and to secure. Many aspects of securing procedures and processes have been explored in the accounting and finance literature. One example is the **separation of duties**. In the granting of exceptions to corporate rules, for example, it is common to forbid the person requesting an exception from authorizing the exception as well.

Organizational Structure and Roles An organizational system has an organizational structure in which individual people have specific reporting patterns and specific roles. Their training must include security aspects of both. On the positive side, if people understand their place and role in the system, they can be given responsibility and accountability for their work. If they are effective, they will discover actions they must take to achieve good security. At least they know that it is up to them to do so.

Policies, Priorities, and Culture At a broad level, there are policies, priorities, and culture. These must have appropriate security content and must be taught to all members of the organizational system. Having a strong security culture is important because if security culture is strong, people are likely to figure out how to work securely even when security flaws are present. When security culture is weak, the best technical, procedural, and other protections are likely to be circumvented. It is important to realize that in most organizations, the security culture is only moderate and that some circumvention must be expected.

Communication Communication must be constant and effective. If there is effective communication among everyone, problems are likely to be resolved and mistakes discovered. In the attack on Pearl Harbor, Admiral Kimmel believed that the Army had the capacity to protect his ships at anchor. The Army did not. In fact, it was focusing on an entirely different problem, the sabotage of Army airplanes. Admiral Kimmel relied on the Army without sufficient justification because of lack of deep communication between the Army and Navy.

Active Management Above everything is active management. Unless the management personnel at the top of the organizational system are engaged and effective, the system has little chance of being secure. They must create and communicate security requirements and disciplines effectively, they must ensure that problems in security elements are identified and worked through, and they must actively look for lapses in security. Above all, they must demonstrate security discipline in their own behavior frequently and consistently.

Test Your Understanding

10. a) What is an organizational system? b) Why is it necessary to train employees in security relevant to their jobs? c) Distinguish between procedures and processes. d) which are harder to secure? Why? e) Explain why the separation of duties may be necessary and how it is done. f) How does establishing a clear organizational structure and roles tend to lead to better security? g) Which is hardest to create—good security policies, priorities, or culture? h) Why is security culture important? i) What is the value of communication in good security? j) Why is active management necessary for good organization system security?

POLICY-BASED SECURITY MANAGEMENT

We have discussed the importance of security planning and major security principles. Now we look at how plans are implemented in well-run organizations.

Policies versus Implementation

Policies (What to Do) The heart of security management is the creation and implementation of security policies by high-level policy makers. Figure A-16 illustrates how policies should be used. **Security policies** are broad statements that specify what should be accomplished in terms of security. They specify "what to do" rather than "how to do it." For example, the security policy might be, "All information on user PCs must be strongly encrypted." Policy makers have the overview knowledge that operational people do not have. For instance, policy makers may know that new compliance regimes create serious liabilities unless all mobile phone data is strongly encrypted. Operation-level people may not realize this. Or, policy-level people, who scan the horizon broadly, may realize that a serious new threat can only be stopped with the encryption of mobile phone data. Again, operation-level people may not realize this. It is not that operational-level employees are incompetent. They have extensive operational expertise. However, they do not have the broad view that planners have.

	Policies	Implementation
Characterized by:	What to do	How to do it
More formally	Broad statements about what should be done	Specific decisions about how to implement the policies
Example	Policy that all mobile phones must be encrypted strongly.	Decision to use functions built into mobile phone operating systems or added encryption software with added features.
Created by	Planners with superior knowledge of regulatory requirements and trends in attack patterns.	Implementers with superior knowledge of security and specific technologies.
Problems avoided	Implementation actions inconsistent with the broad environment.	Micromanagement by policy makers.

FIGURE A-16 Policy-Based Management

Policies are broad statements that specify what to do, not how to do it,

Implementation (How to Do It) Note that the policy does not specify *how* the encryption should be done. That is for implementers to decide within the dictates of the policy. Implementers may decide to use the implementation of encryption built into mobile phone operating systems, or they might adopt an encryption product that offers superior features, such as the ability to decrypt encrypted PCs using a central database that stores encryption keys securely. This would prevent the user's loss of the encryption key from rending the data useless and permit forensic analysis of the PC if phone user is suspected of misuse. These decisions should be left to implementers because they have superior knowledge of security software for mobile phones. Separating policies from implementation prevents senior security policy professionals from micromanaging implementers and forcing implementers to use suboptimum choices for mobile phone encryption.

Implementation decisions specify how to do it.

Policies, in turn, prevent the implementers from overlooking mobile phone encryption because they do not realize that encryption is mandatory under a particular compliance regulation or necessary in light of a new threat

The key point is that the separation of policies from implementation uses the specific different strengths of both policy makers and implementers.

Policy makers have the overview knowledge that operational people do not have. Implementers know about specific technologies and the local situation that policy makers do not. Separating policies from implementation uses the specific different strengths of both policy makers and implementers.

Test Your Understanding

11. a) What is a policy? b) Distinguish between policy and implementation. c) Why do companies separate policies from implementation?

Oversight

Figure A-17 notes that policy makers cannot merely toss policies and implementation guidance out and ignore how implementation is done. It is essential for management to exercise **oversight**, which is a collection of methods for ensuring that policies have been implemented appropriately in a particular implementation. Policies do not give protection by themselves. Nor do unexamined implementations. Protection is most likely to be effective when excellent implementation based on excellent policies is subject to strong oversight.

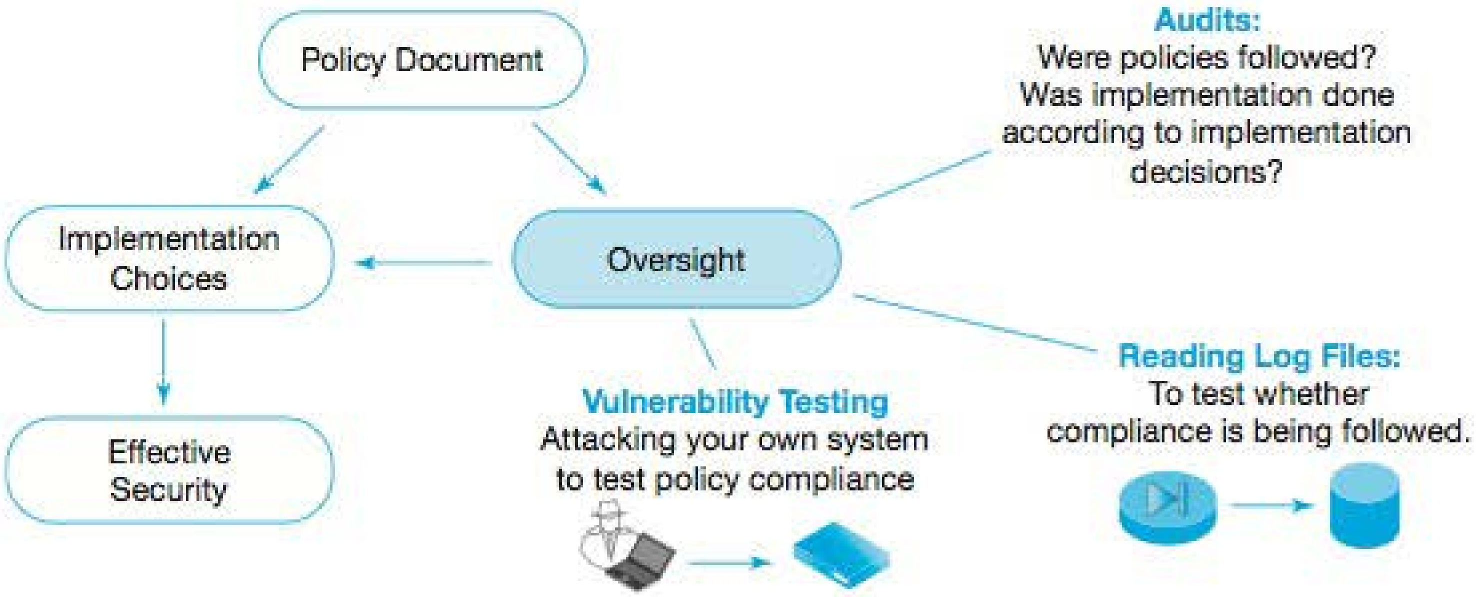
Policy Document
Audits:
Were policies followed?
Was implementation done according to implementation decisions?
Implementation Choices
Oversight
Effective Security
Vulnerability Testing
Attacking your own system to test policy compliance
Reading Log Files:
To test whether compliance is being followed.

FIGURE A-17 Oversight

Oversight is a collection of methods for ensuring that policies have been implemented appropriately in a particular implementation.

Audits One form of oversight is the audit. An **audit** samples and analyzes actions taken during development (and use) to ensure that policies are being implemented properly. Note that an audit only *samples* actions. It does not look at everything, which would be impossible. However, if the sampling is done well, the auditor can issue an opinion on whether a policy is being carried out appropriately based on well-considered data.

An audit samples actions taken within the firm to ensure that policies are being implemented properly.

Reading Log Files Another form of oversight is **reading log files**. Whenever users take actions, their actions should be recorded in log files. Reading log files can also reveal whether policies were implemented successfully. Of course, if these log files are not read, they are useless. Log files should be read daily or even several times each day. Nobody enjoys reading log files, so reading log files is an important thing to audit.

Reading log files can reveal improper behavior.

Vulnerability Testing The most important oversight mechanism is vulnerability testing. Simply put, **vulnerability testing** is attacking your systems before attackers do. Vulnerability testing identifies weaknesses so that you can fix them before they are exploited by attackers. Nearly every implementation will have security vulnerabilities, so testing should be mandatory before an implementation is used operationally. In addition, fixing one security vulnerability may create unexpected vulnerabilities in other parts of the system. Systemwide vulnerability testing should be done regularly and frequently.[11]

Vulnerability testing is attacking your systems before attackers do so that you can identify weaknesses and fix them before they are exploited by attackers.

Test Your Understanding

12. a) Why is oversight important? b) List the three types of oversight described in the text. c) What is vulnerability testing, and why is it done?

Implementation Guidance

In many cases, the policy maker will only specify the broad policy, such as "encrypt all mobile data strongly." However, in some cases, policy makers also will provide **implementation guidance**, which consists of instructions that are more specific than policies but more general than implementation (see Figure A-18). An example of implementation guidance is to specify that strong encryption for confidentiality requires keys that are150-bits or longer. This implementation guidance ensures that when implementers use encryption, they will specify encryption that policy makers have deemed to be strong.

Implementation guidance consists of instructions that are more specific than policies but less specific than implementation.

There are two forms of implementation guidance: standards and guidelines.[12]

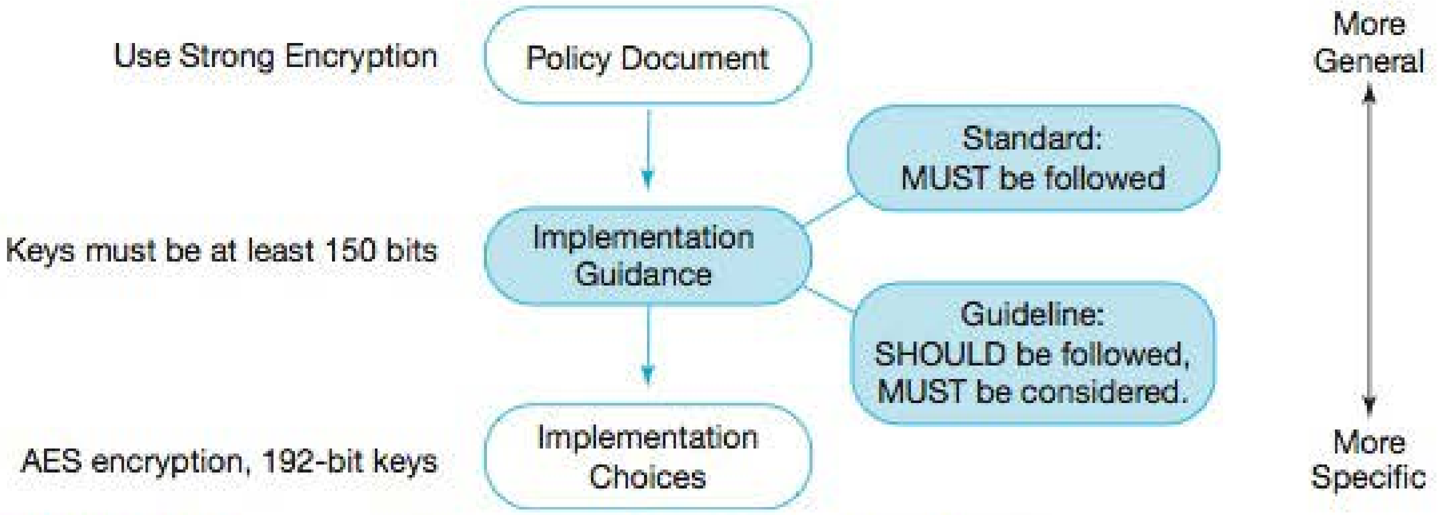
Use Strong Encryption
Policy Document
More General
Standard: MUST be followed
Keys must be at least 150 bits
Implementation Guidance
Guideline: SHOULD be followed, MUST be considered.
AES encryption, 192-bit keys
Implementation Choices
More Specific

FIGURE A-18 Implementation Guidance: Standards and Guidelines

[11] Before doing a vulnerability test, the tester must have explicit written permissions for each test based on a detailed description of what will be done and what damage might be done accidentally. Vulnerability testers who do not take these precautions have been accused of making malicious attacks. This has resulted in firings and even jail terms.

[12] When do firms use guidelines instead of standards for implementation guidance? They use guidelines for situations that are not amenable to black-and-white rules. Encryption strength is relatively easy to specify. The quality of work experience requires human judgment.

Standards Standards *MUST* be followed. If a compliance regulation governing the system requires at least 128-bit keys, this should be specified as a standard for the system being developed. The implementer might not know of the regulatory requirement, so mandatory implementation guidance is justified.

Standards are mandatory directives that must be followed.

Guidelines In contrast, **guidelines**[13] are directives that *SHOULD* be followed. This gives the implementer guidance but also some leeway in deciding whether to follow the guidance if there is good reason not to. For example, guidance might be, "Use SHA-512 hashing in authentication where feasible." The developer receiving the SHA-512 guideline may know that SHA-512 authentication is not possible because the technology in use only allows 384-bit hashing. After due consideration, he or she may use SHA-384 in this situation.

But *Considering* Guidelines Is Mandatory The fact that guidelines are not mandatory does not mean that implementers can ignore guidelines. They *must consider* them carefully. For example, a guideline that security staff members should have three years of security work experience indicates that someone hiring a security staff member must consider that having at least three years of experience is an expectation. If the person doing the hiring selects someone with only two years of security work experience, he or she should have a very good reason for doing so, typically in the form of offsetting relevant experience in other IT jobs.

Guidelines are implementation guidance directives that should be followed but can be dispensed with if circumstances warrant it.

Following guidelines is optional, but seriously considering guidelines is mandatory.

Test Your Understanding

13. a) Compare the specificity of policies, implementation guidance, and implementation. b) Distinguish between standards and guidelines. c) Which must be followed? d) Must guidelines be considered?

Policy-Based Centralized Management

Given the critical importance of policies, it is important to have a way to disseminate and enforce policies. For example, many companies have dozens or even hundreds of firewalls in their network. It would be easy to accidentally misconfigure a few of the firewall access control lists inconsistently with policies. Figure A-19

[13] In the *Pirates of the Caribbean* movies, there was a running joke that the Pirate Code is "more of a guideline, really."

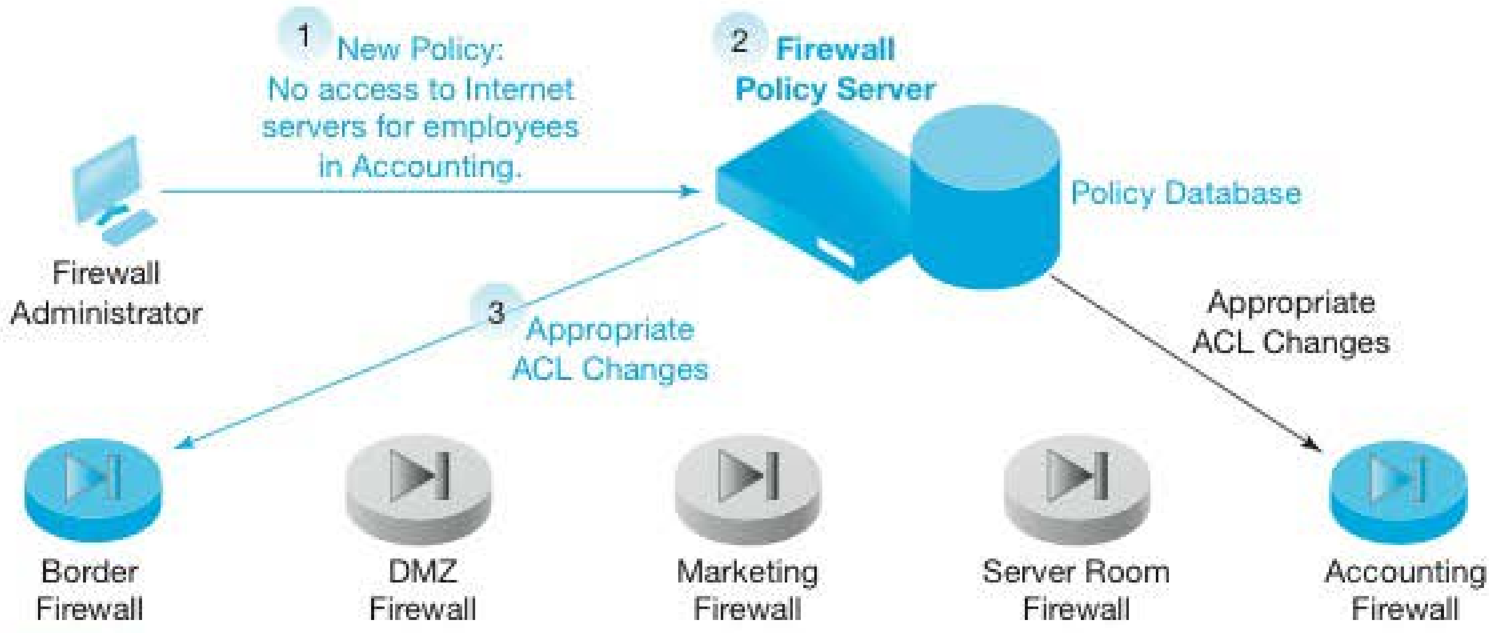

FIGURE A-19 Centralized Firewall Management System

shows a system designed to reduce such errors. All major firewall vendors offer these systems.

The firewall administrator creates high-level policies. In Figure A-19, the firewall policy is that no IP address in the accounting department may access an external webserver. The firewall administrator sends this policy to the **firewall policy server**, which places the policy in its **policy database**.

The firewall policy server then modifies the firewall access control lists (ACL) of affected firewalls. In the figure, the only affected servers are the border firewall and the accounting firewall. The firewall policy server then pushes these ACL changes to the affected firewalls. Note that the policy server does not simply push policies out to firewalls. It creates detailed ACL changes and moves these changes out to the firewalls.

Separating firewall policies from ACL rules is a good example of policy-based security. The firewall administrator sets a high-level policy. The firewall policy server converts this policy into individual firewall ACL rules. The firewall policy server will not make human mistakes such as forgetting to add a particular rule to a particular firewall. Furthermore, if there is a question about a particular firewall rule on a particular firewall, the firewall administrator can ask what policy it implements. Policies are usually easier to understand than specific firewall rules.

The process of creating specific firewall rules for different firewalls based on firewall policies is based on technology rather than magic. Consequently, errors are inevitable. This makes vulnerability testing mandatory. The test suite must include both things that should be forbidden to ensure that they are forbidden and everything else that should not be forbidden to ensure that it gets through.

Test Your Understanding

14. a) Distinguish between firewall policies and firewall ACL rules. b) After a firewall administrator sends a policy to the policy server, what does the policy server do? c) Which is easier to understand—a firewall policy or a firewall rule?

RESPONSE

The final phase in the plan–protect–respond cycle is responding to security **breaches**, which are also called **incidents** and **compromises**. These protection failures take place in every firm all too frequently. As noted earlier, response is responding to incidents according to plan (Figure A-19). If no plan is in place, losses will be more than they need to be.

Response always takes place under conditions of stress, and people do not think well under stress. In addition, there are almost always time pressures to stop the attack and restore systems to running order. These pressures expand in major attacks by sophisticated hackers. Time is always of the essence because time gives the attacker more opportunity to do damage and hide his or her efforts from the security staff. However, rushing response will backfire if the security team misdiagnoses the root cause of the attack. This will result in time-consuming backtracking once the misdiagnosis is discovered.

There are two keys to effective fast response. The first, of course, is to have a plan for how to respond to compromises. A plan helps focus the security team on key steps and prevents steps from being overlooked. A good plan focuses the team in a way that benefits from experience. In fact, response is best defined as "responding to compromises according to plan." Specific situations will require modifying plans during the response process, but modifications within a plan are far more likely to succeed than unplanned reaction.

Response is best defined as "responding to compromises according to plan."

The second key is practice. When a football team gets a new play, they try it first in practice, usually with disastrous results. They then see what went wrong, try it again,

Responding to breaches, also called incidents or compromises

Response time and accuracy are both critical

- Goal is to find the root cause of the failure, stop the attack quickly, and restore operation
- The attack must be stopped quickly because the attacker will use delays to increase damage and hide
- The analysis must also be correct, or much time will be spent backtracking

Keys to Success

- Have a good plan
  - Response is "responding to compromises according to plan"
  - People do not think well under stress
  - No plan will be perfect, but it is best to improvise within a good plan
- Practice
  - Speed and accuracy require rehearsal
  - Only after many repetitions can they execute the play correctly and rapidly

FIGURE A-20 Response (Study Figure)

and make new mistakes. Only after many repetitions can they execute the play correctly and rapidly in a game. Response to security incidents is the same. Having a plan that has not been rehearsed extensively is useless.

Test Your Understanding

15. a) What are the three names for successful attacks? b) What are the two keys to fast and effective response?

Normal Incidents

Most security incidents are relatively minor, such as several PCs becoming infected with a virus. Security teams see these **minor incidents** frequently. Even more common are **false alarms**, which are apparent compromises that turn out to be legitimate actions. Both are treated the same way initially because there is no way of knowing in the beginning which is occurring (Figure A-20).

Most importantly, these are IT and security issues. Both minor incidents and false alarms are handled by the on-duty IT and security staff as part of their normal work. There is rarely a need to call in an external consulting company.

	Normal Incidents	Major Incidents
Example	Normal incidents: Malware compromise of a few dozen PCs False alarms: Nondamaging but time-consuming	Breaking into the company's financial system Terrorist attack that closes one data center
Corporate-wide consequences	None	Medium to Very High
Should be viewed as	An IT and security issue	A business issue
Are handled by	The on-duty IT and security staff	The Computer Security Incident Response Team (CSIRT)
Led by	The IT manager on duty	A senior business manager
Team includes IT and security people	Yes	Yes
Team includes other line managers	No	Legal, public relations, etc. Only PR should talk to outsiders Unless the CEO steps into that role
Calling in an external consulting company	Rarely	Most of the time

FIGURE A-21 Normal Incidents versus Major Incidents

Minor incidents are handled by the on-duty IT and security staff.

There are no broad corporate issues involved. There is no need to call in line managers from other departments during the incident, although affected departments should be kept informed of the situation.

For both minor incidents and false alarms, occurrences are frequent enough to give responders the experience they need to handle them. There usually is no need for additional rehearsals.

Test Your Understanding

16. a) What are the two types of normal incidents? b) Who handles normal incidents? c) Why do normal incidents typically require no rehearsal?

Major Incidents

Major incidents are those that cannot be handled by the on-duty security staff. They require special handling because major incidents have business implications that beyond IT, sometimes involving the entire company. The Target breach is an example of a major incident. Damage to the firm and its reputation were high, and top management needed to be involved. Target also brought in security consultants whose specialized knowledge and experience allowed the company to identify root causes and take appropriate actions (although it sometimes failed to do so).

Major incidents are those that cannot be handled by the on-duty security staff.

They require special handling because major incidents have business implications that extends beyond IT, sometimes involving the entire company.

When major incidents occur, companies activate their **computer security incident response team (CSIRT)**. The team consists of members of the IT and security staff, but it goes well beyond that. In fact, the CSIRT is headed by a high-level business manager, not by a member of the IT and security staff. During response to a major incident, technical decisions are also business decisions. For example, taking the company's e-commerce server offline for five hours is a decision that must be considered from all business and technical viewpoints.

During response to a major incident, most technical decisions are also business decisions.

The CSIRT will have several other business managers on the team. One will certainly be from the Legal Department. Major compromises always involve complex legal questions and decisions. Input from Legal must be built into the CSIRT's work at every step.

Another critical department is Public Relations. Rumors and unauthorized pronouncements can be extremely damaging, and they often present information that is plainly wrong. CSIRT members and other employees must not spread rumors or speak to the press. Only the public relations director should speak for the company—unless the CEO elects to do so. Whoever speaks for the company must have excellent information about what is transpiring.

Typically, a CSIRT calls on external security experts to speed the work and bring in higher-level expertise. These are typically consulting firms with experience in large breaches. These consultants should be held on retainer and should be familiar with the company's IT system so they can provide immediate help.

Test Your Understanding

17. a) What is the definition of a major incident? b) What group handles a major incident? c) What characteristics should its leader have? d) What major departments will almost always be involved? e) Who is the only person who should talk about the incident with the media and other outsiders? f) Who may elect to speak instead?

Rehearsing for Major Incidents

It is difficult to rehearse for major incidents because most team members view their participation as time consuming and irrelevant to their main work. However, the CSIRT must rehearse a few times each year. Rehearsal improves reaction time and quality. It also integrates new members to the team. Given normal organizational changes, new membership is almost always present.

Rehearsal Improves Response Time and Quality

- Rehearsals are mandatory

Desktop Exercises

- Sit around a table
- Presented with a situation
- Walk through the analysis step by step
- Each member tells what he or she would do and what help is needed from others
- The moderator throws in surprises
- Postmortem on what worked, failed, and was learned

Live Exercises

- Go through a breach step by step
- Actually do each step on the system
- Expensive compared to desktop exercises
- Can find things desktop exercises never can
- E.g., The planned steps for moving control to another data center fail to transfer the security credentials needed to run the system there

FIGURE A-22 Rehearsing for Major Incidents (Study Figure)

Desktop Exercises The least costly way for CSIRTs to rehearse is the **desktop exercise**, in which the team members sit in a room and run through scenarios. In these exercises, the moderator leads the team through a situation step by step. At each step, the members discuss what each must do, how they will do it, and what help they need from others. A good moderator adds new facts to the exercise abruptly, changes previous data, and does other things to make desktop exercises more realistic.

Live Exercises Desktop exercises are valuable, but they are not as good as **live exercises**, in which team members do the work at each step on the live system or a close facsimile. Live exercises reveal that confident statements made during desktop exercises are far more problematic than the team members believe. They also get to a level of detail that exposes practical problems. For example, one live exercise involved a terrorist attack that filled the company's main server room with toxic fumes. The personnel were evacuated safely, and the company switched the server room's processing to another data center and moved the staff to this center. Only when they got there did they realize that a critical list of passwords was in the old server room, which was now impossible to enter. Although a firm will undertake more desktop exercises than live exercises, live exercises are still needed occasionally. No football coach would diagram a play on the backboard, have each player discuss what he or she would do during the play, and use the play in the next game without practice.

Test Your Understanding

18. a) Distinguish between desktop exercises and live exercises in CSIRT rehearsals. b) Why are desktop exercises important? c) Why are live exercises necessary?

Real-Time Fail-Over

If there is a disruption in a major data center, critical corporate transaction processing will grind to a halt. Years ago, it was common to recommend that a firm maintain a "hot site" with equipment, software, and power. If the main site failed, the hot site would be turned on and personnel transferred to it. However, this took days to implement. By then, the company could be near bankruptcy.

Today, the norm is to use **real-time fail-over** with **synchronized data centers**. Figure A-23 illustrates this situation. The company has two data centers. One is in New York City. The other is in Denver. Each handles four of the company's eight major applications. They are connected by an ultra-high-speed transmission link. This allows them to copy and store one another's data in real time. Application and operating system software changes to running programs also are transferred to the corresponding backup versions of the programs at the other data center in real time. Both data and software, in other words, are fully synchronized.

If the Denver data center fails, the New York data center can take over immediately. This is called **fail-over**. Immediate response is a crucial advantage. Downtime is enormously expensive. Fail-over with synchronized data centers can limit downtime to seconds.

The downside of fail-over with synchronized data and programs is cost. Each data center needs to have extra capacity to handle the increased workload after fail-over.

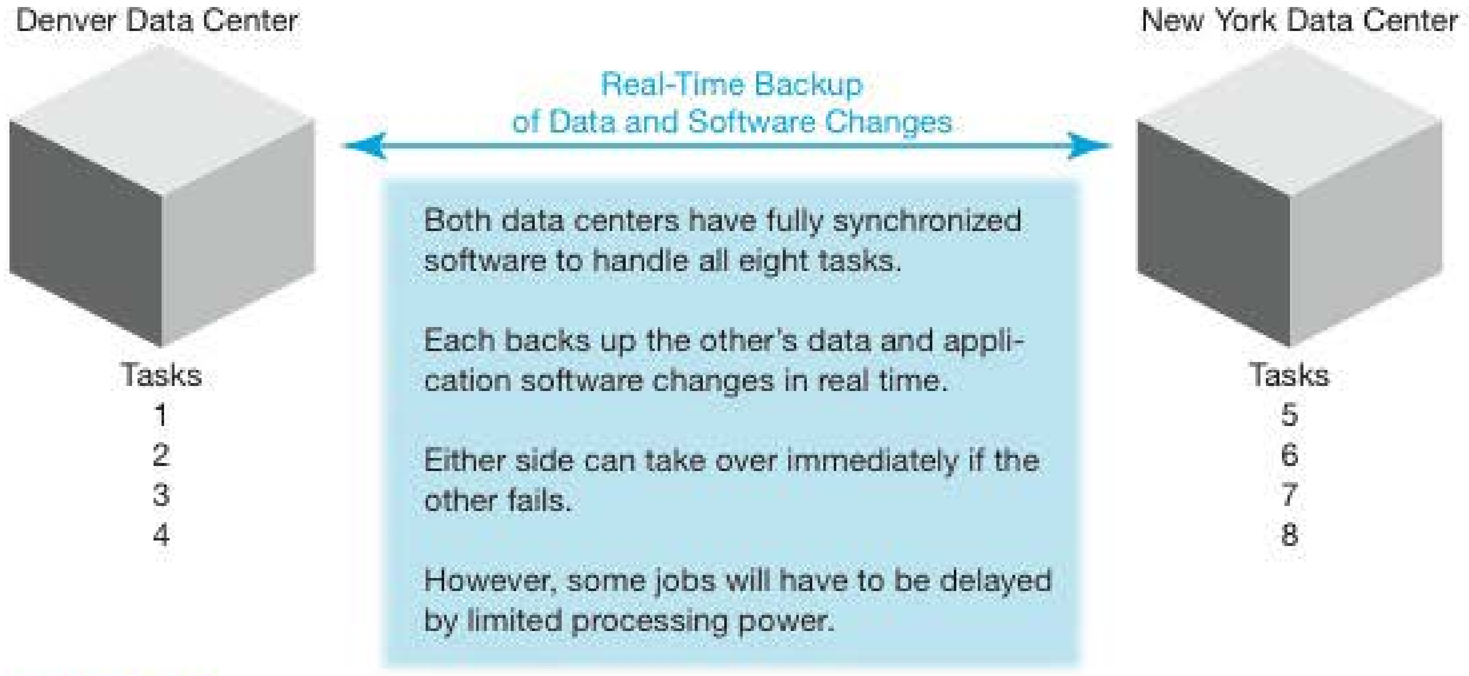
Denver Data Center
New York Data Center
Real-Time Backup
of Data and Software Changes
Both data centers have fully synchronized software to handle all eight tasks.
Each backs up the other's data and application software changes in real time.
Either side can take over immediately if the other fails.
However, some jobs will have to be delayed by limited processing power.
Tasks
1
2
3
4
Tasks
5
6
7
8

FIGURE A-23 Real-Time Fail-Over with Synchronized Data Centers

More directly, synchronizing data requires a massive transmission link between the two sites. Although falling networking costs have made immediate fail-over economically feasible, they have not made it cheap. Setting up and maintaining this link is a central issue in network design, implementation, and operation. Most obviously, there must be backups in place for failures in this ultrafast data pipe.

Test Your Understanding

19. a) What is the advantage of real-time fail-over with synchronized data centers? b) Why is it expensive?

Intrusion Detection Systems (IDSs)

Your car has a lock to keep people out. Some cars also have car alarms that warn you when someone is trying to break into your car. (Or when a cat walks across the hood of your neighbor's car.) In the security world, firewalls and other countermeasures are like car locks or security guards. They stop people who are trying to break in.

Intrusion Detection Systems Security professionals also need something like car alarms to tell them when someone is trying to hack the system. These are **intrusion detection systems (IDSs)**. When suspicious things happen on a network, IDSs note them and create alarms for the security staff. The security staff can then block the attacker. Otherwise, the assailant can attack a resource repeatedly until he or she succeeds. If someone is trying to break into your car, you certainly want to know about it when they start to attempt the break-in.

The problem with firewalls is that they only look at provable attack packets. If a packet is suspicious, the firewall ignores it. IDSs, as Chapter 4 noted, specifically target suspicious attack packets, raising an alarm if they are serious.

The problem with IDSs is the same problem that car alarms have. Many of their alerts are false alarms. Out of a hundred notifications, only one or two may signal a real

	Firewall	Intrusion Detection System
Automobile Analogy	Security guard who prevents a break-in of your car	A car alarm that notifies you if someone is attempting to break in
Prevents a Break-In	Yes	No
Requires Near-Certainty to Act	Yes	No
Target	Provable (definite) attack packets	Suspicious packets
Logging	Logs dropped packets; may log all packets	Collects a broad spectrum of event data and puts it into an integrated log file for analysis
Actions	Passes or drops packet	Generates alarms, allows security administrator to query the integrated log file to understand patterns
Major Problem	Failure to stop suspicious packets	Many false alarms
Vigilance is a Major Problem	No	Yes

FIGURE A-24 Firewall versus Intrusion Detection System (Study Figure)

problem. Security professionals usually find themselves unable to follow up on each alarm. In other cases, many check alarms cursorily, perhaps missing a real attack.

The Problem of Vigilance It is possible to "tune" IDSs to ignore many alarms that do not make sense. For example, if you have no Linux devices outside your data center, you may disable Linux-only alarms. In addition, IDSs normally report events with some indication of likely severity. Many security staffs only follow up on high-level security alerts. To go back to the car alarm case, the owner might reduce the sensitivity of the alarm. (This does not seem to be possible with my neighbor's car alarm.) Even in the best situations, however, there will be far more false alarms than dangerous incidents. Vigilance tends to flag under such conditions, and it is all too easy when vigilance flags to dismiss an important alert. Nearly all the time, ignoring an alarm will be both the right choice and a time-saving measure. This creates powerful negative reinforcement to ignore alerts.

Log Files Each IDS constantly records information about its device's operation. For example, a router may report every packet's source and destination IP address and other packet header information. It may also report errors. Each event's data is stored as a **record**.

Distributed and Integrated Log Files A company may have hundreds of IDSs on individual devices. However, attack analysis requires combining the IDS records on different devices into a single combined log file. This **integrated log file** permits event correlation, which is analyzing events from different devices that

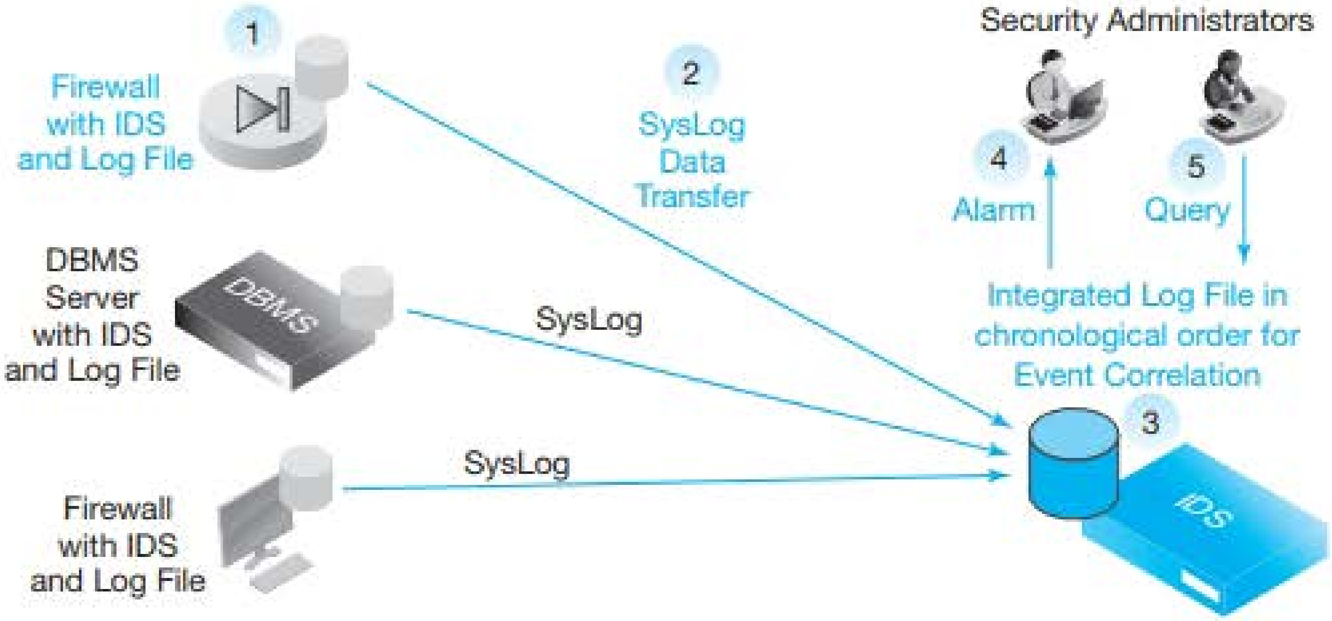
1
Firewall with IDS and Log File
2
SysLog Data Transfer
Security Administrators
4
Alarm
5
Query
DBMS Server with IDS and Log File
DBMS
SysLog
Integrated Log File in chronological order for Event Correlation
3
SysLog
Firewall with IDS and Log File
IDS

FIGURE A-25 Distributed IDSs and Integrated IDS Log Files in Chronological Order for Event Correlation

together present a clearer picture of what happened. It should end in **root cause analysis** to identify the attack and attacker perfectly. Figure A-25 notes that this transmission follows the **SysLog standard**.

Alarms Once the data are on the Integrated Log File, they are in useable form. In Figure A-25, the security administrator on the left is receiving an alarm from the system-wide IDS. The alarm gives a brief description of the situation. It is probably wrong but should not be ignored.

Querying the Integrated Log File After receiving an alarm, a security administrator will conduct queries on the integrated log file to get a better understanding of patterns that touched off the alarm.

Figure A-26 shows the results of a query. The query returns the requested event records. For simplicity, irrelevant entries have been removed in the figure.

The first three log file entries tell a combined story. First, a packet from Host 1.15.3.6 goes to Host 60.3.4.5. Internal host IP addresses are in the range 60.x.x.x. Therefore, Host 1.15.3.6 is an external host. Second, Host 60.3.4.5 records a failed login for the account of Lee. Third, a packet goes from the internal host to the external host.

These packets need to be interpreted. An obvious interpretation is that someone at Host 1.15.3.6 seems to have attempted to log into Host 60.3.4.5 under the username Lee. The login attempt failed, and a notification was sent to the external host. (ACKs have been removed.) This might be the sign of an attack, or it may simply be Lee forgetting his or her password or typing it incorrectly. Following this logic, the next six records indicate that there is one more failed attempt and then a success.

This may be an attack using password guessing, or it may be normal human memory failure or poor typing. More analysis is needed. However, even the first nine records contain potentially useful hints. There is some time between notification of failure and the next login attempt, so the actor at the external host appears to be human. (An automated password cracker would send the next guess much faster.) This is more evidence for the "bumbling human" interpretation. Are we convinced that this pattern

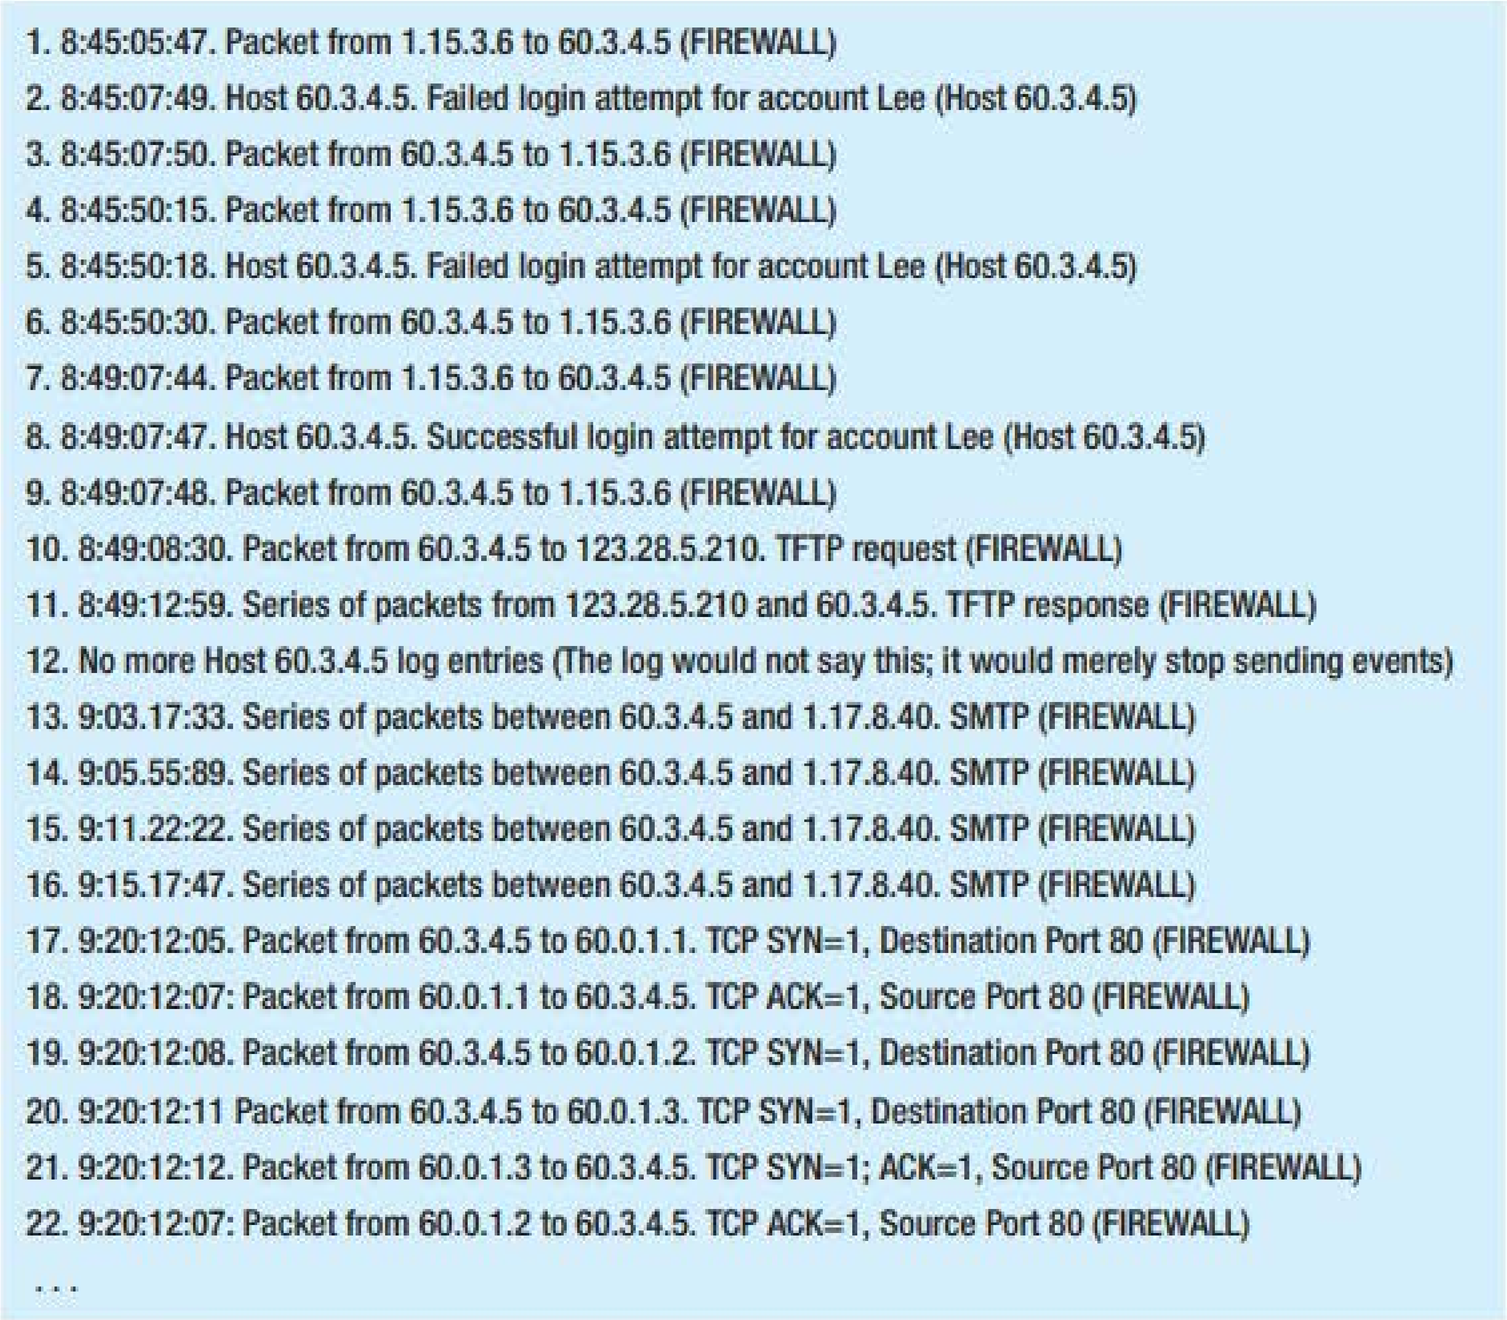

```
1. 8:45:05:47. Packet from 1.15.3.6 to 60.3.4.5 (FIREWALL)
2. 8:45:07:49. Host 60.3.4.5. Failed login attempt for account Lee (Host 60.3.4.5)
3. 8:45:07:50. Packet from 60.3.4.5 to 1.15.3.6 (FIREWALL)
4. 8:45:50:15. Packet from 1.15.3.6 to 60.3.4.5 (FIREWALL)
5. 8:45:50:18. Host 60.3.4.5. Failed login attempt for account Lee (Host 60.3.4.5)
6. 8:45:50:30. Packet from 60.3.4.5 to 1.15.3.6 (FIREWALL)
7. 8:49:07:44. Packet from 1.15.3.6 to 60.3.4.5 (FIREWALL)
8. 8:49:07:47. Host 60.3.4.5. Successful login attempt for account Lee (Host 60.3.4.5)
9. 8:49:07:48. Packet from 60.3.4.5 to 1.15.3.6 (FIREWALL)
10. 8:49:08:30. Packet from 60.3.4.5 to 123.28.5.210. TFTP request (FIREWALL)
11. 8:49:12:59. Series of packets from 123.28.5.210 and 60.3.4.5. TFTP response (FIREWALL)
12. No more Host 60.3.4.5 log entries (The log would not say this; it would merely stop sending events)
13. 9:03.17:33. Series of packets between 60.3.4.5 and 1.17.8.40. SMTP (FIREWALL)
14. 9:05.55:89. Series of packets between 60.3.4.5 and 1.17.8.40. SMTP (FIREWALL)
15. 9:11.22:22. Series of packets between 60.3.4.5 and 1.17.8.40. SMTP (FIREWALL)
16. 9:15.17:47. Series of packets between 60.3.4.5 and 1.17.8.40. SMTP (FIREWALL)
17. 9:20:12:05. Packet from 60.3.4.5 to 60.0.1.1. TCP SYN=1, Destination Port 80 (FIREWALL)
18. 9:20:12:07: Packet from 60.0.1.1 to 60.3.4.5. TCP ACK=1, Source Port 80 (FIREWALL)
19. 9:20:12:08. Packet from 60.3.4.5 to 60.0.1.2. TCP SYN=1, Destination Port 80 (FIREWALL)
20. 9:20:12:11 Packet from 60.3.4.5 to 60.0.1.3. TCP SYN=1, Destination Port 80 (FIREWALL)
21. 9:20:12:12. Packet from 60.0.1.3 to 60.3.4.5. TCP SYN=1; ACK=1, Source Port 80 (FIREWALL)
22. 9:20:12:07: Packet from 60.0.1.2 to 60.3.4.5. TCP ACK=1, Source Port 80 (FIREWALL)
. . .
```

FIGURE A-26 Query Results from the Integrated IDS Log File (Irrelevant Entries Omitted)

indicates an attack? However, if we follow this up by analyzing the rest of Figure A-26, we may change our mind. We have only seen 9 entries. There are 22 in total, so we still have work to do.

Test Your Understanding

20. a) What do firewalls do that IDSs do not? b) What do IDSs do that firewalls do not? c) Compare what is logged by IDSs and firewalls. d) Why are IDS false alarms a problem? e) What standard do device IDSs normally use to send their log data to the central IDS?

21. a) Continuing the analysis of the IDS query in Figure A-26, scrutinize Events 10 and 11. You need to know that the Trivial File Transfer Protocol is a way to download a file from a remote computer. What do these two records suggest? b) After Event 11, there are no more log entries in the IDS log file for Host 60.3.4.5. What does this suggest? c) If you combine this with what you learned in the first nine records, what do you conclude, at least tentatively?

SOME FINAL PICTURES

FIGURE A-27 Rack Server

FIGURE A-28 Corporate Access Point

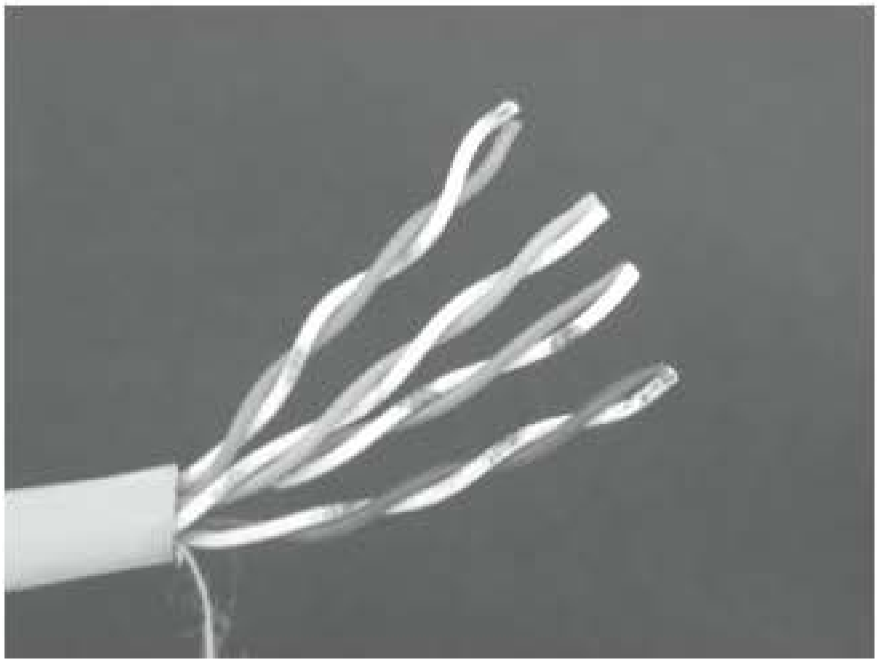

FIGURE A-29 Four-Pair UTP

END-OF-CHAPTER QUESTIONS

Thought Questions

A-1. List the nine security principles named in the appendix.

A-2. For each, say whether it was violated at Target, according to what you read here and in Chapter 4. If so, describe how it is related.

A-3. Some companies abandon their IDSs. Why do you think they do so?

A-4. Some companies are outsourcing the examination of IDS event logs to other companies. a) Why do you think they do so? b) Why was doing so ineffective in the case of Target?

A-5. During the American Revolutionary War, the British landed troops on Long Island. Their goal was to march west to New York City. George Washington arrayed his troops near the western end of Long Island to intercept the British. There were two passes for crossing the mountains between the western end of the island from the British landing point. Washington put half his troops at each pass. However, Loyalists on Long Island knew that there was a third, smaller pass through the mountains; they told the British. The British decided to attack through that smaller gap, although it was a more difficult route. Washington learned of this third route just before the attack. He positioned some riders there to give warning if the British took that more difficult road. Unfortunately, the British captured Washington's scouts. They descended without warning on Washington's flank, and the battle was a rout. It almost ended the Revolutionary War. That night, however, the American troops quietly retreated to the western end of the Island, then evacuated in boats to New York City. When the British rose the next morning, Washington was gone—defeated and chastised but with his army mostly intact and a bitter lesson learned. What security principle caused this failure? Justify your answer.

A-6. Castles are often surrounded by moats or other protections that will forestall attacks or reduce the speed of attacks. Then come thick walls that must be breached. If these fail, the defenders fall back to an inner keep with more defenses. If attackers manage to break into the inner keep's doors, they must ascend spiral "murder steps" that are uneven and require concentration that attackers need to avoid being killed by defenders. These steps rise counterclockwise so that the attacker's right arm is next to the wall, making it difficult for most knights to swing their swords effectively. What principle do these protections embody? Justify your answer.

A-7. Edward Snowden wanted files that he did not have permissions to see. He asked another employee who did have permissions to show him a file. The other employee had Snowden walk away while the person logged in and downloaded the file. Snowden looked at the file for the particular paragraph he "needed to do his job." Snowden then walked away, and the other employee logged out of the system. Only then did Snowden come back. What the other employee did not know was that Snowden had installed a key logger on his computer. It had captured the other employee's login credentials. Snowden used these credentials to log back in and steal massive numbers of files. What security principles were violated? Justify your answer.

A-8. For her dorm room, a student bought a 20-pound safe for $300 to store her laptop, tablet, and phone when she is not using them. a) What principle should have been involved in the purchase? b) Do you think it was justified? c) What other security principles may be relevant?

A-9. Bob Eckert had a heavily automated home. One night, while he was watching TV, the television set shut down, various lights started blinking wildly, and other weird device behavior ensued. What security planning principle was probably violated?

Log File Analysis Questions

A-10. Interpret lines 13 through 16 in Figure A-26.

A-11. Interpret lines 17 through 22 in Figure A-26.

A-12. Do you think an attack is happening? Justify your answer based on what the log file has revealed.

A-13. Do you think that Lee is the culprit? Weigh the evidence for and against Lee's guilt.

Perspective Questions

A-14. What was the most surprising thing you learned in this chapter?

A-15. What was the most difficult thing for you in this chapter? Why was it difficult?

This page intentionally left blank

GLOSSARY

2.4 GHz Unlicensed Service Band: Unlicensed frequency band around 2.4 GHz. Used for Wi-Fi, Bluetooth, and other services.

4-Pair Unshielded Twisted Pair (UTP): The type of wiring typically used in Ethernet networks. 4-pair UTP contains eight copper wires organized as four pairs. Each wire is covered with dielectric insulation, and an outer jacket encloses and protects the four pairs.

5 GHz Unlicensed Service Band: Unlicensed radio band around 5 GHz. Used for Wi-Fi and other services.

64-bit modified extended unique identifier (EUI-64): In most IPv6 addresses, the interface ID that specifies a particular device on a subnet is 64 bits long. Typically the 64-bit address is derived from a 48-bit EUI-48 address (formerly called a MAC address). If so, it is a modified extended unique identifier (EUI-64).

802 LAN/MAN Standards Committee: The IEEE committee responsible for Ethernet standards.

802.1X Initial Authentication Mode: An initial authentication mode used in 802.11i. Requires the use of an authentication server. Called enterprise mode by the Wi-Fi Alliance.

802.1X Authentication Server: Authentication server used in 802.1X initial authentication mode.

802.1X Authenticator: In Wi-Fi transmission, the wireless access point.

802.1X Mode: See 802.1X Initial Authentication Mode.

802.3 Working Group: Working group in the IEEE 802 LAN/MAN Standards Committee that creates Ethernet standards.

802.3 MAC Layer Standard: The data link layer standard for Ethernet.

802.11ac: In W-Fi, the fastest physical transmission standard for sale today.

802.11ad: 802.11 physical layer standard in the 60 Ghz unlicensed band. Has a theoretical top speed of 7 Gbps.

802.11ax: Planned successor to the 802.11ac standard. Will be able to accommodate a substantial increase in density—the number of wireless hosts that can be served by an access point.

802.11ay: In the 60 GHz band, the Wi-Fi successor to 802.11ad. Adds MU-MIMO and other improvements, should raise the basic speed to 20 to 30 Gbps and is likely to allow bonding for even higher speeds. Under development.

802.11i: An advanced form of 802.11 wireless LAN security.

802.11n: Version of the 802.11 WLAN standard that uses MIMO and sometimes doubled bandwidth to achieve a rated speed of 100 Mbps or more and longer range than earlier speed standards.

802.1X Port-Based Network Access Control: In Ethernet, a standard for access control on switch ports.

Absorptive Attenuation: In wireless transmission, the attenuation of a signal but water along the way absorbing its signal power. In optical fiber, attenuation due to the absorption of signal strength as a signal propagates.

Access Card: Small card with a magnetic stripe or microprocessor that gives you access to your computer or to a room.

Access Control: Limiting who may have access to each resource and limiting his or her permissions when using the resource.

Access Control List (ACL): An ordered list of pass/deny rules for a firewall or other device.

Access Link: 1) In networks, a transmission line that connects a station to a switch. 2) In telephony, the line used by the customer to reach the PSTN's central transport core.

ACK: See Acknowledgment.

Acknowledgment (ACK): 1) An acknowledgment message, sent by the receiver when a message is received correctly. 2) An acknowledgment frame, sent by the receiver whenever a frame is received; used in CSMA/CA+ACK in 802.11.

Acknowledgment Number Field: In TCP, a header field that tells what TCP segment is being acknowledged in a segment.

ACL: See Access Control List.

ADSL: See Asymmetric Digital Subscriber Line.

ADSL Modem: Modem used in Asynchronous Digital Subscriber line service. Terminates the carrier’s connection.

Advanced Persistent Threat (APT): Attack occurring over a long period of time. The user employs many advanced methods to get deeper and deeper into the target system.

Advanced Research Projects Agency (ARPA): An agency within the U.S. Department of Defense that funded the creation of the ARPANET and the Internet.

Advertisement Message: Bluetooth LE clients periodically send this type of message to advertise their presence.

Ad-hoc wireless network: A self-organizing wireless network.

Aggregate Throughput: Throughput shared by multiple users; individual users will get a fraction of this throughput.

Agility: The ability to rapidly change how the network operates when conditions change.

Alphanumeric: Strictly speaking, letters and numbers. However, often used to refer to all keyboard characters and, often, some control codes.

Alternative States: In physical transmission, a change in a transmission medium that can signal one data pattern that represents a particular bit pattern. Different (alternative) states signal different bit patterns.

Amazon Web Services (AWS): A cloud service provided by Amazon.

American Standard Code for Information Interchange (ASCII): Code for representing all American keyboard characters plus some control codes.

Amplitude: The maximum (or minimum) intensity of a wave. In sound, this corresponds to volume (loudness).

Antenna: A physical structure that transmits radio signals.

Antivirus (AV) Program: Program to remove malware from arriving messages and from the computer’s disk drive.

API: See Application Program Interfaces.

Application Architecture: The arrangement of how application layer functions are spread among computers to deliver service to users.

Application-Aware Firewall: A firewall that can identify and manage the application that creates a stream of packets.

Application Messages: A message sent from one networked application to another over a network.

Application Program Interfaces (API): A standardized interface between programs.

API. See Application Program Interface.

Apps: Small applications created for mobile devices.

APT: See Advanced Persistent Threat.

ARP cache poisoning: Sending false information to a host to place in its ARP cache. This will cause it to send messages to a particular IP address to the wrong data link address and therefore the wrong destination host.

ARP update: A command to tell a host to send messages to a particular IP address to a particular data link layer address. Useful if the data link address is the correct one. Causes the receiving host to send frames to the wrong host if the ARP update is false.

ARPANET: A packet-switched network created by the Advanced Research Projects Agency.

ASCII Code: A code for representing letters, numbers, and punctuation characters in 7-bit binary format.

Asymmetric: Different in two directions.

Asymmetric Digital Subscriber Line (ADSL): The type of DSL designed to go into residential homes, offers high downstream speeds but limited upstream speeds.

Audit: Collecting data about events to examine actions after the fact.

AUP: See Acceptable Use Policy.

Authentication: The requirement that someone who requests to use a resource must prove his or her identity.

Authentication Header: In IPsec, a header that protects part or all of the packet with authentication.

Authoritative DNS Server: DNS server that manages host names for a particular domain.

Authorizations: A rule that determines what an account owner can do to a particular resource (file or directory).

Availability: The ability of a network to serve its users.

AWS: See Amazon Web Services.

Backward-Compatible: Able to work with all earlier versions of a standard or technology.

Base Case: In a risk analysis, the case in which the organization does nothing.

Basic Printer Profile (BPP): Bluetooth profile that allows a device to print wireless to a printer without needing to download a particular printer driver for that printer.

Beacon: Bluetooth LE advertising message that transmits potentially useful information.

Beamforming: In radio transmission, directing energy toward a wireless device without using a dish antenna.

BGP: See Border Gateway Protocol.

Binary Numbers: The base two counting system where 1s and 0s used in combination can represent whole numbers (integers).

Binary Signaling: Digital signaling that uses only two states.

Biometrics: The use of bodily measurements to identify an applicant.

Bits per Second (bps): The measure of network transmission speed.

Bluetooth: A wireless networking standard created for personal area networks.

Bluetooth LE: See Bluetooth Low Energy

Bluetooth Low Energy (LE): New form of Bluetooth designed for low-energy devices such as Internet of Things devices.

Bluetooth Profile: An application layer standard designed to allow devices to work together automatically, with little or no user intervention.

Bluetooth Special Interest Group: The organization that creates Bluetooth standards.

Bonding: See Link Aggregation.

Border Gateway Protocol (BGP): The most common exterior routing protocol on the Internet. Recall that gateway is an old term for router.

Border Router: A router that sits at the edge of a site to connect the site to the outside world through leased lines, PSDNs, and VPNs.

Bot: A type of malware that can be upgraded remotely by an attacker to fix errors or to give the malware additional functionality.

Botmaster: Attacker who control a botnet.

Botnet: A large number of computers infected with bot malware.

BPP: See Basic Printer Profile.

Breach: A successful attack.

Broadband: 1) Transmission where signals are sent in wide radio channels. 2) Any high-speed transmission system.

Broadband Channels: Strictly speaking, a radio channel with large bandwidth. This permits high-speed transmission. More broadly, the term is used for any fast transmission system.

Brute-Force Attack: A password-cracking attack in which an attacker tries to break a password by trying all possible combinations of characters.

CA: See Certificate Authority.

Cable Modem: 1) Broadband data transmission service using cable television. 2) The modem used in this service.

Cable Modem Service: Asymmetrical cable data service offered by a cable television company.

Cable Television: Form of television delivery that distributes signals to the home over coaxial cable.

Cache: A limited amount of memory to hold data for a very short period of time until the device can deal with it.

Caching: In general, storing very temporary information for retrieval. In routing, storing routing decisions for particular IP addresses that were very recently handled instead of going through the whole routing process again.

Carriage Return: Takes the cursor back to the start of the current line.

Carrier: A transmission service company that has government rights of way.

Carrier Ethernet: Ethernet service provided in a MAN or WAN by a carrier to user organizations.

Carrier WAN: Wide area networking service offered by a carrier.

CDN: See Content Delivery Network.

Cell: In cellular telephony, a small geographical region served by a cellsite.

Cellphone: A cellular telephone, also called a mobile phone or mobile.

Cellsite: In cellular telephony, equipment at a site near the middle of each cell, containing a transceiver and supervising each cellphone's operation.

Certificate Authority (CA): Organization that provides public key–private key pairs and digital certificates.

Challenge Message: Message sent by a verifier to a supplicant. The supplicant is challenged to transform the message and return it. The transform will prove the supplicant's identity.

Channel: A small frequency range that is a subdivision of a service band.

Channel Bandwidth: The range of frequencies in a channel; determined by subtracting the lowest frequency from the highest frequency.

Channel Reuse: The ability to use each channel multiple times in different cells in the network.

Cipher: An encryption method.

Class 5 Switch: In telephony, the switch at the lowest level of the switching hierarchy. Subscribers connect to these switches.

Classic Bluetooth: Early version of Bluetooth that operated at speeds of 2 to 3 Mbps.

Clear Line of Sight: An unobstructed radio path between the sender and the receiver.

CLI: See Command Line Interface.

Client Host: In client/server processing, a server program on a server host provides services to a client program on a client host.

Client Program: Program that receives service from a server program on a server host.

Client/Server Architecture: The form of client/server computing in which the work is done by programs on two machines.

Clock Cycle: A period of time during which a transmission line's state is held constant.

Cloud: An image that indicates that the user does not need to know what goes on within the problem. A general name for services provided by companies over the Internet.

Cloud Service Provider (CSP): A company that provides cloud computing services.

Coaxial Cable: Copper transmission medium in which there is a central wire and a coaxial metal tube as the second connector.

Co-channel Interference: In wireless transmission, interference between two devices transmitting simultaneous in the same channel.

Coin Battery: Small round batter about the size of a coin. Produces little power but has a long battery life.

Collision: When two simultaneous signals use the same shared transmission medium, the signals will add together and become scrambled (unintelligible).

Command and Control Server: In a distributed denial of service attack, an intermediate server to which the botmaster sends commands. The command and control server sends commands to individual bots on compromised hosts.

Command Line Interface (CLI): Software interface in which the user types commands on a single line. Communication in both directions is limited to keyboard character.

Comprehensive Security: Security in which all avenues of attack are closed off.

Compression: Reducing the number of bits needed to be transmitted when the traffic has redundancy that can be removed.

Compromise: A successful attack.

Computer Security Incident Response Team (CSIRT): A team convened to handle major security incidents, made up of the firm's security staff, members of the IT staff, and members of functional departments, including the firm's legal department.

Computing Infrastructure: In Infrastructure as a Service, servers and their operation, database management systems, and related services.

Command mode: In Cisco's Internet Operating System, an interaction mode in which the device gives the user a prompt and the user types a command. This is a primitive but efficient interaction mode that consumes few resources. However, commands usually have complex syntax.

Confidentiality: Assurance that interceptors cannot read transmissions.

Connection: An enduring communication session with a start, individual message exchanges, and a close.

Connectionless Protocol: A protocol in which there is no enduring communication session between two devices. Messages are sent individually with no prior agreement to communicate.

Connection-Oriented Protocol: Type of conversation in which there is a formal opening of the interactions, a formal closing, and maintenance of the conversation in between.

Content Delivery Network (CDN): An Internet delivery system that stores content near the user in order to reduce latency.

Continuity Testers: UTP tester that ensures that wires are inserted into RJ-45 connectors in the correct order and are making good contact.

Control Agility: The ability to change the control function quickly and easily.

Control Function: In SDN, the function that determines how the control function acts. Traditionally, the control function was implemented on individual switches, routers, and access points. In SDN, the control function is centralized.

Core: 1) In optical fiber, the very thin tube into which a transmitter injects light. 2) In a switched network, the collection of all core switches.

Core Switch: A switch further up the hierarchy that carries traffic between pairs of switches. May also connect switches to routers.

Corporate Access Point: Access point used in an organization. Has higher quality than a home access router and is centrally manageable.

Country Top-Level Domain (cTLD): First-level domain name that specifies the owner's country (.UK, .AU, .CN, etc.)

Credentials: Proof of identity that a supplicant can present during authentication.

Credit Card Number Theft: Stealing a credit card number, and usually related information, in order to commit fraud.

Crimeware: Software used to commit crime. Often built by a third party and sold to the attacker.

Crimp: A device that presses the connector onto wires in a cord. To prevent the wires from being pulled out of the connection.

Crimping Tool: Tool used to compress an RJ-45 connector onto the untwisted wires of a UTP cord.

Cross-Site Scripting (XSS): Attack in which the application program reflects user input back in a way that permits the execution of a malicious script in the application program.

Cryptography: Mathematical methods for protecting communication.

CSIRT: See Computer Security Incident Response Team.

CSMA/CA+ACK: See Carrier Sense Multiple Access with Collision Avoidance and Acknowledgments. See definitions of the individual components.

CSP: See Cloud Service Provider.

CSU/DSU: Device on a customer premises that terminates a carrier's transmission line. Carrier service unit/data service unit.

Customer Premises: The property owned by the organization that uses the network.

Customer Premises Equipment (CPE): Equipment owned by the customer, including PBXs, internal vertical and horizontal wiring, and telephone handsets.

Cybercriminal: Criminal who commits crimes using a computer.

Cyberterror: A computer attack made by terrorists.

Cyberwar: A computer attack made by a national government.

Data Link: The path that a frame takes across a single network (LAN or WAN).

Data Link Layer Addresses: Device address at the data link layer. The source and destination data link layer addresses in in the frame's header and are used by switches or access points to forward the frame.

Data Miner: Malware that actively searches a victim computer's data files to information that can be used in a crime.

Datagram: Generic name for a message in a connectionless protocol.

Dead Zone (Dead Spot): A location where a receiver cannot receive radio transmission due to an obstruction blocking the direct path between sender and receiver.

Decibels (dB): A way of expressing the ratio between two power levels, P_1 and P_2 on a logarithmic basis.

Decision Cache: In routing, a list a router keeps of recent routing decisions for specific IP addresses so that it does not have to go through an entire routing decision again if another packet to that IP address arrives. This is nonstandard and somewhat risky.

Decrypt: Conversion of encrypted ciphertext into the original plaintext so an authorized receiver can read an encrypted message.

Dedicated Link: Unshared transmission link dedicated to the use of a single device.

Default Router: The next-hop router that a router will forward a packet to if the routing table does not have a row that governs the packet's IP address except for the default row.

Default Row: The row of a routing table that will be selected automatically if no other row matches; its value is 0.0.0.0.

Defense in Depth: The use of successive lines of defense.

Demilitarized Zone (DMZ): A subnet that holds servers that must be freely accessible by the outside world, such as public webservers and mail servers. Hosts in the DMZ will be under constant attack and must be hardened exceptionally well. Access from the DMZ to the internal network should be rare and very tightly controlled.

Denial-of-Service (DoS) Attack: The type of attack whose goal is to make a computer or a network unavailable to its users.

Density: In Wi-Fi, the number of wireless devices that use an access point.

Destination Host: Host that receives a message from another host, the source host.

Destination IP Address: The IP address of the host that receives a packet.

Destination IP Address Field: In a packet, a field that gives the IP address of the destination host.

Destination Port Number Field: In a TCP segment or UDP datagram, a field that gives the port number on the destination device.

DHCP: See Dynamic Host Configuration Protocol.

Differentiated Services Control Point: Field that specifies the quality of service that a packet should receive.

Diffserv (Differentiated Services): The field in an IP packet that can be used to label IP packets for priority and other service parameters.

Digital Certificate: A document that gives the name of a true party, that true party's public key, and other information; used in authentication.

Digital Certificate Authentication: Authentication in which each user has a public key and a private key. Authentication depends on the applicant knowing the true party's private key; requires a digital certificate to give the true party's public key.

Digital Signaling: Signaling that uses a few states. Binary (two-state) transmission is a special case of digital transmission.

Digital Subscriber Line (DSL): A technology that provides digital data signaling over the residential customer's existing single-pair UTP voice-grade copper access line.

Directory Search: In telephony, the searching for the address of a peer to which a peer wishes to connect to. In the domain name system, searching for the IP address associated with a host name.

Directory Server: Server that stores information about an organization's resources hierarchically.

Directly Propagating Worms: A type of worm that tries to jump from the infected computer to many other computers without human intervention.

Dish Antenna: An antenna that points in a particular direction, allowing it to send stronger outgoing signals in that direction for the same power and to receive weaker incoming signals from that direction.

Disassociate Message: In Wi-Fi, a frame that tells a wireless device that is associated with an access point to disassociate itself. This has legitimate uses, but it can also be used to create a denial-of service attack against the wireless devices associated with the access point.

Distributed Computing Architecture: An application architecture in which a program running on one machine calls multiple programs on other machines, which may call programs on yet other machines. After calling other programs, the calling program uses results from the called programs in its own logic flow.

Distribution System: In 802.11 W-Fi, the transmission system that connects different Wi-Fi access points. In LANs, this is almost always Ethernet.

DMZ: See Demilitarized Zone.

DNS: See Domain Name System.

Domain: In DNS, a group of resources (routers, single networks, and hosts) under the control of an organization.

Domain Name System (DNS): A server that provides IP addresses for users who know only a target host's host name. DNS servers also provide a hierarchical system for naming domains.

Domain Name Registrar: An organization that sells or allocates second-level domain names.

Domain registrars: Companies that allow individuals and organizations to purchase the right to use a particular second-level domain name on the Internet.

Dotted Decimal Notation: The notation used to ease human comprehension and memory in reading IPv4 addresses.

Drive-By Hacker: A hacker who parks outside a firm's premises and eavesdrops on its data transmissions; mounts denial-of-service attacks; inserts viruses, worms, and spam into a network; or does other mischief.

Drop Cable: A thin coaxial cable access line that runs from the cable television company line in a neighborhood to individual homes.

DSL: See Digital Subscriber Line.

DSL Access Multiplexer (DSLAM): A device at the end office of the telephone company that sends voice signals over the ordinary PSTN and sends data over a data network such as an ATM network.

DSLAM: See DSL Access Multiplexer.

Dual Mode: In Bluetooth, a device that implements both Classic Bluetooth and Bluetooth LE

Dynamic Host Configuration Protocol (DHCP): The protocol used by DHCP servers, which provide each user PC with a temporary IP address to use each time he or she connects to the Internet.

Dynamic IP Address: A temporary IP address that a client PC receives from a DHCP server.

Dynamic Routing Protocol: A protocol that allows routers to exchange routing table information.

D-Wave: An Internet of Things protocol similar to Zigbee.

Echo Reply Message: In ICMP, a message that responds to an Echo message.

Echo Request Message: ICMP message that asks a host to send back an echo reply message. This lets the sender know that the other devices is reachable and also gives the round-trip latency to that device.

ECN: See Explicit Congestion Notification.

Economically Feasible: Whether the benefits of a choice outweigh the costs. If they do, then the choice is economically feasible.

Edge Router: A router at the edge of the network between an organization and its Internet service provider.

EIGRP: See Enhanced Interior Gateway Routing Protocol.

Electromagnetic Interference (EMI): Unwanted electrical energy coming from external devices, such as electrical motors, fluorescent lights, and even nearby data transmission wires.

Electronic Signature: A bit string added to a message to provide message-by-message authentication and message integrity.

E-LAN Service: In carrier Ethernet, a service that gives Ethernet connections between multiple sites, effectively connecting them into a single Ethernet network.

E-Line Service: In carrier Ethernet, a service that gives an Ethernet connection between two sites, effectively connecting them into a single Ethernet network.

Encoding: Converting messages into bits.

Encapsulated Security Protocol (ESP): In IPsec, the standard that adds encryption, authentication, and message integrity to IPv4 or IPv6 packets.

Encryption for Confidentiality: To encrypt a message so that an eavesdropper who intercepts it cannot read it; however, the intended receiver can decrypt it and read it.

End Office Switch: The nearest switch of the telephone company to the customer premises.

End-to-End Encryption: The encryption of traffic all the way between two end devices, such as the source and destination host.

End-to-End Security: The encryption of traffic all the way between two end devices, such as the source and destination host.

Enhanced Interior Gateway Routing Protocol (EIGRP): Interior routing protocol used by Cisco routers.

Enterprise Mode: In WPA and 802.11i, operating mode that uses 802.1X.

Ephemeral Port Number: The temporary number a client selects whenever it connects to an application program on a server. According to IETF rules, ephemeral port numbers should be between 49153 and 65535.

Error! Reference Source Not Found: An HTTP response message status code that is returned to the browser if the requested webpage could not be found.

ESP: See Encapsulating Security Payload.

ESP Header: The part of the ESP content that goes before the data to be protected.

Espionage: To steal the trade secrets of a company.

ESP Trailer: The part of the ESP content that goes after the data to be p; works with the ESP header to provide security to the data.

Error Rate: In biometrics, the normal rate of misidentification when the subject is cooperating.

Ethernet: Switched network standard; dominates in LANs; also used in WANs. Standardized by the IEEE 802.3 Working Group.

Ethernet Cord: A physical cord used for Ethernet transmission. The term is normally used for 4-pair UTP wiring.

Ethernet Connector: Connector that terminate a 4-pair UTP cord so that it can be plugged into an Ethernet jack.

Ethernet Frame: A message at the data link layer in an Ethernet network.

Ethernet Jacks: Port in an Ethernet switch or host used by Ethernet. The term is normally used for RJ-45 ports for 4-pair UTP cords.

Ethernet II Frame: The Ethernet frame syntax that was in use prior to the 802.3 Working Group taking control of Ethernet standards. Simpler than the 802.3 Ethernet frame. However, the Internet Protocol standard calls for the use of Ethernet II frames rather than 802.3 Ethernet frames, and this is normal practice.

EtherType Field: In an Ethernet II frame, the field that specifies the contents of the data field—usually an IPv4 or IPv6 packet.

EUI-48: See Extended Unique Identifier-48.

Evil Twin Access Point: Attacker access point outside a building that attracts clients inside the building to associate with it.

Evil twin attack: Wi-Fi attack in which the attacker intercepts encrypted frames from a host, decrypts and reads them in the clear, and then reencrypts them and passes them on.

Exit Node: In a TOR network, the node that transmits the final packet to the destination host. It is the final node in the TOR network transmission—the node from which the packet exits the TOR network.

Explicit Congestion Notification (ECN): Field that notifies the receiver that there is congestion on the network. The receiver may respond by reducing its transmission rate.

Exploit: Term used variously for the act of breaking into a computer, the method used to break in, or the crimeware software used to break in.

Extended Unique Identifier-48 (EUI-48): A common data link address format with a length of 48 bits. Formerly called a MAC address.

Extension Header: In IPv6, a header that follows the main header.

Exterior Dynamic Routing Protocol: Routing protocol used between autonomous systems.

Facial Recognition: Biometric authentication method that uses the shape of a person's face as proof of identity.

Fail-Over: When one system will take over the work load immediately if another system fails.

False Alarm: An apparent incident that proves not to be an attack.

False Positive: A false alarm.

Fiber Cord: Optical fiber cord. Used for longer Ethernet physical links.

Fiber to the Home (FTTH): Optical fiber brought by carriers to individual homes and businesses.

Field: (1) A subdivision of a message header or trailer. (2) An IPv6 field is a group of four lowercase hexadecimal symbols. Each field represents 16 bits. Fields are separated by colons.

File Format Standard: Standard for the format of files delivered by a network application. Most network applications have two standards—one to control delivery, the other for the file format of the delivered file or message.

File Storage Services: Cloud services such as Dropbox and iCloud, which store user files in the cloud for backup and to provide access anywhere.

Fin Bit: One-bit field in a TCP header; indicates that the sender wishes to open a TCP connection.

Fingerprint Recognition: The use of fingerprints to identify a person.

Firewall: A security system that examines each packet passing through it. If the firewall identifies the packet as an attack packet, the firewall discards it and copies information about the discarded packet into a log file.

Firewall Filtering Mechanisms: The methods used by a firewall to identify provable (definite) attack packet; if identified, the packet is dropped and logged.

Firewall Log File: A file that contains summary information about packets dropped by a firewall.

Firewall Policy Server: Server that stores firewall policies. It sends access control list changes to individual firewalls to implement these policies.

Flag Field: A one-bit field.

Flow Label Field: In IPv6, all packets in a stream of packets are given the same flow label number.

Forwarding Function: In SDN, the switch, router, or access point function that sends incoming frames or packets back out.

Frame: 1) A message at the data link layer. 2) In time division multiplexing, a brief time period, which is further subdivided into slots.

Frame Check Control Sequence Field: A four-octet field used in error checking in Ethernet. If an error is found, the frame is discarded.

Frequency: The number of complete cycles a radio wave goes through per second. In sound, frequency corresponds to pitch.

Frequency Spectrum: The range of all possible frequencies from zero hertz to infinity.

FTTH: See Fiber to the Home.

Full-Duplex Transmission: A type of communication that supports simultaneous two-way transmission. Almost all communication systems today are full-duplex systems.

Gateway: An obsolete term for "router"; still in use by Microsoft.

Gbps: Gigabit per second.

Generic Top-Level Domain (gTLD): First-level domain name that specifies the type of organization that owns the domain (.com, .edu, etc.).

Get: An SNMP command sent by the manager that tells the agent to retrieve certain information and return this information to the manager.

gTLD: See Generic Top-Level Domain.

Guideline: A directive that should be followed but that need not be followed, depending on the context.

Hacking: The intentional use of a computer resource without authorization or in excess of authorization.

Hacktivists: Hackers who are motivated to steal information by politics rather than monetary gain.

Handoff: a) In wireless LANs, a change in access points when a user moves to another location. b) In cellular telephony, transfer from one cellsite to another, which occurs when a subscriber moves from one cell to another within a system.

Head End: The cable television operator's central distribution point.

Header Checksum Field: The UDP datagram field that allows the receiver to check for errors.

Hertz (Hz): One cycle per second, a measure of frequency.

Hex Notation: See Hexadecimal Notation.

Hexadecimal (Hex) Notation: The Base 16 notation that humans use to represent 48-bit MAC source and destination addresses.

Hierarchical Topology: A network topology in which all switches are arranged in a hierarchy, in which each switch has only one parent switch above it (the root switch, however, has no parent); used in Ethernet.

Hierarchy: 1) The type of topology wherein there are multiple layers of switches organized in a hierarchy, in which each node has only one parent node; used in Ethernet. 2) In IP addresses, three multiple parts that represent successively more specific locations for a host.

Hop Limit Field: In IPv6, the field that limits the number of hops an IPv6 packet may make among routers.

Host: Any computer attached to a network.

Host Name: An unofficial designation for a host computer.

Host Part: The part of an IP address that identifies a particular host on a subnet.

Host-to-Host VPN: Virtual private network that creates cryptographically protected connections between two individual hosts.

Host-to-Site VPN: Virtual private network that creates cryptographically protected connections between an individual host and a corporate site.

HTTP: See Hypertext Transfer Protocol.

HTTP Request Message: In HTTP, a message in which a client requests a file or another service from a server.

HTTP Response Message: In HTTP, a message in which a server responds to a client request; contains either a requested file or an error message explaining why the requested file could not be supplied.

Human Interface Device (HID) Profile: In Bluetooth, this profile is used for mice, keyboards, and other input devices.

Hybrid TCP/IP–OSI Architecture: The architecture that uses OSI standards at the physical and data link layers and TCP/IP standards at the internet, transport, and application layers; dominant in corporations today.

Hypertext Transfer Protocol (HTTP): The protocol that governs interactions between the browser and the webserver application program.

IaaS: See Infrastructure as a Service.

IANA: See Internet Assigned Numbers Authority.

ICMP: See Internet Control Message Protocol.

ICMP Control Message: Internet Control Message Protocol message that directs a host to take an action.

ICMP Error Advisement: A message sent in error advisement to inform a source device that an error has occurred.

ICV: See Integrity Check Value.

Identity Theft: Stealing enough information about a person to impersonate him or her in large financial transactions.

IDC: See Insulation Displacement Connection.

IDS: See Intrusion Detection System.

IEEE 802.11 Working Group: IEEE Working Group that creates Wi-Fi (802.11) wireless LAN standards.

IKE: See Internet Key Exchange.

IMAP: See Internet Message Access Protocol.

Implementation Guidance: Instructions that are more specific than policies but less specific than implementation.

Incident: A successful attack.

Internet process The process (hardware or software) that implements the transport layer's functionality.

Individual Throughput: The actual speed a single user receives (usually much lower than aggregate throughput in a system with shared transmission speed).

Infrastructure as a Service: Providing computing infrastructure, which consists of servers and their operation, database management systems, and related services, as a service in which the customer pays by use instead of owning the infrastructure.

Initial Authentication: Authentication at the beginning of a communication session, before the two sides exchange working data. As opposed to message-by-message authentication during data exchange.

Insiders: People within an organization; they are especially dangerous if they attack you. Includes everyone with insider permissions, such as contractors.

Instantiate: See Spawn.

Insulation Displacement Connection (IDC): A connection in which a metal prong is pushed through insulation into another wire.

Integers: Whole numbers.

Integrated Log File: A log file that integrates the data from multiple log files on different devices. Permits a more complete picture of an attack or suspected attack.

Integrity Check Value (ICV): The optional message integrity part of the trailer for the Encapsulating Security Protocol.

Interface: 1) The router's equivalent of a network interface card; a port on a router that must be designed for the network to which it connects. 2) In webservices, the outlet through which an object communicates with the outside world.

Interface ID: The third part of an IPv6 address. Indicates the host on the subnet of the organization on the Internet containing the host.

Interior Dynamic Routing Protocol: Routing protocol used within a firm's internet.

Internal Router: A router that connects different LANs within a site.

International Organization for Standardization (ISO): A strong standards agency for manufacturing, including computer manufacturing.

International Telecommunications Union–Telecommunications Standards Sector (ITU–T): A standards agency that is part of the United Nations and that oversees international telecommunications.

Internet Assigned Numbers Authority (IANA): The organization that allocates blocks of IP addresses to regional assigned number authorities for distribution to organizations and Internet service providers.

Internet Control Message Protocol (ICMP): The protocol created by the IETF to oversee supervisory messages at the internet layer.

Internet Core Routers Router used by an Internet service provider.

Internet Engineering Task Force (IETF): TCP/IP's standards agency.

Internet Key Exchange (IKE): In IPsec, the standard for the initial negotiation stage in establishing a security association.

Internet Layer: The layer that governs the transmission of a packet across an entire internet.

Internet Message Access Protocol (IMAP): One of the two protocols used to download received e-mail from an e-mail server; offers more features but is less popular than POP.

Internet of Things (IoT): Internet use by small devices that talk to one another, with no human involvement.

Internet Layer Process: Hardware or software process that implements internet layer functionality on a host or router.

Internet Service Provider (ISP): Carrier that provides Internet access and transmission.

Interoperate: To be able to work together.

Intrusion Detection System (IDS): A system that warns of a possible attack.

Inverse square law: Radio signal strength declines with the square of transmission distance.

IOS: Operating system used on Cisco switches, routers, access points, firewalls, and other devices. Designed to use a command line interface.

IoT: See Internet of Things.

IP Address: An Internet Protocol address; the address that every computer needs when it connects to the Internet; IP addresses are 32 bits long.

Internet Protocol Security (IPsec): A set of standards that operate at the internet layer and provide security to all upper layer protocols transparently.

IP Version 4 (IPv4): The standard that governs most routers on the Internet and private internets.

IP Version 6 (IPv6): A new version of the Internet Protocol.

IPsec Gateway: Border device at a site that converts internal data traffic into protected data traffic that travels over an untrusted system such as the Internet.

IPv4: See IP Version 4.

IPv4 Addresses: Addresses in the fourth version of the Internet protocol. 32 bits long. In contrast, IPv6 addresses are 128 bits long.

IPv4 Mask: A 32-bit series with a number of 1s followed by the number of 0s. The number of 1s corresponds either to the IKPv4 packet's network part or network plus subnet part. Used by routers to assign routes to all packets going to a particular network or subnet on a network.

IPv6: See IP Version 6.

IPv6 Canonical Text Representation: A standardized way of representing an IP address for condensed human reading.

IPv6 main header: The first header in an IPv6 packet. Other headers may follow. These are extension headers. The data field follows the last extension header.

Iris Recognition: Authentication that uses the pattern in the iris (the colored part of the supplicant' eye)

ISO: See International Organization for Standardization.

Jitter: Variability in latency.

kbps: Kilobits per second.

Keystroke Logger: Type of spyware that captures victim keystrokes and sends them to the attacker.

Kill Chain: The series of steps that must all succeed for an attack to succeed. If defenses can stop a single link in the chain, the attack will fail. A method for visualizing attacks and how to stop them.

Label Header: In MPLS, the header added to packets before the IP header; contains information that aids and speeds routers in choosing which interface to send the packet back out.

Label Number: In MPLS, number in the label header that aids label-switching routers in packet sending.

Label Switched Path: A path that all packets to a particular address will take across and MPLS label-switched network.

Label Switching Router: Router that implements MPLS label switching.

LAN: See Local Area Network.

Latency: Delay, usually measured in milliseconds.

Layers: Standards agencies divide the job of getting two applications on two different hosts into four to seven layers of functionality, each providing service to the layer above it. One layer can be changed without requiring a change in upper layers.

Leased Line: A high-speed, point-to-point, always-on connection.

Least Permissions: The minimum permissions an employee needs to do his or her job. If broader permissions are given, that creates a security vulnerability.

Licensed Service Band: Regulated radio signal band that requires radio devices to be licensed to prevent interference between radios.

Line Feed: Moves the cursor or print head one line down.

Link Aggregation: The use of two or more trunk links between a pair of switches; also known as trunking or bonding.

Link encryption: Providing encryption over a single physical link or data link, instead of over the entire route between the source and destination hosts.

Link Security: Security over *part* of the path between two devices. As opposed to end-to-end security between the devices. For Instance, In 802.11i, security over the link between an access point and a wireless device.

Load Balancing: Dividing traffic across routers in order not to overload any single route.

Local Area Network (LAN): A network within a customer's premises.

Local Loop: In telephony, the line used by the customer to reach the PSTN's central transport core.

Longest Match: The matching row that matches a packet's destination IP address to the greatest number of bits; chosen by a router when there are multiple matches.

MAC: See Media Access Control.

MAC address: Former name for EUI-48 address.

Main IPv6 Header: The primary header in IPv6. Followed by zero or more extension header, then the higher-level content of the packet.

Major Incident: A large security incident with wide repercussions. Must be managed by the computer security incident response team.

Malware: Software that seeks to cause damage.

MAN: See Metropolitan Area Network.

Manageable Switch: A switch that can be managed remotely via the Simple Network Management Protocol.

Managed Device: A device that can be managed remotely via the Simple Network Management Protocol. Examples: printers, switches, routers, and user PCs.

Management Information Base (MIB): A specification that defines what objects can exist on each type of managed device and also the specific characteristics of each object; the actual database stored on a manager in SNMP. There are separate MIBs for different types of managed devices; both a schema and a database.

Manager: The central PC or more powerful computer that uses SNMP to collect information from many managed devices.

Man-in-the-Middle Attack: An attack in which an eavesdropper intercepts message transmissions between two devices in order to read exchanged messages.

Master–Slave Control: Form of transmission in which one host controls the transmission of another host.

Mbps: Megabits per second.

Media Access Control (MAC): The process of controlling when stations transmit; also, the lowest part of the data link layer, defining functionality specific to a particular LAN technology.

Media Gateway: A device that connects IP telephone networks to the ordinary public switched telephone network. Media gateways also convert between the signaling formats of the IP telephone system and the PSTN.

Message Integrity: The assurance that a message has not been changed en route; or if a message has been changed, the receiver can tell that it has been changed.

Message Order: Controlling when one device in a pair may transmit.

Metric: A number describing the desirability of a route represented by a certain row in a routing table.

Metropolitan Area Network (MAN): A WAN that spans a single urban area.

MIB: See management information base.

Millisecond (ms): The unit of time in which latency is measured.

Milliwatt (mW): One thousandth of a watt.

MIMO: See Multiple Input/Multiple Output.

Minor Incident: Security incident that can be managed by the on-duty staff.

Mobile Phone: See Cellphone.

Mobile Telephone Switching Office (MTSO): A control center that connects cellular customers to one another and to wired telephone users, as well as overseeing all cellular calls (determining what to do when people move from one cell to another, including which cellsite should handle a caller when the caller wishes to place a call).

Modal Dispersion: The main propagation problem for optical fiber; dispersion in which the difference in the arrival times of various modes (permitted light rays) is too large, causing the light rays of adjacent pulses to overlap in their arrival times and rendering the signal unreadable.

Mode: An angle at which light rays are permitted to enter an optical fiber core.

Momentary Traffic Peak: A surplus of traffic that briefly exceeds the network's capacity, happening only occasionally.

MPLS: See Multiprotocol Label Switching.

Ms: See Millisecond.

MTSO: See Mobile Telephone Switching Office.

Multimode Fiber: The most common type of fiber in LANs, wherein light rays in a pulse can enter a fairly thick core at multiple angles. Inexpensive but can transmit signals over sufficient distance for LAN usage.

Multipath Interference: Interference caused when a receiver receives two or more signals—a direct signal and one or more reflected signals. The multiple signals may interfere with one another.

Multiuser MIMO (MU-MIMO): Using MIMO to send Wi-Fi frames to multiple hosts simultaneously and also to receive frames simultaneously.

Multiple Input/Multiple Output (MIMO): A radio transmission method that sends several signals simultaneously in a single radio channel.

Multiplex: 1) Having the packets of many conversations share trunk lines; reduces trunk line cost. 2) The ability of a protocol to carry messages from multiple next-higher-layer protocols in a single communication session.

Multiprotocol Label Switching (MPLS): A traffic management tool used by many ISPs.

mW: See Milliwatt.

Nanometer (nm): The measure used for wavelengths; one billionth of a meter (10^{-9} meter).

NAT: See Network Address Translation.

National Institute of Standards and Technology. United States agency that creates security recommendations for Federal agencies. Given the Institute's recommendation, the Institute's recommendations are widely adopted in industry.

Near Field: In radio transmission, the signal very near the antenna. Has unique properties. Can be used to query radio frequency ID circuits that have no power.

Near Field Communication (NFC): Form of radio transmission in which devices within about 4 cm (roughly 2 in.) can communicate peer-to-peer.

Network Address Translation (NAT): Converting an IP address into another IP address, usually at a border firewall; disguises a host's true IP address from sniffers. Allows more internal addresses to be used than an ISP supplies a firm with external addresses.

Network Applications: Those applications that require a network to communicate with one another in order to function.

Network Core: The central part of the network.

Network Management Program (Manager): A program run by the network administrator on a central computer.

Network Operation Center (NOC): Central management point for a network.

Network Part: The part of an IP address that identifies the host's network on the Internet.

Network Segmentation: When the network is divided into different security domains, each with security controls that are appropriate to it. Strict rules for communication between security zones.

Network Stack: Programs on a host that govern communication to and from the Internet.

Network Standard: A rule of operation that governs the exchange of messages between two hardware or software processes.

Network Visibility: A type of tool that helps managers comprehend what is going on in their networks.

Network Working Group: The original ad hoc standards setting group for the ARPANET. When the ARPANET grew into the Internet, the group matured into the Internet Engineering Task force, which now sets standards on the Internet.

Next-Generation Firewall (NGFW): Firewall that can detect applications, not simply port numbers. Permits much finer control over network traffic.

Next Header Field: In an IPv6 main or extension header, the field that specifies the next header's type or specifies that the payload follows the header.

Next-Hop Router: A router to which another router forwards a packet in order to get the packet a step closer to reaching its destination host.

NGFW: See Next-Generation Firewall.

Nm (nm): See Nanometer.

NOC: See Network Operation Center.

Nonmalicious Insiders: Insiders (employees, etc.) who do not mean to do serious harm yet may do so through ignorance or while underestimating the riskiness of their actions.

Northbound APIs: In SDN, an application program interface between an SDN application and the SDN controller.

Object: In SNMP, an aspect of a managed device about which data is kept.

Octet: A collection of 8 bits; same as a byte.

OFDM: See Orthogonal Frequency Division Multiplexing.

OM: See Optical Multimode.

Omnidirectional Antenna: An antenna that transmits signals in all directions and receives incoming signals equally well from all directions.

One-to-One Connection: Transmission from one host to another. Unicasting.

One-Pair Voice-Grade (1PVG) UTP: The traditional telephone access lines to individual residences.

Open Connect Appliances: Video delivery servers in Netflix's content delivery network *Open Connect*.

Open Connect Network: Netflix's content delivery network.

Open Shortest Path First (OSPF): Complex but highly scalable interior routing protocol.

Optical Fiber: Cabling that sends signals as light pulses.

Optical Multimode (OM): Quality standard for multimode fiber.

Organizational system security: A name for all of the non-technological aspects needed for the protection of a business systems such as a department or project team.

Orthogonal Frequency Division Multiplexing (OFDM): A form of spread spectrum transmission that divides each broadband channel into subcarriers and then transmits parts of each frame in each subcarrier.

OSI: The Reference Model of Open Systems Interconnection; the 7-layer network standards architecture created by ISO and ITU-T; dominant at the physical and data link layers, which govern transmission within single networks (LANs or WANs).

OSPF: See Open Shortest Path First.

Oversight: A collection of methods to ensure that policies have been implemented properly.

Packet: A message at the internet layer.

Pairwise Session Key: A session key for encrypted transmission between two devices. This key is not known by other devices.

PAN: See Personal Area Network.

Parallel Transmission: A form of transmission that uses multiple wire pairs or other transmission media simultaneously to send a signal; increases transmission speed.

Passphrase: A series of words used to generate a key.

Password Dictionaries: Dictionary of common names and passwords and common variants of these. If a password is in the dictionary, it will be cracked immediately no matter how long it is.

Patch: An addition to a program that will close a security vulnerability in that program.

Payload: 1) In security, a piece of code that can be executed by a virus or worm after it has spread to multiple machines. 2) In IPv6, all of the packet after the main packet header.

Payload Length Field: In IPv6 packets, a field that gives the length of everything following the main header, including subsidiary headers.

Peers: In peer-to-peer applications, devices that traditionally were called clients.

Peer-to-Peer (P2P) Applications: Applications that operate between devices traditionally considered to be clients, with little or no server involvement.

Peer-to-Peer (P2P) Computing: Most or all of the work is done by cooperating user computers, such as desktop PCs. If servers are present at all, they serve only facilitating roles and do not control the processing.

Peer-to-Peer Traffic: Traffic between peers in peer to peer applications.

Permission: A rule that determines what an account owner can do to a particular resource (file or directory).

Personal Area Network (PAN): A small wireless network used by a single person.

Personal Identification Number (PIN): A four- or six-digit number a cardholder types to authenticate himself or herself.

Personal Mode: Pre-shared Key Mode in WPA or 802.11i.

Physical standard: The process (hardware or software) that implements the transport layer's functionality.

Phishing: Social engineering attack that uses an official-looking e-mail message or website.

Physical Standard: Standard at the physical layer, the lowest layer in networking.

Piconet: In Bluetooth, a personal area network with up to eight devices.

PIN: See Personal Identification Number.

Ping: Sending a message to another host and listening for a response to see if it is active.

Planning: The first step in the plan-protect-respond cycle for cyberdefense. Creating plans for protections and responses.

Plan–Protect–Respond Cycle: The basic management cycle in which the three named stages are executed repeatedly.

Point-to-Point Network: A network that directly connects two devices. Often used to connect two routers on the Internet that are many miles apart.

Point-to-Point Protocol (PPP): The most widely used data link layer protocol in point-to-point networking.

Policy-Based Configuration: In SDN, creating policies that are automatically translated into configuration changes on individual devices.

Policy Database: In SDN, creating policies that are automatically translated into configuration changes on individual devices.

POP: See Post Office Protocol.

Port Number: The field in TCP and UDP that tells the transport process what application process sent the data in the data field or should receive the data in the data field.

Port Spoofing: Using a well-known port number for a different purpose, with malicious intent

Post Office Protocol (POP): The most popular protocol used to download e-mail from an e-mail server to an e-mail client.

PPP: See Point-to-Point Protocol.

Prefix Notation: A way of representing masks. Gives the number of initial 1s in the mask.

Pre-Shared Key (PSK): A mode of operation in WPA and 802.11i in which all stations and an access point share the same initial key.

Pre-Shared Key (PSK) Initial Authentication: An initial authentication mode used in 802.11i. All devices use the same pre-shared key for initial authentication. Used in residences and organizations that only have a single access point. Called personal mode by the Wi-Fi Alliance.

Priority Level: The 3-bit field used to give a frame one of eight priority levels from 000 (zero) to 111 (eight).

Private IP Address Range: An IP address that may be used only within a firm. Private IP addresses have three designated ranges: 10.x.x.x, 192.168.x.x, and 172.16.x.x through 172.31.x.x.

Private Key: A key that only the true party should know. Part of a public key–private key pair.

Profile Wave: The Wi-Fi alliance creates profiles, which are subsets of a particular standard. The alliance bases interoperability testing on specific profiles.

Prompt: In a command line interface, characters at the start of a line to indicate that the system is awaiting your input. May give Information on what type of input you may type.

Propagate: In signals, to travel.

Protection: Implementing a security plan; the most time-consuming stage in the plan–protect–respond management cycle.

Protocol: 1) A standard that governs interactions between hardware and software processes at the same layer but on different hosts. 2) In IP, the header field that describes the content of the data field.

Protocol Field: In IP, a field that designates the protocol of the message in the IP packet's data field.

Provable Attack Packet: A packet that is provably an attack packet.

PSK: See Pre-Shared Key.

PSTN: See Public Switched Telephone Network.

PSTN Core: The public switched telephone networks central transmission lines and switches. Does not include end office switches that serve users or transmission lines to users.

Public-facing servers: Servers that provide services to clients on the Internet. Clients must be able to access it. This can lead to attacks, so public-facing servers must be especially well protected.

Public Key: A key that is not kept secret. Part of a public key–private key pair.

Public Switched Telephone Network (PSTN): The worldwide telephone network.

QoS Guarantee: A guarantee that certain traffic will get through regardless of network congestion. Requires reserving capacity on each device.

Quality of Service (QoS) Metrics: Numerical service targets that must be met by networking staff.

Rack Server: Server that fits in a standard equipment rack. Each rack can hold several rack servers positioned one on top of another.

Radio Frequency ID (RFID): A tag that can be read at a distance by a radio transmitter/receiver.

Rapid Spanning Tree Protocol (RSTP): A version of the Spanning Tree Protocol that has faster convergence.

Rate Limited: Traffic that is limited to a certain small percentage of a network's total traffic in order to reduce congestion.

Rated Speed: The official standard speed of a technology.

RBAC: See Role-Based Access Control.

Real Time Fail-Over: Two data centers that are synchronized so that if one data center fails, the other can take over in real time (immediately).

Real Time Protocol (RTP): The protocol that adds headers that contain sequence numbers to ensure that the UDP datagrams are placed in proper sequence and that they contain time stamps so that jitter can be eliminated.

Recognized Organization: An organization recognized by the Internet Assigned Number Authority; it receives a network part.

Recommendation System: A system that recommends a product that a user might like based upon the user's past pattern of selections.

Record: In a file or database, information about a single entity.

Redundancy: Duplication of a hardware device in order to enhance reliability.

Reading log files: Many devices create log files that list each operational or security-relevant event. The organization must read these log files constantly to detect attacks. Early identification of an attack may mitigate its damage.

Reference Model of Open Systems Interconnection: Standards architecture created by the ITU-T and ISO. Acronym is OSI. Rarely spelled out.

Reflection: In cross-site scripting, when an application executes a script sent in a user's input. This can be a malicious script.

Regenerate: In a switch or router, to clean up a signal before sending it back out.

Remote Access VPN: Virtual private network that allows a remote host to communicate securely with a site.

Request for Comment (RFC): A document produced by the IETF that may become designated as an Official Internet Protocol Standard.

Request Message: In request–response cycles, a message a client programs sends to request service from a server application program.

Reserved Capacity: On routers, switches, and transmission lines, reserving a certain amount of capacity for a particular application so that messages in the application will always get through even if congestion is severe.

Reset: In TCP, a flag in a TCP segment to inform the other side that the sender will accept no further input.

Residential Access Routers: In a home network, a multifunction devices that is a trivial router but includes an Ethernet switch, a consumer-grade wireless access point, a DHCP server, and often other functionality.

Respond: In security, the act of stopping and repairing an attack.

Response Header field: In HTTP, a header field that follow the status line in an HTTP response message.

Response Message: In Challenge–Response Authentication Protocols, the message that the applicant returns to the verifier.

Reusable Password: Password that is used repeatedly to get access.

RFC: See Request for Comment.

RFID: See Radio Frequency ID.

Right of Way: Permission to lay wires in public areas; given by government regulators to transmission carriers.

Risk Analysis: The process of balancing threats and protection costs.

RJ-45 Connector: The connector at the end of a UTP cord, which plugs into an RJ-45 jack.

RJ-45 Jack: The type of jack into which UTP cords RJ-45 connectors may plug.

Roaming: 1) In cellular telephony, the situation when a subscriber leaves a metropolitan cellular system and goes to another city or country. 2) In 802.11, when a wireless host travels from one access point to another.

Rogue Access Point: An unauthorized access point. If it has no security or poor security, it allows a malicious outsider access even if all regular access points are highly secure.

Role-Based Access Control (RBAC): Assigning access to resources based on roles in the organization rather than assigning them to individual people. Individuals are then assigned to roles.

Root Cause Analysis: The analysis of data in log files to determine the fundamental cause of and observed pattern in the data.

Root DNS Server: One of 13 top-level servers in the Domain Name System (DNS).

Root Privileges: In UNIX systems, complete privileges (authorizations) on the machine, allowing the user to do anything. Also used to refer to similar privileges on non-UNIX machines, such as Windows, Apple, and mobile phone systems.

Route: The path that a packet takes across an internet.

Round-Trip Latency: The time delay between when a message is sent and the response is received.

Router: A device that forwards packets within an internet. Routers connect two or more single networks (subnets).

Routing: 1) The forwarding of IP packets. 2) The exchange of routing protocol information through routing protocols.

Routing Decision: When a router receives a packet, it must make a decision about what port to send the packet back out to get to either the next-hop router or the destination host.

Routing Prefix: The first part of an IPv6 address. Indicates the organization on the Internet containing the host.

RST Bit: In a TCP segment, if the RST (reset) bit is set, this tells the other side to end the connection immediately.

RSTP: See Rapid Spanning Tree Protocol.

RST: A TCP flag field. If set, the TCP segment tells the other party that the sender is breaking the connection.

RTP: See Real Time Protocol.

RTS/CTS: See Request to Send/Clear to Send.

SaaS: See Software as a Service.

SDN: See Software-Defined Networking.

SDN Application Programs: In SDN, a program that implements a control function, such as imposing quality of service rules on one or more devices.

SDN Controller: In SDN, the device that manages the control function for multiple switches, routers, and other forwarding devices.

Searchable Fields: In e-mail and applications, the ability to search for messages or files on the basis of the contents of specific fields such as sender, receiver, date, time, and subject.

Second-Level Domain: The third level of a DNS hierarchy, which usually specifies an organization (e.g., microsoft.com, hawaii.edu).

Security Association: An agreement between two parties on the security methods and parameters they will use in their subsequent interactions.

Security Policies: A security policy is statement of what should be done to achieve a desired level of security. Implementation is actually doing it according to the policy. Takes advantage of the different knowledge of policy makers and implementers.

Self-organizing: A network is self-organizing if it reorganizes itself automatically when devices are added or dropped.

Separation of Duties: Creating procedures or processes that require two (or more) people to complete an action. This prevents a single person from acting alone to take an unsecure or malicious action.

Sequence Number Field: In TCP, a header field that tells a TCP segment's order among the multiple TCP segments sent by one side.

Server Host: In client/server processing, a server program on a server host provides services to a client program on a client host.

Server Program: Program on a server host that provides service to a client program on a client host.

Service Band: A subdivision of the frequency spectrum, dedicated to a specific service such as FM radio or cellular telephone service.

Service Level Agreement (SLA): A quality-of-service guarantee for throughput, availability, latency, error rate, and other matters.

Service set ID (SSID): The name of a Wi-Fi access point or group of access points. A Wi-Fi user must know the SSID to connect to an access point.

Session Initiation Protocol (SIP): Relatively simple signaling protocol for voice over IP.

Session Key: Symmetric key that is used only during a single communication session between two parties.

Set: 1) When a flag's field is given the value 1. 2) An SNMP command sent by the manager that tells the agent to change a parameter on the managed device.

Shadow Zone: See Dead Zone.

Signal Analysis Software: Software that analyzes the characteristics of a radio signal, such as signal strength.

Signal Bandwidth: The range of frequencies in a signal, determined by subtracting the lowest frequency from the highest frequency.

Simple Mail Transfer Protocol (SMTP): The protocol used to send a message to a user's outgoing mail host and from one mail host to another; requires a complex series of interactions between the sender and the receiver before and after mail delivery.

Simple Network Management Protocol (SNMP): The protocol that allows a general way to collect rich data from various managed devices in a network.

Single-Mode Fiber: Optical fiber whose core is so thin (usually 8.3 microns in diameter) that only a single mode can propagate, also the one traveling straight along the axis.

Single Network: A network that uses a single set of standards for all devices. E.g., Ethernet.

Single Point of Takeover: If an attacker can take over a single system, the attacker gains control over a significant portion of your network.

SIP: See Session Initiation Protocol.

Site-to-Site VPN: Virtual private network that secures all communication between two sites.

Site Survey: In wireless LANs, a radio survey to help determine where to place access points.

Skype: A P2P VoIP service that currently offers free calling among Skype customers over the Internet and reduced-costs calling to and from Public Switched Telephone Network customers.

SLA: See Service Level Agreement.

S/MIME Protocol: A security protocol for end-to-end communication between the programs of two e-mail users.

SMTP: See Simple Mail Transfer Protocol.

Sniffer Program: In security, a program that intercepts traffic to read it in order to find information useful to an attacker.

SNMP: See Simple Network Management Protocol.

SNMP Agent: In the Simple Network Management Protocol, the hardware or software functionality on a managed device that communicates with the SNMP manager.

SNMP Get: In the Simple Network Management Protocol, a command sent by the manager that asks an agent for information about its managed device.

SNMP Manager: In the Simple Network Management Protocol, the program that collects data from managed devices and can send commands to managed devices to change their configuration.

SNMP Schema: In the Simple Network Management Protocol, the schematic structure of the management information base.

SNMP Set: In the Simple Network Management Protocol, a command from the manager to the agent to change the configuration of a managed device.

SNMP Trap: In the Simple Network Management Protocol, an alarm sent by an agent to the manager if the agent detects a problem.

Social Engineering: Tricking people into doing something to get around security protections.

Socket: The combination of an IP address and a port number, designating a specific connection to a specific application on a specific host. It is written as an IP address, a colon, and a port number, for instance, 128.171.17.13:80.

Software as a Service (SaaS): Service in which an application service provider supplies an application to customers on demand.

Software-Defined Networking (SDN): A radical change in networking that removes the control function from individual switches, routers, access points, and other devices.

Solid-Wire UTP: Type of UTP in which each of the eight wires really is a single solid wire rather than a collection of strands.

Source Host: Host that transmits a message to another host, the destination host.

Source IP Address: The IP address of the host that transmits.

Source IP Address Field: Field in an IP packet containing the IP address of the host that transmits the packet.

Source Port Number Field: Field in a TCP segment or a UDP datagram containing the IP address of the host that transmits.

Southbound APIs: In SDN, an application program interface between an SDN controller and a switch, router, or other device.

Spawn: To launch a copy of a virtual machine or a new virtual machine. Also called instantiation (creating an instance of).

Spear Phishing: A phishing attack that is highly focused on an individual. Likely to be extremely convincing because it contains content highly familiar to the intended victim.

Splitter: A device that a DSL user plugs into each telephone jack; the splitter separates the voice signal from the data signal so that they cannot interfere with each other.

Spread Spectrum Transmission: A type of radio transmission that takes the original signal and spreads the signal energy over a much broader channel than would be used in normal radio transmission; used in order to reduce propagation problems, not for security.

Spyware: Software that sits on a victim's machine and gathers information about the victim.

SSL/TLS: See Secure Sockets Layer and Transport Layer Security.

Stand-Alone Processing: An application architecture in which all processing is done on a single machine.

Standards Agency: An organization that creates and maintains standards.

Standards Architecture: A family of related standards that collectively allows an application program on one machine on an internet to communicate with another application program on another machine on the internet.

State: In digital physical layer signaling, one of the few line conditions that represent information.

Stateful Packet Inspection: Firewall filtering mechanism that uses different filtering methods in different states of a conversation.

Static IP Address: An IP address that never changes.

Strain Relief: In a UTP connectorization, pressing the RJ-45 connector into the jacket of a UTP cord. This means that even if the cord is pulled, causing strain, the cord will not pull out of the connector.

Strand: In optical fiber, a core surrounded by a cladding. For two-way transmission, two optical fiber strands are needed.

Stranded-Wire UTP: Type of UTP in which each of the eight "wires" really is a collection of wire strands.

Stripping Tool: Tool for stripping the sheath off the end of a UTP cord.

Subnet: A small network that is a subdivision of a large organization's network.

Subnet ID: The second part of an IPv6 address. Indicates the host's subnet in the organization on the Internet containing the host.

Subnet Part: The part of an IP address that specifies a particular subnet within a network.

Supervisory Protocols: A protocol that governs how network devices operate, as opposed to a protocol that is used to send, receive, and forward information.

Supplicant: The party trying to prove his or her identity.

Surreptitiously: Done without someone's knowledge, such as surreptitious face recognition scanning.

Switch: A device that forwards frames within a single network.

Switching Decision: In switched networks, the decision a switch makes when it receives a frame in one port and must decide which other port to send the frame back out to the next device along the data link.

SYN Bit: In TCP, the flags field that is set to indicate if the message is a synchronization message.

Synchronization Profile (SYNCH): Bluetooth profile for synchronizing data on two devices.

Synchronized Data Center: Two (or more) data centers with synchronized software and data. This permits real-time fail-over.

Synchronous DSL: Digital subscriber line with the same speed in both directions. Normally used In businesses.

Syntax: In message exchange, how messages are organized.

SYSLOG Standard: Standard for transmitting data from log files on individual devices to an integrated log file.

Tag: An indicator on an HTML file to show where the browser should render graphics files, when it should play audio files, and so forth.

Tag Field: One of the two fields added to an Ethernet MAC layer frame by the 802.1Q standard.

Tbps: Terabits per second—a thousand billions of bits per second.

TCP: See Transmission Control Protocol.

TCP Reset Segment: TCP segment in which the RST flag bit is set.

TCP Segment: A TCP message.

TCP/IP: The Internet Engineering Tasks Force's standards architecture; dominant above the data link layer.

TDR: See Time Domain Reflectometry.

Test Signals: Signal sent by a high-quality UTP tester through a UTP cord to check signal quality parameters.

Text Standards: Standards for representing keyboard characters plus some control codes. Therefore, not actually limited to text.

Threat Environment: The threats that face the company.

Throughput: The transmission speed that users *actually* get. Usually lower than a transmission system's rated speed.

Time Domain Reflectometry (TDR): A testing system for UTP that can detect breaks in the wire.

Time to Live (TTL) Field: The field added to a packet and given a value by a source host, usually between 64 and 128. Each router along the way decrements the TTL field by one. A router decrementing the TTL to zero will discard the packet; this prevents misaddressed packets from circulating endlessly among packet switches in search of their nonexistent destinations.

Top-Level Domain: The second level of a DNS hierarchy, which categorizes the domain by organization type (e.g., .com, .net, .edu, .biz, .info) or by country (e.g., .uk, .ca, .ie, .au, .jp, .ch).

Tor: A peer-to-peer application designed to keep the sender's IP address anonymous. This increases privacy but also conceals the identity of attackers.

Total Cost of a Countermeasure: All of the costs a firm will encounter if it installs a countermeasure, including technology costs, IT security labor costs, and increases in labor costs in non-IT business units.

Traceroute: Program that gives the round-trip latency to every router along the route to a particular destination host. Identifies links with unusually high latency.

Traffic Analysis: Analysis that asks how much traffic must flow over each of the network's many individual transmission links.

Traffic Class Field: An IPv6 field for specifying special handling for a packet.

Traffic Engineering: Designing and managing traffic on a network.

Traffic Shaping: Limiting access to a network based on type of traffic.

Transceiver: A transmitter/receiver.

Transcoding: Changing a video file into one of many formats that different viewers need to view the video.

Transmission Control Protocol (TCP): The most common TCP/IP protocol at the transport layer. Connection-oriented and reliable.

Transparent: An intermediate process whose workings are invisible to end devices.

Transport Mode: One of IPsec's two modes of operation, in which the two computers that are communicating implement IPsec. Transport mode gives strong end-to-end security between the computers, but it requires IPsec configuration and a digital certificate on all machines.

Transport Layer Process: Internet transmission standard implemented on the source and destination host. Above the internet layer and below the application layer.

Transport Process: The process (hardware or software) that implements the transport layer's functionality.

Traps: The type of message that an agent sends if it detects a condition that it thinks the manager should know about.

Trojan Horse: A program that looks like an ordinary system file but continues to exploit the user indefinitely.

True Party: In authentication, the person the supplicant says that he or she is.

Trunk Link: A type of transmission line that links switches to each other, routers to each other, or a router to a switch.

TTL: See Time to Live Field.

Tunnel Mode: One of IPsec's two modes of operation, in which the IPsec connection extends only between IPsec gateways at the two sites. Tunnel mode provides no protection within sites, but it offers transparent security.

Two-Factor Authentication: A type of authentication that requires two forms of credentials.

Two-Way Amplifier: In cable television, an amplifier that amplifies signals traveling in both directions.

UDP: See User Datagram Protocol.

UDP Checksum Field: Field in the UDP header that the receiver uses to check for errors. If the receiving transport process finds an error, it drops the UDP datagram.

UDP Length Field: Field in the UDP header that gives the length of the UDP data field in octets.

UDP Datagram: Message in the User Datagram Protocol.

UNICODE: The standard that allows characters of all languages to be represented.

Unlicensed Service Band: Unregulated radio band that does not require radio devices to be licensed.

Unreliable Protocol: Protocol that does not do error correction.

User Datagram Protocol (UDP): Unreliable transport-layer protocol in TCP/IP.

Username: An alias that signifies the account that the account holder will be using.

Verifier: The party requiring the supplicant to prove his or her identity.

Version Number Field: In IP packets, the first field; it tells whether the packet in an IPv4 packet or an IPv6 packet.

Video over IP (VoIP): The transmission of video codec data in IP packets.

Virtual LAN (VLAN): A closed collection of servers and the clients they serve. Broadcast signals go only to computers in the same VLAN.

Virtual Machine (VM): One of multiple logical machines in a real machine; to its users, it appears to be a real machine.

Virtual Private Network (VPN): A network that uses the Internet or a wireless network with added security for data transmission.

Virus: A piece of executable code that attaches itself to programs or data files. When the program is executed or the data file opened, the virus spreads to other programs or data files.

VLAN: See Virtual LAN.

VM: See Virtual Machine.

VM Instances: Specific virtual machines.

Voice over IP (VoIP): The transmission of voice signals over an IP network.

VoIP: See Voice over IP and Video over IP.

VPN: See Virtual Private Network.

Vulnerability: A security weakness found in software.

Vulnerability Testing: Testing after protections have been configured, in which a company or a consultant attacks protections in the way a determined attacker would and notes which attacks that should have been stopped actually succeeded.

WAN: See Wide Area Network.

WAN Optimization Device: Network device that optimizes wide area network traffic through compression and other methods. Desirable because WAN traffic is more expensive than LAN traffic per bit transmitted.

Wavelength: The physical distance between comparable points (e.g., from peak to peak) in successive cycles of a wave.

Weakest Link: In a series of protections that must all succeed for a countermeasure to succeed, the protection most likely to fail. If it fails, the entire series of protections is meaningless.

Well-Known Port Number: Standard port number of a major application that is usually (but not always) used. For example, the well-known TCP port number for HTTP is 80. Well-known port numbers range from 0 through 1023.

WEP: See Wired Equivalent Privacy.

Wide Area Network (WAN): A network that links different sites together.

Wi-Fi: A name created by the Wi-Fi Alliance to refer to 802.11 standards.

Wi-Fi Alliance: Trade group that created interoperability tests of 802.11 LANs; actually produced the WPA standard.

Wi-Fi Direct: A form of Wi-Fi in which wireless hosts sends frames to one another directly instead of though and access point.

Wired Equivalent Privacy (WEP): A weak security mechanism for 802.11.

Wireless LAN (WLAN): A local area network that uses radio transmission instead of cabling to connect devices.

Wireless Protected Access (WPA): The 802.11 security method created as a stopgap between WEP and 802.11i.

Wireless Protected Access 2 (WPA2): Another name for 802.11 security.

Workgroup Switch: A switch to which stations connect directly.

Working Group: A specific subgroup of the 802 Committee, in charge of developing a specific group of standards. For instance, the 802.3 Working Group creates Ethernet standards.

Worm: An attack program that propagates on its own by seeking out other computers, jumping to them, and installing itself.

Worst Case: In service-level agreements, the worst service a customer will receive without the service provider paying a penalty. The worst case for speed would be a certain *minimum* speed.

XSS: See Cross-Site Scripting.

Zero-Day Attack: Attack that takes advantage of a vulnerability for which no patch or other workaround has been released.

Zigbee: Popular Internet of Things transmission protocol.

Zigbee controller: In Zigbee, a device that controls end devices that server users, such as light, security cameras, and thermostats.

Zigbee end devices: A Zigbee device that serves users, such as light, security cameras, and thermostats.

This page intentionally left blank

INDEX

Page numbers in bold type indicate where terms are defined or characterized; page numbers in italics indicate tables or figures; page numbers with an "n" indicate a footnote.

This page intentionally left blank

CREDITS

The credits are for the icons/images used.

Chapter 1

1-2D Dima Gorohow/Shutterstock; **1-8** Kjetil Kolbjornsrud/Shutterstock; **1-11a** Anjana23121985/DigitalVision Vectors/Getty Images; **1-15a** JakeOlimb/DigitalVision Vectors/Getty images

Chapter 2

2-4a Thorbjorn66/DigitalVision Vectors/Getty images

Chapter 3

3-10 MathisworksDigitalVision Vectors/Getty Images

Chapter 5

5-10a A Sk/Sutterstock; **5-18a** RealVector/Shutterstock

Chapter 6

6-18a Vladwel/Shutterstock images

Chapter 6a

6a-1 to 6a-8 Courtesy of Xirrus Inc.

Chapter 7

7-1 Golden Sikorka/Shutterstock; **7-13** Soloma/Shutterstock; **7-17** Unggoonk/Shutterstock; **7-23** Granger Wootz/Blend Images/Getty Images; **7-24** Albert Lozano/Shutterstock; **7-25** Andrey Popov/Shuitterstock

Chapter 8

8-3 Lineicons freebird/Shutterstock; **8-16** Adapted from: Statistics of IPv6 Adoption by Google. Retrieved from: https://www.google.com/intl/en/ipv6/statistics.html

Chapter 8a

8a-2 a-1—8a-4 Reprinted with permission from Wireshark Foundation

Appendix

A-27 Kjetil Kolbjornsrud/Shutterstock; **A-28** Magnetic Mcc/Shutterstock; **A-29** Georgios Alexandris/Shutterstock

This page intentionally left blank

OTHER MIS TITLES OF INTEREST

Introductory MIS

Experiencing MIS, 8/e
Kroenke & Boyle ©2019

Using MIS, 10/e
Kroenke & Boyle ©2018

Management Information Systems, 15/e
Laudon & Laudon ©2018

Essentials of MIS, 13/e
Laudon & Laudon ©2019

Processes, Systems, and Information: An Introduction to MIS, 3/e
McKinney & Kroenke ©2019

Information Systems Today, 8/e
Valacich & Schneider ©2018

Introduction to Information Systems, 3/e
Wallace ©2018

Database

Hands-on Database, 2/e
Conger ©2014

Modern Database Management, 13/e
Hoffer, Ramesh & Topi ©2019

Database Concepts, 8/e
Kroenke, Auer, Vandenburg, Yoder ©2018

Database Processing, 15/e
Kroenke & Auer ©2019

Systems Analysis and Design

Modern Systems Analysis and Design, 8/e
Hoffer, George & Valacich ©2017

Systems Analysis and Design, 10/e
Kendall & Kendall ©2019

Decision Support Systems

Business Intelligence, Analytics, and Data Science, 4/e
Sharda, Delen & Turban ©2018

Business Intelligence and Analytics: Systems for Decision Support, 10/e
Sharda, Delen & Turban ©2014

Data Communications & Networking

Applied Networking Labs, 2/e
Boyle ©2014

Digital Business Networks
Dooley ©2014

Business Data Networks and Security, 11/e
Panko & Panko ©2019

Electronic Commerce

E-commerce 2018: Business. Technology. Society, 14/e
Laudon & Traver ©2019

Enterprise Resource Planning

Enterprise Systems for Management, 2/e
Motiwalla & Thompson ©2012

Project Management

Project Management: Process, Technology and Practice
Vaidyanathan ©2013

Pearson

www.pearson.com

ISBN-13: 978-0-13-481712-5
ISBN-10: 0-13-481712-5

EAN

9 780134 817125

90000>